



Modernising neutron spectrum unfolding for fusion applications

WONG, Ocean

Available from the Sheffield Hallam University Research Archive (SHURA) at:

<https://shura.shu.ac.uk/36014/>

A Sheffield Hallam University thesis

This thesis is protected by copyright which belongs to the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Please visit <https://shura.shu.ac.uk/36014/> and <http://shura.shu.ac.uk/information.html> for further details about copyright and re-use permissions.

Modernising neutron spectrum unfolding for fusion applications

Ocean Wong

A thesis submitted in partial fulfilment of the requirements of
Sheffield Hallam University
for the degree of Doctor of Philosophy

May 2024

Abstract

Neutron spectrum unfolding with activation foils is a technique uniquely suited to measuring the neutron spectra of fusion reactors due to its passive nature. However, more research is needed to consolidate its use as a potential diagnostic tool for future fusion power plants, as there is much room for optimization for its use in nuclear fusion applications. This PhD scrutinised this technique thoroughly and improve upon it where possible.

Separating the project into two parts, this thesis first studied the mathematics of unfolding problems, exploring how its characteristics changes with the degree of determination. Gaps in the rigour and suitability of the existing algorithms were identified, and to bridge these gaps, new unfolding algorithms were created, and they were shown to perform well on some example synthetic and real-world data. These algorithms leverage modern advances in linear algebra packages and improvement in gradient optimization algorithms, producing versatile algorithms that are quick to converge and allows for uncertainties to be analytically propagated. These new algorithms, along with some commonly used existing unfolding algorithms, were implemented into a Python module known as the *unfoldingsuite*.

Then, the process of designing activation foil neutron spectrum unfolding experiments was examined. A deficiency was identified, namely, there is currently no standardised method of selecting foils. Therefore, a procedure of choosing activation foils was created and formalised into the foil selector framework, in which the **accuracy** and **precision** of the expected unfolded neutron spectrum were simultaneously maximised, using a discrete multi-objective optimization algorithm.

This framework was then implemented as a Python program known as the foil selector program, which was used to select foils for an experiment at a high-energy neutron beamline. The results of this experiment was presented to examine the strengths and shortcomings of the foil selector program, and of activation foil neutron spectrum unfolding in general. Directions of future development of the foil selector framework and program were identified as a result.

Keywords: foil selection, discrete multi-objective optimization, unfolding algorithms, fusion neutronics, inverse problems.

I hereby declare that:

1. I have not been enrolled for another award of the University, or other academic or professional organisation, whilst undertaking my research degree.
2. None of the material contained in the thesis has been used in any other submission for an academic award.
3. I am aware of and understand the University's policy on plagiarism and certify that this thesis is my own work. The use of all published or other sources of material consulted have been properly and fully acknowledged.
4. The work undertaken towards the thesis has been conducted in accordance with the SHU Principles of Integrity in Research and the SHU Research Ethics Policy.
5. The word count of the thesis (including appendices) is 80689.

Name: Ocean Wong

Date: Wednesday 2nd July, 2025

Award: PhD in Physics

Research Institute: Materials and Engineering Research Institute

Director of Studies: R. Smith

Supervisors: C.R. Shand, A.M. Bruce, A. Alderson

Acknowledgements

I have to thank my supervisory team: Robin, Chantal and Alison, for their encouragement and helpful feedback, as well as their incredible patience, providing guidance to me through my PhD, and creating opportunities for me to broadcast my work when possible. It has been a helpful step forward in my career, and I will make sure to pay the favour forward.

I also have to thank the friendly faces at UKAEA, who made me feel supported and accommodated me through the PhD, and provided me with collaboration opportunities where applicable. It has made completing the PhD very fulfilling.

Special mention to Dr. Marcel Reginatto, who was kind enough to discuss with me over multiple video calls about how he wrote MAXED and how it was used despite his busy schedule; as well as Dr. Mitja Marjerle, whose brief mentorship I greatly enjoyed.

And lastly I would also like to thank my friends, too many to be named, who formed my steadfast support network, strong enough to bring me through even the darkest moments of these few years. In particular, I extend my appreciation to my personal friend Kristen for the many nights of co-working to motivate me to continue writing the thesis, without whom this thesis would not have been completed at all.

Table of contents

1	Introduction	1
2	Background	8
2.1	Review of existing neutron detection techniques	8
2.2	History of activation foil unfolding	12
2.3	The current state of activation foil unfolding in fusion	12
2.4	The role of activation foil on fusion power plants	14
3	Mathematics of unfolding	16
3.1	Formulation	17
3.1.1	Group structure selection	17
3.1.2	The unfolding problem formulated in linear algebra	18
3.1.3	Difficulty of the unfolding problem	21
3.2	Uncertainty consideration	22
3.2.1	Covariance matrix	22
3.2.1.1	The consequence of ignoring the off-diagonal terms of the covariance matrix	24
3.2.2	χ^2 calculation	25
3.3	Effect of degree of determination	28
3.3.1	Re-examining the difficulty of the unfolding problem using the concept of manifolds	30
4	Unfolding algorithms review and improvement	33
4.1	Classification of unfolding algorithms	33
4.1.1	Machine learning	34
4.1.2	Parametric unfolding	36
4.1.3	Regularising	38

4.1.3.1	Linear least-squares	41
4.1.3.2	Roughness of the spectrum	41
4.1.3.3	Information Entropy	42
4.1.4	Purely iterative	45
4.2	Assessing the correctness and suitability of unfolding algorithms	49
4.2.1	The model of ϕ_{true} , ϕ_{sol} , and ϕ_0	51
4.2.2	Choosing the metric	53
4.3	New algorithms developed	54
4.3.1	IMAXED	55
4.3.2	AMAXED	56
4.3.3	AMAXED-Regularization	57
4.3.4	Pseudo-Inverse	58
4.4	Demonstration with synthetic data	62
4.4.1	2-bin example	62
4.4.2	5-bin example	67
4.4.3	172-bin example	69
4.4.3.1	Theoretical resolution	73
4.5	Demonstration with experimental data	74
5	Activation foil unfolding experiment design optimization	78
5.1	Formulation and assumptions	79
5.1.1	Measuring the response vector $\mathbf{N}$ for activation foils	81
5.1.2	Reactant quantity time-invariance	87
5.1.3	Neutron spectrum time-invariance	88
5.1.4	Neutron spectrum position-invariance	89
5.1.5	Limited self-shielding from neutrons	89
5.1.6	Complete retention of radionuclides produced	92
5.1.7	Accurate <i>A priori</i>	94
5.2	Experiment design parameters	94
5.3	Limitations on the feasible range of the experiment design parameters	98
5.4	Parameters fixed by the user	101
5.5	Parameters optimised by the framework	108
5.5.1	Optimising for accuracy and precision	109
5.6	Implementation as the foil selector program	114

5.6.1	Formatting <i>a priori</i> neutron spectrum and group structure	115
5.6.2	Extracting from nuclear data libraries	116
5.6.3	Calculating number of counts	120
5.6.4	Finding the optimal foil set	122
5.6.5	Prediction of the gamma-ray spectrum (optional)	124
5.7	Future development plan for the foil selector program	125
6	Experimental validation	131
6.1	Experiment design	131
6.2	Data analysis	139
6.2.1	Extracting the number of radionuclides decayed using Genie 2000	141
6.2.2	Results	144
6.3	Review of the experiment: Lessons learned	146
6.3.1	How the foil selector's organization impacted the experiment analysis . .	149
6.3.2	Reducing the difficulty of experiment planning	151
7	Conclusion	153
Appendix A	The minimum value of χ^2	158
A.1	Clearing up misconceptions around the minimum χ^2 in an overdetermined v.s. underdetermined problem.	158
A.2	Fixing the solution ϕ_{sol} to a user-specified value of χ^2	160
A.3	A better approach: Tikhonov Regularization	161
Appendix B	IMAXED	162
B.1	Algorithm derivation	162
B.1.1	Newton's method	162
B.1.2	Line search (Wolfe condition)	163
B.2	Guaranteed convergence to a maximum	164
B.3	Uncertainty propagation	164
Appendix C	AMAXED	167
C.1	Algorithm derivation	167
C.1.1	Line search (Wolfe condition)	169
C.2	Guaranteed convergence to a minimum	169
C.3	Uncertainty Propagation	169

Appendix D	AMAXED-Regularization	171
D.1	Algorithm derivation	171
D.1.1	Line search (Wolfe condition)	172
D.2	Uniqueness of minimum	172
D.3	Uncertainty Propagation	173
Appendix E	SAND-II weight matrix	174
Appendix F	Loss function of GRAVEL	176
Appendix G	Decomposing branching decay chains	177
Appendix H	Overlapping decay chains	180
H.1	Future direction of investigation	183
Appendix I	Calculating $\Sigma_{k,\max}$	184
I.1	Future direction of investigation	185
Appendix J	Supplementary plots	186
Appendix K	Notations	196
K.1	Number sets	197
K.2	Generic information about the activation foil	197
K.3	un-discretized fluxes	198
K.4	Information Theory	198
K.5	Linear Algebra operators	199
K.6	Transmutation equations	200
K.7	Gamma-ray spectroscopy	200
Appendix L	Glossary	201
L.1	Nuclear physics	201
L.1.1	Nuclear data	202
L.2	Fusion-specific terminology	202
L.3	Experimental physics terminology	203
L.4	Mathematical terminology	203
L.4.1	Mathematical programming terminology	204
L.5	Unfolding algorithms	204

Table of contents

L.6 Metrics 205

L.7 Computational terminology 205

L.8 Terminology typically used in unfolding experiments 205

L.9 Terminology specific to this thesis 206

References **207**

Chapter 1

Introduction

Due to the ever-increasing demand for energy as well as the necessity to lower carbon emissions [1], there is a growing interest in developing nuclear fusion as a source of baseload power for electrical grids around the world. Compared to other low-carbon energy sources, it has numerous advantages. Most notably it has very high energy density [2]. The supply of fuel for nuclear fusion is also virtually inexhaustible as the isotopes required (deuterium, lithium-6 and lithium-7) are very abundant, both on a cosmic scale and on the terrestrial scale ([3] gives an estimate of enough lithium in landmasses to power the world for thousands of years to come, and enough lithium in the ocean for millions more). Contrasted with nuclear fission, which is the current type of nuclear energy in commercial use, nuclear fusion also has the advantage of being inherently stable due to its self-limiting nature [4], as well as the small volume of radioactive waste it produces compared to nuclear fission [5].

On the other hand, nuclear fusion has the drawback of having a low Technology Readiness Level (TRL) compared to other low-carbon energy sources [6]. Much research is needed to turn nuclear fusion into a commercially viable source of electricity. To this end, there are many ongoing experiments researching various ways to harness the power produced by nuclear fusion. One such approach is to use a device called a “tokamak”, where the fuel is heated into a plasma and confined to a toroidal cavity using a magnetic field. Some example tokamaks include the world’s largest tokamak, the Joint European Torus (JET), operated by the UK Atomic Energy Authority (UKAEA on behalf of EUROfusion), while an even larger tokamak known as ITER is being built through international collaborative effort for further fusion plasma experiments [7], in preparation for building another next-generation power-generating tokamak known as EU DEMO [8]. A variation of the conventional tokamak is the “spherical tokamak”, which is essentially a tokamak with a high aspect ratio in its vertical cross section. An example






Figure 1.1: The reactivity of various reactions listed in Equation 1.1 to 1.4 as given by [10].

of a currently operational spherical tokamak is the MAST-U, operated by the UKAEA on the same site as JET. STEP, the first fusion power plant planned by the UK government to supply power to the electrical grid of the UK, is also a spherical tokamak, whose concept design will be finalised in 2024.

Nuclear fusion releases energy by fusing two light nuclei to release the binding energy of the product nucleus. For most power plant concept designs, a mixture of deuterium and tritium is used as the fuel mixture in the plasma, where several reactions may take place. The four most prevalent reactions are listed below [9]:



Of the four reactions, Equation 1.3 is the dominant fusion reaction that occurs in a plasma of deuterium-tritium mixture due to its high reaction rate at the operational temperature range of tokamaks, as shown in Figure 1.1¹. The α particle and neutron share the full energy released by the reaction in the form of kinetic energy in a ratio inversely proportional to their masses, i.e. in the centre-of-mass reference frame, the α particle carries away a fifth of the energy while the neutron carries away the remaining four-fifths of the energy released per DT reaction. This

¹The reactivity parameter $\langle \sigma v \rangle$ reflects the likelihood of a reaction occurring and is obtained by integrating the (velocity dependent) reaction cross section σ with the velocity distributions of the two species of reactants. Taking the operating temperature of ITER for example, which is $150,000,000 \text{ K} = 12.9 \text{ keV}$, for the same volume and density of reactants, reaction Equation 1.3 would occur at a rate of at least 3 orders of magnitude higher than all other reactions.

gives the neutron about 14 MeV of kinetic energy and the alpha particle about 3.5 MeV of kinetic energy.

The requirement of tritium sets the logistics of nuclear fusion apart from most other energy sources, as it has to create its own fuel. Tritium is a radioactive isotope of hydrogen with a half-life of 12.3 years, thus there is no naturally occurring tritium readily available for use. The current supply of tritium (mainly produced by CANDU reactors) is insufficient for sustainable fusion power generation applications in the future. Thus tritium needs to be produced by neutron-induced reactions from the tokamak itself to make fusion a renewable energy [11]. This is done by embedding tritium-breeding material such as lithium into the first wall (i.e. interior-most wall) of the tokamak, where neutrons carrying the energy of the reaction impinge upon these materials and create tritium, which is then collected and then re-injected into the tokamak as fuel. The Tritium Breeding Ratio ($TBR = \frac{\text{tritium produced}}{\text{tritium consumed}}$) of the fusion reactor must be able to exceed unity for the tritium inventory to remain sustainable.

In magnetic confinement fusion devices such as the tokamak, the ionized fuel and products (e.g. deuterium, tritium, and the α particles in Equation 1.3) are confined by the magnetic field. Contrastingly, the neutrons escape the magnetic field and penetrate various components of the reactor, depositing their kinetic energy into the crystal structures of the components directly through neutron scattering and absorption. Scattering changes the direction of travel and kinetic energy of the neutron, while absorption binds the neutron to a nucleus, thus removing the free neutron. Both processes change the physical properties of materials that make up the fusion reactor, and this is referred to as neutron damage.

The subject of study in this thesis, the neutron spectrum, refers to the distribution of kinetic energies of all free neutrons passing through a specified region in space. The rates of neutron scattering and absorption are dependent not only on the medium that the neutrons are travelling through but also on the neutron spectrum. In turn, the neutron spectrum is modified by both processes as well, forming an intricately interdependent relationship between the neutron spectrum and neutron reaction rate in the medium. Therefore, the knowledge of the neutron spectrum, when combined with research on how neutrons of various energies modify the characteristics of materials used in the fusion reactor, allows us to optimally design and operate future fusion power plants. This is explained in further detail in Figure 1.2.

Our understanding of how neutron damage emerges in a tokamak is informed by a series of models: how the plasma generates neutrons [14], how these neutrons get transported around the tokamak [15], how materials in the tokamak are transmuted as a result [16], and finally how these product nuclides cause dimensional changes and degradation in electrical and mechanical




Figure 1.2: Different ways in which neutron spectrum information will benefit the design and operation of future fusion reactors.

properties of components (e.g. embrittlement of support structures to the magnet coils used to stabilise the plasma [17], reduction of critical field of superconductors [18], and swelling of the joining interface of the tungsten armour used to protect the reactor against the hot plasma [19]). All of these models must be validated against experimental data to extrapolate this understanding of how neutron damage emerges to larger tokamaks, which would then inform better design decisions for future fusion power plants. For example, some of these design decisions would involve using low-activation structural materials and thicker neutron shields, at the expense of reducing the margin of safety in fracture toughness and increasing the construction cost. Knowledge of the neutron spectra of existing tokamaks is therefore an indispensable piece of the puzzle, forming part of the data required to validate these models and hence allow better design decisions to be made (Figure 1.2a).

Similarly, once the fusion power plants are constructed and commissioned, it is desirable to operate them most profitably and safely. This means scheduling its maintenance in such a way that simultaneously minimises loss of revenue due to downtime and the likelihood of failure of components, maximising its operational lifetime, and reducing the cost of end-of-life decommissioning. These would require careful adjustments of the plasma parameters, which would then affect the power output, rate of neutron damage, activation of reactor materials, and

TBR. The calculation of these quantities would benefit from knowledge of the neutron spectrum; therefore if these quantities need to be continuously fed back to the operators, a frequently updated measurement of the neutron spectrum is required (Figure 1.2b). The operator can then take actions such as modifying the plasma parameters or replacing the tritium breeding blanket materials accordingly.

However, despite the emphasis on the importance of knowledge of the neutron spectrum, several obstacles stand in the way of obtaining neutron spectrum measurements. The neutron flux, temperature, and magnetic field strength in future tokamaks are expected to be so high as to render all active neutron spectrum measurement methods unusable. Take ITER as an example. It is expected to have a neutron flux of $> 10^{14}\text{cm}^{-2}\text{s}^{-1}$ [20], comparable to the neutron flux of $\approx 10^{15}\text{cm}^{-2}\text{s}^{-1}$ found in a fast reactor (Table 1 from reference [21]). The temperature of the first wall is expected to reach up to $800\text{ }^{\circ}\text{C}$ (at 4.7MWm^{-1} [22]), while the magnetic field strength is expected to exceed 12.5 T [23]. Any one of these three conditions is enough to disable even the most advanced active neutron detection method currently available. Therefore, despite its inferior energy resolution and time resolution, the passive measurement method of foil activation neutron spectrum unfolding is the prime choice of technology to measure the neutron spectrum at the first wall.

While the use of an activation foil to measure neutron rate (a scalar quantity) is straightforward and often done [24–29], using it to measure the neutron spectrum (a vector quantity) remains a difficult challenge. Conceptually, the challenges of activation foil neutron spectrum unfolding can be broken down into three stages as shown in Figure 1.3:

1. Neutron activation, which includes:
 - (a) foil selection, and
 - (b) choice of irradiation schedule.
2. Gamma-ray spectroscopy to measure the activities of radionuclides present in the foil, and
3. unfolding to reconstruct the original neutron energy spectrum.

The field of gamma-ray spectroscopy is well-developed, with extensive literature available. Therefore challenge 2 will not be a key topic of discussion in this thesis. A useful mental model is to split the remaining two challenges into “hardware” and “software” challenges: Neutron activation (challenge 1) concerns the hardware (physical set-up) that is placed in the neutron field, and unfolding (challenge 3) concerns the software procedure (unfolding algorithms) to use to reconstruct the neutron spectrum.




Figure 1.3: Flow diagram depicting the neutrons from the tokamak (top) activate an array of foils (middle), which are then extracted to have their radioactivities measured by gamma-ray detectors (bottom). Unfolding algorithms then reconstruct the neutron spectrum using the gamma-ray spectrum. Images taken from [30] and [31].

This thesis is arranged in such a way that the reader can familiarise themselves with activation foils neutron spectrum unfolding before learning about its challenges, pitfalls, and proposed solutions to each challenge. The proposed solutions are put to the test in an experiment, and then the results are discussed.

The current chapter (Chapter 1) gives an introduction to nuclear fusion and the role that neutron spectrum measurement plays in helping to commercialise it as an energy source for the electrical grid.

Chapter 2 lays out the background context of why neutron spectrum unfolding with activation foils is favoured in fusion energy generation applications, and what alternatives are available. In addition, a review of the history of this technique and its current usage is included to provide a comprehensive overview.

Chapters 3 and 4 deal with the “software” issue (challenge 3) by explaining the mathematical problem of unfolding, and then showing how it is solved in practice using various algorithms. Several improved algorithms are proposed, and implemented into part of a Python software package along with several other existing algorithms.

Chapter 5 deals with the “hardware” issue (challenge 1), identifying foil selection (challenge 1a) as the more tractable problem to tackle, then proposing a foil selection framework that maximises the accuracy and precision of the unfolded solution. This framework is implemented as a foil selector program written in Python as well.

Chapter 6 demonstrates both the use of this proposed foil selection framework and the improved unfolding algorithms in an experiment that characterises a neutron beam.

Chapter 7 concludes the findings of this PhD and its contribution to knowledge, by summarising what obstacles have been solved, and what obstacles remain ahead, and then proposes future research directions to tackle these obstacles.

Readers may find it useful to have a printed copy of Appendix K and L on hand when reading the main body of the thesis, as the number of technical terms and notations grows with the length of the thesis, and these appendices are useful references to help keep track of them.

Chapter 2

Background

To fully appreciate how improvements proposed in this thesis impact activation foil experiment design procedures, and wider applications of unfolding in general, we must first familiarise ourselves with the technology of activation foil neutron spectrum unfolding. This chapter explores how activation foils neutron spectrum unfolding came to be the most desirable type of neutron spectrum measurement technique to be deployed in future fusion reactors, and its current shortcomings, which will be addressed in Chapters 3 and 5.

2.1 Review of existing neutron detection techniques

To help contextualise the modern neutron spectrometry technology, an inexhaustive list of currently available neutron spectrometers is compared and contrasted below.

One of the neutron spectrometers with the simplest design is the Bonner sphere system (BSS). A bare ^3He proportional counter can count the number of incident thermal neutrons efficiently but has reduced efficiency at counting higher-energy neutrons. By using plastic spheres of increasing radii to act as moderating sheaths around the thermal neutron counter, the efficiency at detecting higher energy neutrons can be increased as neutrons are slowed down sufficiently to register in the detector (See the response matrix in Fig. 4 of [32]). The neutron spectrum can then be reconstructed by an unfolding algorithm by comparing the count rates of the detector when moderated by differently-sized spheres. The shortcoming of the BSS is its limited resolution: the amount of information provided to the unfolding algorithm is limited by the number of spheres used, which is typically fewer than 10 in practice. This in turn limits the accuracy of the unfolded spectrum [33]. The presence of active electronics renders it difficult to deploy at high temperatures and strong magnetic field conditions like those found at the first

2.1 Review of existing neutron detection techniques

wall. The moderator spheres also take up significant space, especially at fusion-relevant energies, making them difficult to deploy at the first wall.

In contrast to the BSS, scintillator and semiconductor detectors produce pulses whose pulse heights are proportional to the energy deposited by the neutron scattering event inside the detector. The detection medium of scintillator detectors is typically composed of organic compounds such as NE213 or stilbene, which releases photons after excitation by neutron collisions. A photomultiplier tube (PMT) collects the scintillation photons to create a signal proportional to the brightness of that scintillation event, which in turn is proportional to the energy deposited by the neutron collision [34]. Meanwhile in semiconductor detectors, which include Si, SiC, and C (i.e. diamond) detectors, the energy deposited by neutron scattering (both elastically and inelastically) into the semiconductor crystal is directly converted into a number of electron-hole pairs that is proportional to the deposited energy [35–37]. As the neutron collision rarely deposits the full neutron energy, unfolding is required to reconstruct the neutron spectrum from the pulse height spectra, i.e. histograms of the brightness of each scintillation event collected by the PMT anode [38] or the signal pulse-height histogram accumulated by the semiconductor device. As with the BSS, both semiconductor and scintillator detectors expose vulnerable active electronics to the harsh conditions near the fusion plasma. It is also worth mentioning that organic scintillators are typically unsuitable for high-temperature applications such as those found near the fusion plasma, as the scintillating media are typically toxic, flammable and have lower melting and boiling points than the semiconductor detectors' media.

The same principle of measuring the scattered charged-particle energy is employed in Magnetic Proton Recoil spectrometers (MPR). When neutrons collide with a hydrogen-rich target, the protons scattered in the forward direction carry the full energy of the neutrons. Using a magnetic field to sort scattered protons, the neutron spectrum can be reconstructed accurately even without post-processing algorithms such as unfolding [39] to sharpen the peaks¹. However, MPR requires a large volume of equipment and space to generate the magnetic field in which momentum separation can occur. It must also be situated far from the first wall to shield it from stray radiation. This restricts its use for measuring the first wall spectra in future fusion reactors, which are already crowded with breeder blankets necessary to maximise the tritium breeding efficiency [40].

¹In these cases where the response functions are simple, researchers often perform forward fitting (forward operation) rather than unfolding (inverse operation), and compare the goodness of fit by comparing the calculated pulse-height spectrum with the measured pulse-height spectrum directly.

2.1 Review of existing neutron detection techniques

Table 2.1: Comparison of neutron detection techniques available for fusion neutron spectrometry. Properties preventing their use at the first wall of tokamaks with high neutron yields are coloured in red.

Spectrometer	Electronic components near plasma	Size	Radiation hardness to neutrons	Resolution	Necessity of unfolding
BSS	detector electronics	large	$< 2 \times 10^4 \text{cm}^{-2}\text{s}^{-1}$ [42] ³	“poor” (Table 1 of [33])	Required
Scintillator	detector electronics ⁴	small	10^{14}cm^{-2} for stilbene [43]	20.49% at 0.498 MeV −5.13% at 10 MeV for stilbene [44]	Required
SiC	detector electronics	small	$1.65 \times 10^{13} \text{cm}^{-2}$ [35]	9.1% at 3.73–5.16 MeV [35]	Required
Diamond	detector electronics	small	$8 \times 10^{14} \text{cm}^{-2}$ [37]	2.2% at 5.2 MeV [45]	Optional
Magnetic Proton Recoil	bending magnets	large	unlimited, controlled by collimator diameter	2.5%–4% at 14 MeV [40]	Optional
Time-of-Flight	bending magnets	large	unlimited, controlled by foil thickness and collimator	7% (dependent on foil thickness) [41]	Optional
Activation foils	Absent	small	unlimited, controlled by foil volume	poor, with room for improvement (See Section 4.4.3.1)	Required

The time-of-flight (ToF) method presents a similar portability issue. The specific design varies with each ToF detector, but as the name suggests, they all measure the time required by neutrons to traverse a fixed distance, hence inferring the distribution of neutron velocities. Again, due to the sharp response function (injective relationship between neutron energy and ToF), unfolding is rarely required to obtain the neutron spectrum [41]². However, due to the large distances required to separate slow neutrons from fast ones, especially up to 14.1 MeV neutrons, its view angle of the plasma is very limited, and its dimensions render the application of ToF infeasible in future fusion power plants.

All of the aforementioned technologies are compared in Table 2.1. The “Size” column summarises the ease of fitting the instrument into the first wall of a fusion reactor. A “small” instrument is defined as one that can fit within a typical irradiation end in a fusion reactor,

²See footnote 1

³Reference [42] shows the maximum count rate is 20000, and counting efficiency reaches > 1 count per 1cm^{-2} fluence.

⁴Other considerations include: highly flammable liquid or solid with low melting point used as detection medium, which is not suitable for the harsh conditions inside a fusion reactor

2.1 Review of existing neutron detection techniques

which is a cylinder of diameter 9 to 18 mm (Figures 2 and 4 of [46] and [47] shows the irradiation end and the capsule that can fit into it). Otherwise it is considered a “large” instrument.

The “Radiation hardness” column lists the maximum neutron fluence (cm^{-2}) or neutron flux ($\text{cm}^{-2}\text{s}^{-1}$) that a typical detector of that kind can withstand while remaining operational, bearing in mind that the neutron flux at the first wall of an operational nuclear fusion power plant is expected to have neutron flux exceeding $1 \times 10^{14} \text{ cm}^{-2}\text{s}^{-1}$ [20].

“Resolution” is typically calculated as $\frac{\text{FWHM}}{\text{Energy of peak}}$ in the neutron spectrum, but in some cases (BSS and activation foils) this number is not available due to the low fidelity at reproducing the spectrum by unfolding as they are typically underdetermined. The resolution achieved would be strongly dependent on the system and as such does not have a typical value that can be quoted.

Neutron spectrum unfolding with activation foils is the only row in Table 2.1 without any red cells, meaning that no significant obstacles are preventing their use in tokamaks when it comes to its size and tolerance for high temperature and magnetic fields. However, it has some significant drawbacks, such as its inferior resolution. In an ideal system with no noise, the theoretical resolution is equal to $\frac{1}{m}$, where m = the number of distinct radionuclides (see Section 4.4.3.1). With the help of the foil selector (Chapter 5), m may reach up to 30, which means the neutron spectrum may have a resolution of 4% at 14 MeV. However, this is purely a theoretical upper bound, and in practice, accounting for noise and underdetermination, the resolution will be much worse for reasons explained in point 4 of Section 3.3, and is likely to be worse than all other entries in Table 2.1 except for Bonner Spheres. The remaining chapters of this thesis will be spent exploring methods to expand its capability, filling this gap in the resolution, and streamlining experiment designs involving activation foils.

Apart from its hypothesized application in fusion power plants, neutron spectrum unfolding with activation foils is also widely adopted as a spectrometry technique in various neutron-sources around the world [48–52] since it requires no bespoke hardware (only requiring foils and gamma-ray detectors, which are readily available in most physics laboratories), and is gamma-ray insensitive, i.e. the presence of gamma-rays does not interfere its neutron measurement. Hence, improvements brought about in this thesis, including reduced uncertainty and improved resolution due to better foil selection and unfolding algorithms, will benefit these applications of activation foil unfolding as well as fusion applications.

2.2 History of activation foil unfolding

Activation foil unfolding emerged in the late 1960s for measuring fast neutrons from fission reactors [53, 54]. SAND-II is the only program created in this period for the express purpose of unfolding activation foil spectra that remains popular to the present day. Since then, active neutron measurement techniques such as the BSS and scintillator detectors have gained popularity much faster and dominated the field of neutron spectrometry, while activation foils only remained useful in niche applications where high radiation hardness and simplicity of the spectrometers are important [55]. The development of other active measurement techniques led to improvements in unfolding that also spilt over to activation foils neutron spectrum unfolding, specifically unfolding programs such as MAXED [56] and GRAVEL [57] that improved the software portability and mathematical rigour of the unfolding algorithms.

Some of these unfolding programs use the terminology of “deconvolution” instead of “unfolding” (such as the D in MAXED [56]). However, we should clarify that this is a misnomer. Deconvolution problems are a subset of unfolding problems which are less pathological in nature. Their response functions (a.k.a. kernel) do not vary significantly across different detection channels, similar to a Toeplitz matrix⁵. While plenty of deconvolution algorithms exist, most of them rely on the simplicity of the deconvolution problem, and hence cannot be generalised to apply to unfolding problems where this symmetry in the response matrix is not guaranteed. However the reverse is possible: adapting unfolding algorithms to apply to deconvolution problems is straightforward, as will be shown in Section 4.5.

2.3 The current state of activation foil unfolding in fusion

After the 2000s, as the design and construction of the first nuclear fusion power plant approaches, the interest in activation foil unfolding has been rekindled as it is seen as the most viable option for monitoring the first wall spectrum [58, 59].

JET is the largest fusion reactor in the world with tritium handling capabilities and the ability to create and control deuterium-tritium fusion plasmas, making it the most intense source of fusion-relevant energy neutrons [15]. When ITER is constructed and begins operation, it will overtake JET as the most intense fusion neutron source on Earth [60]. In both reactors, foils can be irradiated in an “irradiation end” (Figure 2.1), which is a blind-end tube recessed

⁵A Toeplitz matrix is a matrix where each row is a shifted copy of the previous row.

2.3 The current state of activation foil unfolding in fusion




Figure 2.1: Schematics of where irradiation ends are located in the poloidal cross sections of JET and ITER.

into the wall of the vacuum vessel of the tokamak, or inside port plugs, which are openings in the wall for diagnostic and heating equipment.

Despite the recessed position, these irradiation ends still reach very high temperatures (around 300°C in JET [63] and over 500°C in ITER expected [64]), while being subjected to nuclear heating and strong electromagnetic forces [65]. To avoid manual handling of samples in light of these harsh conditions, the activation foils are placed inside capsules which are then pneumatically transported into and out of these irradiation ends, allowing the irradiation duration to be precisely controlled by the experimentalist. Such a system is called a Neutron Activation System (NAS). NAS is also used by other groups experimenting on smaller reactors and beamlines, such as KIT (at the TUD-NG) [66], and KSTAR [67].

The choice of unfolding program used in fusion neutron spectrum unfolding varies across research groups, but the most common choice is the UMG3.3 package produced by Physikalisch-Technische Bundesanstalt (PTB) [68], which implements the most commonly used algorithms, MAXED and GRAVEL, in FORTRAN-90. The results of these FORTRAN-90 codes are dependent on the platform they are compiled and executed on, as an element of stochasticity is required in some of the subroutines. This is affected by platform-dependent parameters such as the byte length of the floating-point random number used in the simulated annealing subroutine, making their result less reproducible. Another common choice is the HEPROW package, which implements GRAVEL, SAND-II and other older algorithms, also in FORTRAN-90 [69]. Some other research groups, such as the Frascati National Laboratory, use their own parametric unfolding code FRUIT [70], and the National Physics Laboratory (NPL, United Kingdom) [71] uses STAY’S L, which is written in FORTRAN-IV [72].

While this was considered state-of-the-art computing two decades ago, this level of software portability and reproducibility is lacklustre compared to modern-day scientific computing, and has room for improvement by writing updated implementations of these algorithms in more modern languages (such as that seen in [73]). Most of these programs (except for UMG3.3 MAXED) also use Monte-Carlo methods to propagate error from the input space (uncertainties of raw measurements at the spectrometers) into the output space (neutron spectrum uncertainties), which is slow and inefficient compared to analytically derived error propagation methods or automatic differentiation (used in neural networks). To tackle these issues, an upgrade to the UMG3.3 software package is attempted in Section 4.3. Other more intricate shortcomings of some of these algorithms (specifically, GRAVEL and SAND-II) include their underutilisation of the user-provided (*a priori*) information. This is explained in Section 4.1.4 in greater detail.

Another issue with the current approach of activation foil is the method of foil selection. The current approach is done *ad hoc*, where foils with threshold reactions familiar to the experiment designer are chosen, and then irradiated. This limits the number of types of materials available, and consequently limits the amount of information that can be extracted from the experiment, limiting the quality of the unfolded spectrum as will be explained in Section 3.1.3. To remove this reliance on anecdotal knowledge of individual experiment designers, a program was created in Chapter 5 to search for all possible foils more systematically.

2.4 The role of activation foil on fusion power plants

Chapter 1 (Figure 1.2) explained the need for neutron spectrum information in fusion devices and Section 2.1 (Table 2.1) showed that despite its poor resolution, no other spectrometers than activation foils are capable of measuring the first wall spectrum. Why, then, can we not simply pivot away from measuring the first wall neutron spectrum, and measure the neutron spectrum at a location far away from the first wall using other higher-resolution spectrometers?

If active neutron spectrometers were to be used, we must move the measurement location from the first wall to several meters outside the vacuum vessel to avoid the challenging conditions at the first wall of power-generating tokamaks (and likely in other fusion devices as well). Regardless of what type of fusion device is used, this solution requires a long collimated channel with a direct line of sight to the fusion plasma to extend out of the reactor (similar to the JET Roof Lab [74]), massively increasing the dose rates outside of the vacuum vessel, putting operators and sensitive equipment at risk and potentially causing further challenges for decommissioning. Adding shielding to compensate for this would only add further cost and engineering complexity

to the design plans for the fusion power plant, and take up precious space that can otherwise be used for breeding tritium. The narrow line of sight into the plasma means that only a small portion of the plasma neutrons and neutrons backscattered from the first wall would be sampled, thus the quality of measurements would be at the mercy of local perturbations due to impurities in the plasma or transient events. Also as neutrons travel through this collimator tube, it will downscatter on the side walls of this tube, further degrading to the measured neutron spectrum.

Alternatively, if the spectrometer used does not have a direct line of sight access to the fusion plasma, and instead only measures the energy spectrum of fusion neutrons after they have passed through many layers of structural and shielding material to reach the instrument, then a different challenge presents itself: the spectrum obtained would be significantly different to the spectrum at the first wall due to neutron scattering and absorption, so much work would be required to reconstruct the first wall neutron spectrum from this measurement. It would also be difficult to calibrate such an instrument in the absence of a fusion plasma producing DT neutrons at operational power levels [24].

In contrast, if activation foils are used, they can be pneumatically pushed to near the first wall using gas injection manifolds⁶ or cooling channels with only minimal design changes, so no appreciable engineering complexity or cost would be added. Being placed at the first wall, the activation foils would be able to integrate the neutron spectrum from a large viewing angle, making them less susceptible to the aforementioned local perturbations in the plasma. Thus, in practice, using activation foils to measure the fusion neutron spectrum is favoured over other choices of spectrometers.

⁶The term “manifold” used in this chapter refers to piping for transporting gas, which unfortunately collides with the mathematical term “manifold” used in Chapter 3. Compare the glossary in Appendix L.2 and Appendix L.4 for disambiguation.

Chapter 3

Mathematics of unfolding

Before we can review the various types of algorithms used for unfolding, we must lay the foundation of the unfolding problems in mathematical terms. The mathematics of unfolding is often presented as a straightforward minimization problem with no detailed explanation on why it is formulated as such, and more importantly, these descriptions make no differentiation between underdetermined unfolding and overdetermined unfolding, two fundamentally different problems. This chapter aims to bridge these gaps in the literature, as well as to make this thesis more coherent.

As was made clear in Section 2.1, most neutron detection methods require an unfolding algorithm to convert from the physically measurable quantities (such as Bonner sphere count rates or gamma-ray counts from activation foils) into the neutron spectrum, therefore this and the next chapter (Chapter 4) are written in a way that is agnostic to the underlying detection technology used. All statements made here can be generalised to other forms of unfolding technology, and in particular, new algorithms inspired by this chapter and proposed at the end of Chapter 4 can equally be applied to other neutron spectrometers by substituting in the appropriate response matrix for that spectrometer.

The bridging of the aforementioned gaps in the literature, along with the systematic review of existing and invention of new algorithms in the next chapter, forms the first notable contribution to knowledge in this PhD thesis.

3.1 Formulation

The energy of a neutron spectrum is a continuous variable¹, which means the neutron energy can take an infinite number of possible values between $E \in (0, E_{\max})$.

3.1.1 Group structure selection

To reduce the complexity of the problem and allow linear algebra techniques and software to be applied, the neutron spectrum is often discretized into a number of intervals, called a group structure, given by a list of $n + 1$ ascending neutron energies, $(E_0, E_1, \dots, E_n)$, such that:

$$\phi_i = \int_{E_{i-1}}^{E_i} \phi(E) dE, \quad (3.1)$$

where ϕ_i is the i^{th} group neutron flux, which contains the integral of the scalar flux ($\int \phi(E) dE$) between the lower bound energy (E_{i-1}) and the upper bound energy (E_i) of that i^{th} group. For the rest of this thesis, n will be used to denote the total number of energy groups, and i will be used as the dummy variable denoting a specific energy group, $1 \leq i \leq n, i \in \mathbb{N}$ where $\mathbb{N}$ denotes the set of all natural numbers (see Appendix K.1). ϕ would then be used to denote the complete neutron energy spectrum from $i = 1$ to n .

An example group structure typically used in fusion is the VITAMIN-J [76, 77], which has 175 groups binned in roughly even bin sizes on a logarithmic scale, with $E_0 = 1.00 \mu\text{eV}$ and $E_{175} = 19.64 \text{ MeV}$. The reason for using logarithmic bin sizes is that neutrons lose energy in logarithmic steps, and the rate of variation in some nuclear reaction's cross sections also increases as the neutrons thermalise (i.e. reduce in speed), causing more rapid changes to the neutron flux at the low energy end of the neutron spectrum.

Upon discretization into a group structure, the neutron flux per unit energy (unit: $\text{eV}^{-1} \text{cm}^{-2} \text{s}^{-1}$) is assumed to be constant within each bin, i.e. it is assumed to be a piecewise constant function, causing the step-like appearance in all neutron spectra plots in this thesis (e.g. Figure 3.5). Any intra-bin variation will be ignored, so an improper selection of group structure may lead to undetected errors in the calculation in Section 3.1.2. While a finer group structure will reduce the error introduced by intra-bin variations, it makes unfolding more difficult as it increases the degree of underdetermination (see Section 3.1.3). An ideal group structure would have as few bins as possible while still capturing all of the fluctuations in $\phi(E)$ w.r.t. E .

¹Not to be confused with the neutron spectrum (scalar flux $\phi(\mathbf{r}, E; t)$) itself, which is not guaranteed to be continuous with respect to neutron energy (E) [75]. This will be discussed in more detail in Section 4.1.3.2.

A previous study suggests that there exist group structures better suited for fusion applications, as those bespoke group structures were shown to have at least two orders of magnitude less error in their neutron spectra than when VITAMIN-J was used for neutronics simulations [78]. While this shows that there is certainly room for improvement in the choice of group structure, the optimization of group structure has been excluded from the scope of this thesis, as the choice of group structure is a confounding variable that will affect the performance of neutron spectrum unfolding and increase the complexity of foil selection. In this thesis, no bespoke group structures will be created, and all will be chosen from existing, currently widely applied group structures.

3.1.2 The unfolding problem formulated in linear algebra

In the specific case of activation foil measurement, for each neutron-induced reaction with a measurable product, the number of nuclei of the product radionuclide produced is given by N_k , where k is an integer between 1 and m ($k \in \mathbb{N}, 1 \leq k \leq m$). Throughout this thesis m will be used to denote the number of types of measurable radionuclides produced via neutron-induced reactions in the foil set. The value of N_k is given by Equation 3.2:

$$N_k \propto \int_0^\infty \sigma_k(E) \phi(E) dE, \quad (3.2)$$

where $\sigma_k(E)$ is the microscopic cross section for the production of the k^{th} reaction's product nuclide at energy E , and has units of barns (10^{-24} cm^2). Discretizing σ_k into the same group structure as ϕ_i gives:

$$\sigma_{ki} = \frac{\int_{E_{i-1}}^{E_i} \sigma_k(E) \phi(E) dE}{\int_{E_{i-1}}^{E_i} \phi(E) dE}, \quad (3.3)$$

but we can simplify this by assuming $\phi(E)$ is piecewise constant over the range $E_{i-1} - E_i$ (as stated at the end of Section 3.1.1) and then cancelling out $\phi(E)$ from the numerator and

3.1 Formulation

denominator, such that σ_{ki} is the mean value of $\sigma_k(E)$ between E_{i-1} and E_i :

$$\begin{aligned}\sigma_{ki} &= \frac{\int_{E_{i-1}}^{E_i} \sigma_k(E) dE}{\int_{E_{i-1}}^{E_i} 1 \cdot dE} \\ &= \frac{\int_{E_{i-1}}^{E_i} \sigma_k(E) dE}{E_i - E_{i-1}} = \text{group averaged cross section.}\end{aligned}\tag{3.4}$$

By wrapping the proportionality constant in Equation 3.2 along with the microscopic cross section σ_{ki} into a single variable R_{ki} which has unit cm^2s (number of nuclei produced per unit flux), we obtain the following set of summation equations,

$$\begin{aligned}N_1 &= \sum_{i=1}^n R_{1i} \phi_i, \\ &\vdots \\ N_k &= \sum_{i=1}^n R_{ki} \phi_i, \\ &\vdots \\ N_m &= \sum_{i=1}^n R_{mi} \phi_i.\end{aligned}\tag{3.5}$$

The values of R_{ki} will be fully defined in Equation 5.8 in Section 5.1. For the purpose of this chapter and the next chapter, it is sufficient to know that R_{ki} is dependent on the microscopic cross section σ_{ki} .

In fact, by abstracting away the proportionality constants between N_k and ϕ_i into R_{ki} , we restore the generality to the unfolding equation, as R_{ki} can then be substituted by the response function in the k^{th} channel of the instrument, regardless of what underlying instrument was used, so Equations 3.5 can be used to describe the new instrument with equal validity. When generalised in this manner, m is then used to describe the number of measurement channels in the neutron spectrometer.

By collecting N_k into an m -dimensional vector and the neutron spectrum ϕ into an n -dimensional vector, we can rewrite the set of Equations 3.5 into Equation 3.6,

$$\mathbf{N} = \underline{\mathbf{R}}\phi.\tag{3.6}$$

3.1 Formulation

Here, $\underline{\mathbf{R}}$ is a matrix with shape m rows $\times$ n columns, and is known as the “response matrix”. The operation of left-multiplying $\underline{\mathbf{R}}$ onto ϕ to obtain $\mathbf{N}$ is called “folding” the neutron spectrum (ϕ) with the response matrix ($\underline{\mathbf{R}}$) to obtain the response ($\mathbf{N}$), hence the inverse operation of Equation 3.6 is known as “unfolding”. In other neutron spectrometers, the vector $\mathbf{N}$ may instead represent the pulse-height spectrum from a scintillator or diamond detector, a time-of-flight spectrum from a ToF spectrometer, or the list of count rates from each Bonner sphere.

The neutron spectrum that experimentalists seek, ϕ_{true} , is not directly measurable. Instead, what could be measured is the response vector $\mathbf{N}_{\text{meas}}$,

$$\mathbf{N}_{\text{meas}} = \underline{\mathbf{R}}\phi_{\text{true}} + \epsilon. \quad (3.7)$$

The extra error term ϵ is an m -dimensional vector, i.e. in the same shape as $\mathbf{N}_{\text{meas}}$, and represents the stochastic noise inherent to the measurement system [79]. Because both ϵ and ϕ_{true} are unknowns, we can only assume ϵ to be small and negligible, so that ϵ may be omitted (as it often is in the literature), such that Equation 3.7 remains solvable. The measured response $\mathbf{N}_{\text{meas}}$ is contrasted with the calculated response, $\mathbf{N}_{\text{sol}}$. This is calculated using the “solution” spectrum ϕ_{sol} outputted by unfolding algorithms. ϕ_{sol} can be folded through Equation 3.6 to give $\mathbf{N}_{\text{sol}}$, a quantity representing what the measured response would have been if the neutron spectrum was truly equal to ϕ_{sol} (and the measurement errors $\epsilon = \mathbf{0}$).

$$\mathbf{N}_{\text{sol}} = \underline{\mathbf{R}}\phi_{\text{sol}}. \quad (3.8)$$

Since many terms have been introduced, the two types of vectors, ϕ and $\mathbf{N}$, are summarised in Table 3.1.

Table 3.1: Meanings of ϕ and $\mathbf{N}$ with different subscripts

Vector space	Reality, not known to the experimentalist (signal without noise)	Measured	Hypothesized by the unfolding algorithm	<i>a priori</i> information provided by the user
ϕ	ϕ_{true}		ϕ_{sol}	ϕ_0
$\mathbf{N}$		$\mathbf{N}_{\text{meas}}$	$\mathbf{N}_{\text{sol}}$	

The final column of the table introduces another quantity — ϕ_0 , the *a priori* spectrum, which is a piece of information that unfolding programs often require users to input to aid them with the unfolding process and reducing the difficulty of the unfolding problem, which is explained in the following Section 3.1.3.

3.1.3 Difficulty of the unfolding problem

Using the quantities introduced in Table 3.1, we can begin to dissect the mathematical properties of unfolding problems that make them one of the most challenging types of inverse problems to solve. Unfolding problems are difficult due to their ill-posed and often ill-conditioned nature [80, 81]. These issues are particularly serious for underdetermined unfolding problems. These terminologies are defined below.

The term “determined/determination” in the mathematical sense denotes whether there is sufficient information present to determine the answer. Take the set of Equations 3.5 for example. There are m simultaneous equations and n unknowns (ϕ_i , $1 \leq i \leq n$). The problem is considered “fully-determined” when $m = n$ as there is exactly enough information to solve it; “overdetermined” when $m > n$ as there is redundant/repeated information, and “underdetermined” when $m < n$ as there is insufficient information. When the problem is underdetermined, which is fairly often the case for activation foil unfolding (where $m \approx 10$ while $n \approx 175$ for the VITAMIN-J group structure), Equation 3.6 has infinitely many solutions, so a unique solution will never exist. The degree of determination is defined as $m - n$; and to make discussion in underdetermined situations easier, we will also introduce the term “degree of underdetermination” $= n - m$, which is essentially $(-1) \times (\text{degree of determination})$.

The term “ill-conditioned” means errors in the input space $\mathbf{N}$ (i.e. ϵ) are easily magnified and projected into the output space ϕ , and can be numerically quantified as having a response matrix $\underline{\mathbf{R}}$ with a large condition number C defined by $C = \frac{\text{largest singular value of } \underline{\mathbf{R}}}{\text{smallest singular value of } \underline{\mathbf{R}}}$. The smallest condition number possible for any inverse problem is 1, which indicates a well-conditioned problem. Typically the condition number of the response matrix of an underdetermined activation foil unfolding problem is larger than the inverse of the precision of a double-precision floating-point number, i.e. $C > \frac{1}{\text{float precision}} = 10^{15}$, which is undoubtedly a large condition number and hence meets the definition of an ill-conditioned problem.

Meanwhile, “ill-posed” problems are more loosely defined than ill-conditioned problems. A well-posed problem must be well-conditioned and have a unique solution. Otherwise, if either of these criteria are violated, it is considered an ill-posed problem. Hence the set of ill-posed problems is a superset of ill-conditioned problems (ill-posed problems $\supset$ ill-conditioned problems). As seen in Equation 3.7, when ϵ is omitted, there may be an unreconcilable difference between the L.H.S. and R.H.S., making it unsolvable and ill-posed. Another reason for its ill-posed nature is obvious in cases of underdetermined unfolding problems, where there is no unique solution to the set of Equations 3.5.

3.2 Uncertainty consideration

Since finding an exact solution to Equation 3.6 is an ill-posed problem, all unfolding algorithms instead attempt to solve the easier, related, optimization problem of

$$\underset{\phi_{\text{sol}}}{\operatorname{argmin}}(\text{distance between } (\mathbf{N}_{\text{meas}}, \mathbf{N}_{\text{sol}})). \quad (3.9)$$

or an extension of this equation with more terms inside argmin . While our true intention is to minimise the distance between the solution spectrum ϕ_{sol} and the true spectrum ϕ_{true} in ϕ space, i.e. $\underset{\phi_{\text{sol}}}{\operatorname{argmin}}(\text{distance between } (\phi_{\text{true}}, \phi_{\text{sol}}))$, we cannot do so because ϕ_{true} remains unknown to us. Instead, using the known quantity $\mathbf{N}_{\text{meas}}$, the best approximation achievable is by minimising the proxy quantity in the parentheses in Equation 3.9 in $\mathbf{N}$ space. The definition of “distance” may differ across algorithms. The inclusion of *a priori* information in the various manners discussed in Section 4.1 helps to provide extra information for the unfolding algorithms. In fully- and overdetermined cases this helps improve the quality of the solution, while in underdetermined cases, the inclusion of *a priori* information is critical as all unfolding algorithms cannot run on underdetermined unfolding problems without it.

3.2 Uncertainty consideration

The last piece of information that must be provided to the unfolding algorithms is the uncertainty, so that the algorithm can weight the inputted quantities according to their reliability.

3.2.1 Covariance matrix

The form of uncertainty most familiar to experimental physicists is the standard deviation (σ_x), which is used to quantify “how wide is the normal distribution of the possible values of x ”, assuming the probability distribution of the scalar variable x can be modelled as a normal distribution to begin with. A closely related concept, variance $\operatorname{var}(x) = \sigma_x^2$, is simply the square of the standard deviation. This sub-section gives a quick introduction to generalise this concept of variance from scalar variables to vector variables so that we can quantify the uncertainties on $\mathbf{N}_{\text{meas}}$ and ϕ_{sol} in later chapters. This generalization is called the “covariance matrix”. The definition of the i^{th} row j^{th} column element in the covariance matrix can be given by the following formula:

$$\operatorname{cov}(x_i, x_j) = \lim_{N \rightarrow \infty} \sum_{k=1}^N \frac{((x_{i,k}) - \overline{x_i})(x_{j,k}) - \overline{x_j})}{N} \quad (3.10)$$

3.2 Uncertainty consideration

where k = the number of samples

$$= \mathbb{E}[(x_i - \mathbb{E}(x_i))(x_j - \mathbb{E}(x_j))], \quad (3.11)$$

where $(x_{i,k}, x_{j,k})$ is the k^{th} pair of sample of (x_i, x_j) , N is the total number of samples, $\mathbb{E}(\dots)$ denotes the expected value (i.e. mean value) when the quantity inside the parentheses is sampled many times. It is trivial to substitute in $j = i$ to prove that the main diagonal elements $= \text{var}(x_i)$. Therefore $\text{cov}(x_i, x_i) = \text{var}(x_i) = \sigma_{x_i}^2$.

Let us say that there is a two-dimensional vector $\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$. The covariance matrix is given as:

$$\underline{\underline{\mathbf{cov}}}(\mathbf{x}) = \begin{pmatrix} \text{cov}(x_1, x_1) & \text{cov}(x_1, x_2) \\ \text{cov}(x_2, x_1) & \text{cov}(x_2, x_2) \end{pmatrix}, \quad (3.12)$$

$$= \begin{pmatrix} \sigma_{x_1}^2 & \text{cov}(x_1, x_2) \\ \text{cov}(x_2, x_1) & \sigma_{x_2}^2 \end{pmatrix}, \quad (3.13)$$

since $\text{cov}(x_2, x_1) = \text{cov}(x_1, x_2)$ (see symmetry of Equations 3.10 and 3.11),

$$\underline{\underline{\mathbf{cov}}}(\mathbf{x}) = \begin{pmatrix} \sigma_{x_1}^2 & \text{cov}(x_1, x_2) \\ \text{cov}(x_1, x_2) & \sigma_{x_2}^2 \end{pmatrix}. \quad (3.14)$$

The off-diagonal terms $\text{cov}(x_1, x_2)$ express how skewed the distribution should be in the $\begin{pmatrix} \sigma_{x_1} \\ \sigma_{x_2} \end{pmatrix}$ direction, and can take values within the closed interval of $[-\sigma_{x_1}\sigma_{x_2}, \sigma_{x_1}\sigma_{x_2}]$. Figure 3.1 shows the simulated result of sampling 100 measurements from three different two-dimensional probability distributions, and a $\chi^2 = 1$ contour line around each of them (see Section 3.2.2 for further explanation on the contour line). Each measurement is shown as a blue cross (+). The probability distribution is largest when the off-diagonal covariance terms = 0, and progressively reduces in size when this covariance deviates in either the positive or the negative direction. This description holds in higher dimensions as well, i.e. the probability distribution of $\mathbf{x}$ is most expansive when the off-diagonal elements of $\underline{\underline{\mathbf{cov}}}(\mathbf{x})$ are 0, implying that the uncertainty of $\mathbf{x}$ is largest when the covariance matrix is diagonal.




Figure 3.1: Plots of different distributions of $\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$, where $\text{var}(x_1)$ and $\text{var}(x_2)$ are kept constant but $\text{cov}(x_1, x_2)$ is varied.

3.2.1.1 The consequence of ignoring the off-diagonal terms of the covariance matrix

When using unfolding programs, the off-diagonal elements of the covariance matrix are typically assumed to be zero out of convenience, or because the option of inputting a full covariance matrix is not available to the program user. However in reality, the measured values in different elements of $\mathbf{N}_{\text{meas}}$ are often correlated with each other. To ignore these correlations by setting the covariances to zero would overestimate the uncertainty on $\mathbf{N}_{\text{meas}}$ leading to a more difficult unfolding and overestimating the uncertainty of the final unfolded result.

This is especially pertinent in the case of activation foil neutron spectrum unfolding where strong (off-diagonal) covariances may be present. One of the most obvious ways off-diagonal covariance arises in activation foil neutron spectrum unfolding is shown in the following hypothetical scenario: consider radionuclides A and B which emit γ -rays of very similar energies, forming a double-peak on the gamma-ray spectrum. The experimentalist might be very certain about the total number of decays of A and B (i.e. σ_{A+B} is small), but the individual uncertainty on the numbers of decays of A (σ_A) and B (σ_B) is large as the gamma-ray spectrum analysis software has difficulty differentiating whether each gamma-ray came from the decay of A or B . Therefore, $\text{var}(A) = \sigma_A^2$ and $\text{var}(B) = \sigma_B^2$ are both large (positive) numbers, but $\text{cov}(A, B)$ should be a large negative number, denoting that the overall uncertainty is quite small (akin to the left plot of Figure 3.1)). If this negative correlation is ignored and the off-diagonal elements of the covariance matrix are set to zero when performing the unfolding (akin to the middle plot of Figure 3.1), then the unfolding program will overestimate the uncertainty. This will lead




Figure 3.2: A plot showing the χ^2 of an arbitrary sample point (+) calculated using Equation 3.15. This χ^2 value can act as a metric to measure the deviation (represented by the black arrow ($\rightarrow$)) of any sample point from the mean ($\times$).

to the optimization process taking longer to search through a larger solution space, and the outputted solution ϕ_{sol} deviating further from ϕ_{true} and having a larger associated uncertainty.

This highlights one of the shortcomings of the majority of the unfolding programs: most existing unfolding programs are unable to account for this covariance (including the most frequently used GRAVEL [82], SAND-II [83] and MAXED [56]) as they are derived using the simplified definition of χ^2 (Equation 3.17) rather than the definition that includes covariances (Equation 3.15). Therefore, it has the same effect of ignoring covariances as the user will not be able to input the relevant covariance quantities even when they are present.

3.2.2 χ^2 calculation

When the covariance matrix $\underline{\text{cov}}(\mathbf{x})$ is known, the deviation of a particular sample $\mathbf{x}$ from the mean of a sample $\mathbf{x}_{\text{meas}}$ can be concluded with a scalar χ^2 as demonstrated by Figure 3.2. The $\chi^2 = \alpha$ contour line is drawn by colouring all points where $\chi^2 = \alpha$. The concentric, evenly spaced contour lines demonstrate that the χ^2 landscape in two dimensions forms a quadratic basin.

In general, for every point $\mathbf{x}$ in the n -dimensional vector space, the χ^2 quantifying its deviation from the sample mean $\mathbf{x}_{\text{meas}}$ can be calculated using Equation 3.15:

$$\chi^2(\mathbf{x}, \mathbf{x}_{\text{meas}}) = (\mathbf{x} - \mathbf{x}_{\text{meas}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{x}}(\mathbf{x} - \mathbf{x}_{\text{meas}}), \quad (3.15)$$

3.2 Uncertainty consideration

where $\underline{\underline{\mathbf{S}}}_x = (\underline{\underline{\mathbf{cov}}}(\mathbf{x}))^{-1}$ denotes the inverse of the covariance matrix, which is also a positive definite matrix,

$$\chi^2(\mathbf{x}, \mathbf{x}_{\text{meas}}) > 0 \quad \forall \mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{x}_{\text{meas}}\}, \quad (3.16)$$

i.e. except for the point $\mathbf{x} = \mathbf{x}_{\text{meas}}$, χ^2 is positive everywhere else.

One can verify that, when the off-diagonal elements of the covariance matrix are zero, this simplifies to the familiar χ^2 equation (Equation 3.17) often used in entry-level experimental physics,

$$\chi^2 = \sum_{i=1}^m \frac{(x_{\text{meas},i} - x_{\text{model},i})^2}{\text{var}(x_{\text{meas},i})}, \quad (3.17)$$

where $x_{\text{meas},i}$ is the value of the i^{th} datapoint, $\text{var}(x_{\text{meas},i})$ is the variance associated with that datapoint, and $x_{\text{model},i}$ is the value that same datapoint should have according to the model's prediction.

As stated in Section 3.1.3, while the ideal expression to minimise would be

$$\chi^2 = (\boldsymbol{\phi}_{\text{true}} - \boldsymbol{\phi}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\boldsymbol{\phi}_{\text{true}}} (\boldsymbol{\phi}_{\text{true}} - \boldsymbol{\phi}_{\text{sol}}), \quad (3.18)$$

it is not possible to calculate this because both $\boldsymbol{\phi}_{\text{true}}$ and its associated covariance matrix $\underline{\underline{\mathbf{cov}}}(\boldsymbol{\phi}_{\text{true}}) = (\underline{\underline{\mathbf{S}}}_{\boldsymbol{\phi}_{\text{true}}})^{-1}$ are not known.

In lieu of this, some algorithms calculate the goodness-of-fit in the form shown by text Equation 3.9,

$$\chi^2 = (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}), \quad (3.19)$$

where $\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} = (\underline{\underline{\mathbf{cov}}}(\mathbf{N}_{\text{meas}}))^{-1}$ is the inverse of the covariance matrix associated with $\mathbf{N}_{\text{meas}}$. Some algorithms extend the χ^2 definition to include the deviation of $\boldsymbol{\phi}_{\text{sol}}$ from $\boldsymbol{\phi}_0$ as well,

$$\chi^2 = (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) + (\boldsymbol{\phi}_0 - \boldsymbol{\phi}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\boldsymbol{\phi}_0} (\boldsymbol{\phi}_0 - \boldsymbol{\phi}_{\text{sol}}), \quad (3.20)$$

where $\underline{\underline{\mathbf{S}}}_{\boldsymbol{\phi}_0} = (\underline{\underline{\mathbf{cov}}}(\boldsymbol{\phi}_0))^{-1}$ is the inverse of the covariance matrix associated with the *a priori* spectrum $\boldsymbol{\phi}_0$ as provided by the user.

The most rigorous algorithms make use of uncertainty information available on all 3 variables of Equations 3.6. [82]

$$\chi^2 = (\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}}' \phi_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}}' \phi_{\text{sol}}) + (\phi_0 - \phi_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\phi_0} (\phi_0 - \phi_{\text{sol}}) + (\underline{\underline{\mathbf{R}}} - \underline{\underline{\mathbf{R}}}') \cdot \underline{\underline{\mathbf{S}}}_R (\underline{\underline{\mathbf{R}}} - \underline{\underline{\mathbf{R}}}') , \quad (3.21)$$

where, instead of assuming that the response matrix is accurate, the algorithms generate another response matrix $\underline{\underline{\mathbf{R}}}'$, and use that in place of the originally provided $\underline{\underline{\mathbf{R}}}$. By simultaneously adjusting all elements in ϕ_{sol} and $\underline{\underline{\mathbf{R}}}'$, the algorithm finds a solution $(\phi_{\text{sol}}, \underline{\underline{\mathbf{R}}}')$ that would fit into Equation 3.6 to produce the minimum χ^2 as defined by Equation 3.21. The covariance tensor (a straightforward generalization of Equation 3.12) associated with $\underline{\underline{\mathbf{R}}}$ is a 4-tensor whose tensor inverse is given as $\underline{\underline{\mathbf{S}}}_R$. In theory, an algorithm that follows Equation 3.21 is perfectly suited for use in activation foil neutron spectrum unfolding. Because of the large uncertainty in the nuclear data, especially at fusion-relevant energies (e.g. 14.1 MeV, which can only be reached by very few neutron irradiation facilities currently), the response matrix is often unreliable [84]. In practice, however, such an algorithm would have $(m + 1) \times n$ adjustable parameters. If the algorithm requires multiple matrix-inverse per step, it would take up an inordinate amount of computing resources, especially in fusion applications where n can reach up to 709 bins [85]. There is no guarantee that such an algorithm could outperform one that does not account for the uncertainty on $\underline{\underline{\mathbf{R}}}$, as the number of free parameters grow to $(m + 1) \times n$ (compared to the number of constraints m), the problem becomes extremely underdetermined. The solution space also increases in size exponentially, potentially requiring an infeasible amount of computational time and memory to find a solution. Thus the time required to derive, implement and test it as a program may exceed the length of this PhD. Inputting the covariance tensor $\underline{\underline{\mathbf{S}}}_R$ in the correct group structure into this hypothetical program would be a tedious task [71], making it user-unfriendly. Hence a decision was made to exclude nuclear data uncertainty ($\underline{\underline{\mathbf{S}}}_R$) from the unfolding algorithms produced in this PhD.

On a more positive note, various facilities such as IFMIF-DONES [86] and facilities under NNUF may be commissioned for neutron irradiation experiments at fusion-relevant energies in the near future, which will help fill in these gap in the nuclear data and reduce their uncertainties, reducing the impact of the last term in Equation 3.21 and improving the precision of the outputted neutron spectra. If that is successful, then the mathematical frameworks derived in Chapter 4 may be extended to cover the definition of χ^2 as described by Equation 3.21 to create even more comprehensive unfolding algorithms that account for $\underline{\underline{\mathbf{S}}}_R$.




Figure 3.3: The geometry of the χ^2 landscape depending on the degree of determination.

3.3 Effect of degree of determination

The mathematics underlying unfolding problems fundamentally differ depending on whether they are overdetermined, fully-determined, or underdetermined. This section heavily relies on the concept of confidence intervals in Section 3.2.2 (where $\chi^2 < \text{some positive real number}$ as shown in Figures 3.1 and 3.2) to visualize how the sign of $m - n$ (degree of determination) affects the geometry of the solution space.

Consider a very simple 2-bin neutron spectrum example $\boldsymbol{\phi}^T = (\phi_1, \phi_2)$, represented by the x-axis and y-axis of each of the three plots in Figure 3.3. Let us say that the true neutron spectrum is $\boldsymbol{\phi}_{\text{true}}^T = (20, 22)$. After measurement by three neutron spectrometers, which have respective response matrices of $\underline{\mathbf{R}}_1$, $\underline{\mathbf{R}}_2$, and $\underline{\mathbf{R}}_3$ (left, middle and right plots respectively), they yield three $\mathbf{N}_{\text{meas}}$ vectors (with some error $\boldsymbol{\varepsilon}$ on each of them), as described by Equation 3.7. To keep the example simple we assume that the standard deviations of $\mathbf{N}_{\text{meas}}$ are equal to its square roots, and the covariance matrices are all diagonal (i.e. $\text{cov}(N_{\text{meas},i}, N_{\text{meas},j}) = \delta_{ij}\sigma_i^2 = \delta_{ij}N_{\text{meas},i}$).

There are four points of interest here:

1. $\boldsymbol{\phi}_{\text{true}} \neq \boldsymbol{\phi}_{\text{sol}}$ — The true spectrum $\boldsymbol{\phi}_{\text{true}}$ does not coincide with the point(s) of lowest χ^2 as long as $\boldsymbol{\varepsilon} \neq \mathbf{0}$. This is a fairly common occurrence throughout experimental physics, where the complete minimization of the χ^2 may only lead us *close to*, but not exactly on, the true solution. In a real experiment, we do not know where is $\boldsymbol{\phi}_{\text{true}}$. However, (assuming the covariance matrix of $\boldsymbol{\phi}_{\text{sol}}$ = the covariance matrix of $\boldsymbol{\phi}_{\text{true}}$) we can still say

3.3 Effect of degree of determination

that ϕ_{true} has a 68% chance of lying inside the $\chi^2 \leq 1$ region, 95% chance of lying inside the $\chi^2 \leq 4$ region, 99.7% chance of lying inside the $\chi^2 \leq 9$ region, etc.

2. No unique solution is available in the case of underdetermined systems — The χ^2 reaches 0 (the minimum χ^2 value) at multiple (infinitely many) locations for the underdetermined ($m < n$) case, therefore no unique solution is available. Bearing in mind that the black cross $+$ is invisible to the experimentalist since ϕ_{true} is unknown, thus in the underdetermined case, there must be infinitely many solutions, each of them equally valid, unless further regularising conditions are applied. In this example (Figure 3.3 left), there is an infinitely long straight line of points where χ^2 reaches the minimum of zero. This 1-dimensional object (line) embedded in 2 dimensions (plane) is called a manifold, which will be explained in greater detail in Section 3.3.1. The direction where the χ^2 stays the same (direction parallel to the contour lines drawn) is called the singular direction, and the direction where the χ^2 varies (perpendicular to contour lines) is called the nonsingular direction. The χ^2 value decreases and then increases again in a parabolic manner as we move along the nonsingular direction, but remains constant in the singular direction. The fact that there are infinitely many possible solutions is at the heart of why different unfolding algorithms can be given the same problem and produce very different solutions, all of which are considered as valid as each other.
3. Contrastingly, there is one and only one point of minimum χ^2 in any fully- or overdetermined ($m \geq n$) system — even though there is no guarantee that the point of minimum χ^2 corresponds to a physically sensible spectrum (i.e. negative flux might be present in some bins at that point) unless regularising conditions have been added to modify this χ^2 landscape. If the experiment is repeated many times, the values of χ^2 found at this minimum follow the χ^2 distribution with $m - n$ degrees of freedom (Figure 3.4) — this means that the minimum χ^2 is, on average, equal to $(m - n)$, and the probability of having χ^2 equal exactly 0 is vanishingly small when $m > n$. In contrast, when $m \leq n$, the minimum χ^2 will always be 0. There are no singular directions when $m \geq n$, and the χ^2 profile of any straight line through this ϕ_{sol} space, in any arbitrary direction, traces out a quadratic curve.
4. The uncertainty envelope has infinite measure in the case of underdetermined systems — in Figures 3.1 to 3.3, one can take the uncertainty envelope as the space in which $\chi^2 \leq 1$, which has a finite area. As evident from the left plot of Figure 3.3, when the




Figure 3.4: The probability distribution functions of minimum χ^2 values achievable at different degrees of determinations $\nu = m - n$. Note that there are *no valid* χ^2 distributions when the system is fully- or underdetermined ($\nu \leq 0$).

system becomes underdetermined, this area value becomes infinite. Generalised to higher dimensions than 3, we can simply say that the “measure” of the uncertainty envelope is always infinite when the system is underdetermined. This makes comparing the precisions of two underdetermined systems impossible, as we cannot meaningfully compare one infinity against another. This has implications in the optimization step of the foil selection process detailed in Section 5.5.1, where we need to compare the precision of different systems against each other to find the optimal one. To ensure that step works with underdetermined activation foil unfolding systems as well, an alternative function than the measure of the $\chi^2 \leq 1$ envelope needs to be developed to act as the loss function for evaluating the quality of the choice of foils.

These fundamental differences in under-, fully-, and overdetermined problems lead to unfolding algorithms exhibiting different behaviours in the respective regimes. The most pronounced difference is exemplified in Section 4.3.4, showing that the pseudo-inverse algorithm (one of the simplest unfolding algorithms) behaves equivalently to Equations 4.44, 4.45, and 4.46 in these situations respectively.

3.3.1 Re-examining the difficulty of the unfolding problem using the concept of manifolds

For the purpose of this thesis, a manifold is a lower-dimensional object embedded in a higher-dimensional space and is a useful concept on which we can build our understanding of unfolding problems. For example, in Figure 3.3 left, all points where minimum χ^2 is achieved form a line,

3.3 Effect of degree of determination

which is a line embedded in a plane, so it is called a 1-D manifold embedded in 2-D space, and the surface of a 3-D sphere can be considered as a 2-D manifold embedded in 3-D space.

The set of all physically realistic spectra should form a manifold of many fewer dimensions than the set of all outputs allowed by the algorithms (the solution space ϕ). The former set excludes physically impossible spectra (e.g. any spectra where some of the fluxes are negative, or have unphysical oscillations in the neutron spectrum), forming a manifold with dimensionality $= p$, where $p < n =$ the dimensionality of the solution space ϕ .

This has two implications:

The first implication acts as a guideline when designing the response matrix. A well-chosen response matrix should be able to preserve all information about the ϕ_{true} . In mathematical terms, the response matrix should act as an injective function that uniquely maps every point on the manifold (a.k.a. domain of ϕ_{true}) of physically realistic neutron spectra uniquely into a unique point in the domain $\mathbf{N}_{\text{meas}}$. More formally, a good choice of response matrix $\underline{\mathbf{R}}$ should satisfy

$$\begin{aligned} \underline{\mathbf{R}} : \text{domain of } \phi_{\text{true}} &\mapsto \text{domain of } \mathbf{N}_{\text{meas}}, \text{ such that} \\ \forall (\phi_{\text{true1}}, \phi_{\text{true2}}) \in \text{domain of } \phi_{\text{true}}, \quad \underline{\mathbf{R}}\phi_{\text{true1}} = \underline{\mathbf{R}}\phi_{\text{true2}} &\implies \phi_{\text{true1}} = \phi_{\text{true2}}. \end{aligned} \quad (3.22)$$

Only if such an injective relationship exists, then in the absence of noise ϵ , an ideal unfolding algorithm would be able to act as the inverse of Equation 3.22, such that

$$\text{Unfolding algorithm} : \text{domain of } \mathbf{N}_{\text{meas}} \mapsto \text{domain of } \phi_{\text{true}}. \quad (3.23)$$

This means that the true complexity of the unfolding problem is proportional to the dimensionality of this manifold P , not the dimensionality of the group structure n chosen. In order to preserve all information about ϕ_{true} , then the response matrix only needs to have rank $= \min(m, n) \geq p$. This demonstrates the importance of selecting foils properly to maximise m , which, in the case of activation foil neutron spectrum unfolding, is equal to the number of distinct species of radionuclides that can be measured.

Of course, to find out the dimensionality of P , we must have sufficient *a priori* knowledge of the type of ϕ_{true} that we may encounter. This brings us to the second implication, which is a revelation that there is insufficient *a priori* information about what kind of neutron spectrum we would encounter in the first wall of a commercial fusion power plant. Currently, not much is known about this manifold P for fusion spectra, as there are no studies that comprehensively capture the range of neutron spectra that may arise in the first wall of a power-generating fusion




Figure 3.5: A sample of possible neutron spectra at the first wall of fusion reactors, collected from [85]. These can be thought of as members of the set P , where $P =$ all physically realistic fusion neutron spectra at the first wall of fusion reactors.

reactor while also accounting for downscattering in the first wall. While existing studies have generated points on this manifold P (Figure 3.5), the overall shape of P is yet unknown. This absence of *a priori* knowledge places a restriction on the choices of algorithms applicable to fusion neutron spectra unfolding, discussed in greater detail in Section 4.1.2.

As this manifold of physically realistic fusion neutron spectra and its dimensionality p are yet indeterminate, we can only assume that p takes the largest possible value $= n =$ the number of bins in the group structure. This assumption means that when designing the ideal response matrix $\underline{\mathbf{R}}$, it would have rank $= n$, which as we can see in Chapter 6, is difficult to achieve.

Chapter 4

Unfolding algorithms review and improvement

Armed with the basic mathematical understanding of unfolding established by Chapter 3, we now have the language necessary to delve into the approaches taken by different algorithms and understand why there are so many of them. In Section 4.1, unfolding algorithms are divided into several categories to allow in-depth discussions into the nature of each. Where relevant, the specific unfolding procedures used by an example algorithm or two in that category of unfolding algorithms are detailed, and then compared and contrasted with each other. Then, Section 4.2 explains the reason for the numerousness of types of unfolding algorithms, and acknowledges the insufficiencies of existing algorithms, particularly when applied to fusion neutron spectra measurements. Using knowledge gained by dissecting existing algorithms, Section 4.3 produces several new algorithms suitable for fusion neutron spectrum unfolding (but can equally be applied to general-purpose unfolding) to address these insufficiencies.

4.1 Classification of unfolding algorithms

As mentioned in Section 3.1.3, the inclusion of *a priori* information is crucial for unfolding algorithms. While all algorithms attempt to minimise the χ^2 (as defined by one of the three Equations 3.19, 3.20, and 3.21) in very similar ways, their different ways of using *a priori* information can be categorised into four different groups as shown in Figure 4.1. This scheme is an original work of this PhD. Compared to the categorization scheme created by Matzke [82] and often followed by subsequent in-depth research projects [87], the following scheme is more up-to-date and exhaustive, and more cleanly divides unfolding algorithms according to how




Figure 4.1: Categories of unfolding algorithms, with examples at the bottom level.

uncertainty propagation can be achieved for each category. The following four sections will explain each category, and list example algorithms for each where relevant.

4.1.1 Machine learning

Sometimes called Artificial Intelligence (AI), machine learning algorithms are algorithms that apply supervised machine learning techniques to minimise the χ^2 value as their loss function.

The most common machine learning technique is using Neural Networks (NN) to learn from a large set of measured or simulated neutron spectra [88]. Typically the responses $\mathbf{N}$ is inputted along with the neutron spectra ϕ as feature-label pairs so that the neural network can adjust the weights and biases inside the network and learn to replicate the relationship between $\mathbf{N} \longleftrightarrow \phi$ without having explicit knowledge of the response matrix $\underline{\mathbf{R}}$ or calculating the inverse of the response matrix $\underline{\mathbf{R}}$. The neural network implicitly learns the shape of the manifold of solution space ϕ_{sol} and how any given $\mathbf{N}_{\text{meas}}$ should be mapped to that manifold in the training phase, such that when a new $\mathbf{N}_{\text{meas}}$ is given in the testing phase, it returns a corresponding point ϕ_{sol} that should also be on that manifold. The pre-requisite of this technique is that a sufficiently large database of validated neutron spectra must be provided as the training set, and this minimum size requirement scales with the dimensionality of the label (n) and feature (m). While there is no shortage of interest from smaller research groups with

less fusion-focused research interest [48, 89–91] to use neural networks for neutron spectrum unfolding, all of their attempts are domain-specific, so a neural network trained to unfold neutron spectra in a specific accelerator facility cannot be used to unfold the neutron spectra in a different facility where neutrons are generated differently; and no attempts have been made at transfer learning to bridge this gap between domains. Currently, the only promising applications of neural network unfolding in fusion devices are limited to real-time neutron spectra unfolding of larger diagnostics such as MPR¹ and TOFOR² in JET, where the neutron spectra data are abundant, well-known and consistent [92]. However the neutron spectrum at the first wall of a future fusion reactor remains hypothetical until it is built, so no measurements of these spectra currently exist. The only neutron spectra which are available for training are based on wildly different designs in first wall materials, tritium breeding blankets, and cooling systems [85]. Due to the curse of dimensionality [93], the possible solution space of ϕ_{sol} may be very large due to the large number of groups that are used (e.g. 175 for the VITAMIN-J group structure), so there is no guarantee that the current dataset of simulated fusion neutron spectra will cover the manifold of physically-realistic spectra ϕ_{true} densely. This problem is particularly bad in cases of large degrees of underdetermination, where the manifold of ϕ_{sol} (which has a maximum dimensionality of m) is much smaller than the space that it is embedded in (of dimensionality n). While the dimensionality of the solution space can hypothetically be reduced by the use of physics-informed neural networks (PINNs) to make the training of neural networks more tractable, there is a noticeable gap in the literature around the use of PINNs in particle spectra unfolding. Therefore, we consider the use of neural networks too immature for fusion applications at this stage and remains an active area of research [94].

Another commonly used machine learning technique is the Genetic Algorithm (GA) [95–97], where neutron spectra that fit Equation 3.19 well (smaller χ^2 = higher fitness value) are mutated or spliced with each other, selecting for the best fitting neutron spectrum ϕ_{sol} in a similar way that chromosomes of organisms are mutated and spliced together to select for optimal fitness in nature [98]. The use of GA in fusion neutron spectrum remains an active area of research as well [98].

To the author’s knowledge, unlike the other classes of unfolding algorithms, there are no machine learning algorithms that are adopted by more than one research group (likely due to the lack of portability and domain-specificity of machine learning algorithms), thus there are no famous examples that can be given for this class of unfolding algorithms.

¹Magnetic Proton Recoil spectrometers

²Time-Of-Flight neutron spectrometer Optimised for high count Rate

4.1 Classification of unfolding algorithms

Compared to traditional deterministic algorithms, machine learning algorithms are less explainable and hence cannot guarantee the correctness of the solution. This makes uncertainty propagation (from $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}})$ to $\underline{\underline{\text{cov}}}(\phi_{\text{sol}})$, or from $\underline{\underline{\underline{\text{cov}}}}(\mathbf{R})$ to $\underline{\underline{\text{cov}}}(\phi_{\text{sol}})$) difficult. They all have an element of stochasticity, such as the random mutations in GA and random initialization of the weights and biases in NN. Hence, for a budding field of research such as neutron spectrum unfolding in fusion reactors, where it is paramount that we understand its details to ensure correctness and have a reliable algorithm that generates reproducible results, machine learning is not a suitable choice when more deterministic traditional unfolding algorithms exist; at least not until the technology has matured and we have an extensive collection of neutron spectrum data and a robust understanding of them.

4.1.2 Parametric unfolding

As explained in Section 3.3.1, the solution space to be searched through may be reduced using the fact that the physically allowed ϕ_{true} lies on a p -dimensional manifold, instead of the full n -dimensional solution space of ϕ . We assume that this manifold can be fully captured by a parametric model $\mathbf{P}$ that takes in a list of p parameters arranged into the vector $\boldsymbol{\pi}$ where $\boldsymbol{\pi} \in \mathbb{R}^p$, and outputs a physically realistic spectrum ϕ :

$$\mathbf{P} : \text{domain of } \boldsymbol{\pi} \mapsto \text{domain of } \phi_{\text{true}} . \quad (4.1)$$

The optimally fitting list of parameters $\boldsymbol{\pi}_{\text{sol}}$ can simply be given as

$$\boldsymbol{\pi}_{\text{sol}} = \underset{\boldsymbol{\pi}}{\text{argmin}} \left(\chi^2(\mathbf{N}_{\text{meas}}, \underline{\underline{\mathbf{R}}}\mathbf{P}(\boldsymbol{\pi})) \right) , \quad (4.2)$$

where χ^2 is defined by Equation 3.19, such that the solution spectrum is given by

$$\phi_{\text{sol}} = \mathbf{P}(\boldsymbol{\pi}_{\text{sol}}) . \quad (4.3)$$

This simplifies the problem from optimization of ϕ into optimization of $\boldsymbol{\pi}$, which is a smaller solution space to search through. Since we expect the number of parameters p to be much smaller than the number of bins in the group structure n , the solution space to search through is much smaller, so the optimization problem becomes more tractable. p is also assumed to be smaller than m , making the unfolding problem simpler.

4.1 Classification of unfolding algorithms

Unlike the response matrix $\underline{\mathbf{R}}$ which turns the ϕ into $\mathbf{N}$ via a simple linear combination (see Equations 3.5), the parametric model typically involves many nonlinear operations. For example, the intensity of neutrons by the plasma near the DT peak may be related to the average plasma ion temperature T_i by the following relation: number of neutrons with energy $E \propto \frac{1}{\sqrt{T_i}} \exp(-\frac{(E-14.06 \text{ MeV})^2}{2T_i})$. Such nonlinear operations in the function $\mathbf{P}$ would make the χ^2 landscape over π space more complex, such that this landscape is no longer made of simple quadratic components like those pictured in Figure 3.3. Optimization can no longer be done by simple gradient descent or Newton’s method, as it may overshoot the minimum point or get trapped in local (i.e. non-global) minima. However the decrement in the complexity of the problem due to the reduction in the dimensionality from n to p may make up for, or even outweigh, the increment in complexity due to the nonlinearity of the χ^2 landscape. The parametrised solution space is usually small enough (i.e. p is of low enough dimension) that sophisticated Bayesian parameter estimation algorithms such as Markov Chain Monte Carlo (MCMC) can be applied [81]. The additional benefit of MCMC is that it performs uncertainty propagation as part of the calculation, without the need for writing/using auxiliary code or derivation to obtain the covariance matrix associated with the solution ϕ_{sol} . Two examples of these algorithms are WinBUGS and FRUIT [99, 70]. This type of unfolding has the benefit of always producing a physically realistic neutron spectrum.

A drawback of these algorithms is that they require good knowledge about the space of possible spectra ϕ_{true} when constructing the parametric model, a luxury that is not available to the cutting-edge field that is fusion neutronics. Currently, some parametric models of the plasma as a neutron source exist, but there are no parametric models of the first wall spectrum. One of the most prominent examples of programs that simulate the neutron spectrum given out by the plasma is DRESS [100], which has been used to simulate neutron spectra in JET [101], but over-predicted the neutron yield in MAST by up to 67% [29] due to the guiding-centre approximation breaking down near the central solenoid of the spherical tokamak; even though such problem did not appear when DRESS was used to model JET, a tokamak with conventional aspect ratio. This shows that our understanding of plasma neutron sources is still too limited for extrapolation to reactors of different designs and sizes.

Therefore, we conclude that parametric fusion neutron spectrum unfolding cannot be readily applied in commercial fusion power plants. It would require further collaboration between plasma physicists and neutronics scientists to create a model that accurately simulates the neutron generation, breeding, and downscatter processes from the plasma to the breeding blanket of the reactor before parametric unfolding algorithms can become viable candidates.

Until such a study has been made, we can only guess for what such a parametric representation of the fusion neutron spectrum should look like. It can be broken down into several components:

- a DT neutron peak centred at 14.06 MeV, described by a minimum of two parameters (amplitude and width, possibly skew and kurtosis³);
- a DD neutron peak centred at 2.45 MeV, described by a minimum of two parameters (amplitude and width, possibly skew and kurtosis);
- a high-energy tail ($> 14.06 \text{ MeV}$) caused by alpha knock-on reactions that re-distribute energy into the plasma, described by at least one variable (decay rate of this tail w.r.t. to increasing E);
- a low energy tail of both peaks as neutrons get downscattered and then absorbed by materials in the first wall, divertor, the breeder blanket, and further beyond [105, 106]. The rate at which these two tails decay away as E decreases is dependent on the material of the tokamak; they can be described by a minimum of two variables (the neutron absorption rate in the thermal region should be very different to the absorption rate in the non-thermal region, thus at least 2 variables are needed. The more variables used, the more faithfully we can replicate this low-energy tail using the parametric model.)

Regardless, the development of such a parametric model lies beyond the scope of this PhD, therefore it will not be discussed any further in this thesis.

4.1.3 Regularising

The remaining two classes of algorithms (the right-hand half of the classification tree in Figure 4.1) use only the information found in a single user-provided *a priori* neutron spectrum, ϕ_0 , to overcome the ill-posed nature of the unfolding problem. Regularising unfolding algorithms minimise a loss function that is a sum of χ^2 , as defined by Equation 3.19, 3.20, or 3.21, and optionally includes an extra regularising function that takes in ϕ_{sol} as the only parameter,

$$\text{loss} = \chi^2(\mathbf{N}_{\text{meas}}, \mathbf{N}_{\text{sol}}) + \text{RegularisingFunction}(\phi_{\text{sol}}). \quad (4.4)$$

Different regularising unfolding algorithms differ only in their choice of regularising function.

³If the fusion device is sufficiently small such that neutral beam injection and radio-frequency heating plays a significant role in heating the plasma, then the shape of the DD and DT peak would become more complicated due to the contribution to fusion yield by fast deuteron and triton ions, leading to an anisotropic plasma neutron spectrum that shift the peaks' centroids depending on the viewing direction [28, 102–104].

4.1 Classification of unfolding algorithms

The purpose of the regularising function is to simplify the underdetermined χ^2 minimization problem into an overdetermined loss function minimization problem, so that there will always be a finite number of local minima in the loss function landscape, and also to include extra information about the possible ϕ_{true} so that the loss function minima should be as close to ϕ_{true} as possible. The simplification of the loss function landscape allows even the simplest minimization algorithms, such as Newton's method or the (unmodified) Gradient Descent algorithm, to find a solution ϕ_{sol} effectively.

Some algorithms take Equation 4.4 one step further by giving the user some control over the strength of the regularization function, using a positive parameter τ known as the Tikhonov regularization parameter:

$$\text{loss}_\tau = \chi^2(\mathbf{N}_{\text{meas}}, \mathbf{N}_{\text{sol}}) + \tau \cdot \text{RegularisingFunction}(\phi_{\text{sol}}) , \text{ where } \tau \in \mathbb{R}^+ . \quad (4.5)$$

This is an often employed tactic outside of spectra unfolding [107, 108]. τ represents the ratio of the weight of the first term (deviation of $\mathbf{N}_{\text{sol}}$ from $\mathbf{N}_{\text{meas}}$) compared to the weight of the second term (deviation of ϕ_{sol} from ϕ_0), which can be adjusted. If τ is too big, too much weight would be placed on the *a priori* spectrum, giving a solution ϕ_{sol} that barely deviates from ϕ_0 without fitting to $\mathbf{N}_{\text{meas}}$; while a τ that is too small would lead to overfitting to the $\mathbf{N}_{\text{meas}}$, amplifying the effect of noise (ϵ) to yield a nonsensical spectrum. The regularization parameter τ can also be thought of as a unit-conversion factor to reconcile the fact that the two quantities may be inequivalent in nature in cases where the regularising function produces an output that is not dimensionless. To find the correct Tikhonov regularization parameter τ , instead of guessing and picking a single value of τ and performing the optimization only once, multiple optimizations are performed, each one with a different value of τ . Each optimization result will yield two components in the loss function, the χ^2 value and the regularising function value. These values can then be plotted on a 2D-Cartesian coordinate, with χ^2 value on the x-axis and the regularising function on the y-axis. Since increasing the weight should only cause the solution to deviate further from the $\mathbf{N}_{\text{meas}}$ and approach ϕ_0 , the point of (χ^2 , regularising value) is expected to move right and downwards as τ increases, tracing out a monotonically decreasing curve. This curve is called an L-Curve [108], an example of which is shown in Figure 4.2. It is no coincidence that the L-curve shown in Figure 4.2 has a monotonically increasing 1st derivative and a relatively stable 2nd derivative: its loss function was carefully picked out in Section 4.3.3 such that it cooperates nicely with χ^2 . However, in general, L-curves can be quite pathological in nature, sometimes possessing multiple sharp changes in their 1st derivatives (i.e. does not appear

4.1 Classification of unfolding algorithms




Figure 4.2: The L-curve produced using AMAXED-Regularization, an algorithm that will be discussed in Section 4.3.3. D_{KL} is defined in Section 4.1.3.3.

to be smooth). Often, users would choose a value of τ where the value, 1st derivative, and 2nd derivatives of the L-curve change sharply or reach certain thresholds, thus obtaining their final solution ϕ_{sol} . A notable unfolding algorithm that utilises this L-curve method is a variant of MAXED called MAXED-L [109], where MAXED is an algorithm that uses information entropy as the regularising function, discussed in Section 4.1.3.3.

One of the advantages of regularising algorithms is that the uncertainty of its solution can be analytically derived. We can find the lowest-order approximation of the uncertainty on the solution spectrum ϕ_{sol} as the $n \times n$ covariance matrix $\underline{\underline{\text{cov}}}(\phi_{\text{sol}})$, given by the following formulae:

$$\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}} = \nabla_{\phi}(\nabla_{\phi}(\text{loss})), \quad (4.6)$$

$$(\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}})_{ij} = \left. \frac{\partial^2(\text{loss})}{\partial \phi_i \partial \phi_j} \right|_{\phi=\phi_{\text{sol}}}, \quad (4.7)$$

such that

$$\underline{\underline{\text{cov}}}(\phi_{\text{sol}}) = \underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}^{-1}. \quad (4.8)$$

The 2nd derivative of a scalar variable with respect to a vector variable is known as the Hessian matrix. In this case, the Hessian matrix (Equation 4.6) equals the inverse of the covariance matrix associated with ϕ_{sol} (Equation 4.8), meaning the higher the curvature of the loss function around ϕ_{sol} is, the less uncertainty we have about the solution ϕ_{sol} . A good choice of regularising

4.1 Classification of unfolding algorithms

function would ensure that the loss function does not have any singular directions at ϕ_{sol} , such that the Hessian matrix $\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}$ is invertible, allowing Equation 4.8 to be computed.

4.1.3.1 Linear least-squares

For some algorithms, the regularising function is simply the mean squared deviation from the *a priori* spectrum $(\phi_0 - \phi_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\phi_0} (\phi_0 - \phi_{\text{sol}})$, so Equation 4.4 becomes identical to Equation 3.20 (e.g. [77]). More sophisticated algorithms might adopt Tikhonov regularization (Equation 4.5) to improve upon this (e.g. [110, 111]), such that the loss function becomes:

$$\text{loss} = (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) + \tau (\phi_0 - \phi_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\phi_0} (\phi_0 - \phi_{\text{sol}}). \quad (4.9)$$

The extra degree of freedom granted by the Tikhonov parameter τ gives the user to control the relative weight of the second term against the first. While they are conceptually simple and their Hessian matrices easily derived, none of these algorithms guard against unphysical spectra (negative fluxes).

4.1.3.2 Roughness of the spectrum

A slightly more sophisticated class of algorithm uses not just the neutron spectrum ϕ_{sol} itself, but also its derivatives. These algorithms directly measure either the self-information (e.g. [112, 113]) ($H(\mathbf{p})$ in Equation 4.13 in the following Section 4.1.3) or Fisher information content (e.g. [114]) of the neutron spectrum ϕ_{sol} , without comparing it to an *a priori*. The definition of Fisher information content of the neutron spectrum is given by:

$$\text{FisherInformation}(\phi) = \int_0^\infty \frac{1}{\phi(E)} \left(\frac{d\phi(E)}{dE} \right)^2 dE. \quad (4.10)$$

This is equivalent to measuring the mean squared 1st derivative in ϕ_{sol} . A typical Fisher information regularising algorithm would use a loss function of the form:

$$\text{loss} = (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) + \text{FisherInformation}(\phi_{\text{sol}}). \quad (4.11)$$

A closely related class of algorithms are those that attempt to minimise the sum of squares of the 2nd derivative of the log of neutron spectrum [115].

$$\text{loss} = (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) + \sum_{i=2}^{n-1} \left(\log \left(\frac{\phi_{\text{sol},i-1} \phi_{\text{sol},i+1}}{\phi_{\text{sol},i}^2} \right) \right)^2, \quad (4.12)$$

4.1 Classification of unfolding algorithms

where the smoothness of the neutron spectrum is maximised if “Roughness” in Equation 4.12 is used as the regularising function. These types of algorithms that maximise the smoothness of the spectrum have the advantage of not requiring the use of an *a priori* spectrum at all, so they are also one of the few classes of algorithms that allow prior-less unfolding.

However, in fusion applications, Fisher information regularization has only shown limited success in unfolding liquid scintillator spectra [38] (which is usually overdetermined and therefore easier to solve than activation-foil neutron-spectrum unfolding problems), with much poorer results in unfolding fusion spectrum using activated materials [59]. This demotivating result, in addition to the presence of a physical process that promotes discontinuity (inelastic scattering creates a discontinuity in the neutron spectrum [75]), casts doubts about the validity of assuming future fusion power plants to have smooth neutron spectra. This steered this thesis’ investigation towards a direction that encourages the use of less presumptuous regularising functions.

4.1.3.3 Information Entropy

The final class of regularising algorithms utilises information theory more rigorously to obtain the neutron spectrum. In the context of information theory, “Entropy” is used to measure the information content that a probability distribution holds. Consider two different discrete normalized probability distributions $\mathbf{p}$ and $\mathbf{q}$, such that $\sum_{i=1}^n p_i = 1 = \sum_{i=1}^n q_i$. There are three important quantities that can be derived from $\mathbf{p}$ and $\mathbf{q}$: [116–118]

$$\text{Entropy} = H(\mathbf{p}) = - \sum_{i=1}^n p_i \log(p_i). \quad (4.13)$$

$$\text{CrossEntropy} = H(\mathbf{p}, \mathbf{q}) = - \sum_{i=1}^n p_i \log(q_i). \quad (4.14)$$

$$\text{Kullback–LeiblerDivergence} = D_{KL}(\mathbf{p}, \mathbf{q}) = - \sum_{i=1}^n p_i \log\left(\frac{q_i}{p_i}\right). \quad (4.15)$$

All three quantities above are nonnegative, i.e. $H(\mathbf{p}, \mathbf{q}), H(\mathbf{p}), D_{KL}(\mathbf{p}, \mathbf{q}) \geq 0$ regardless of what distributions are chosen as $\mathbf{p}$ and $\mathbf{q}$. All three quantities have unit “nat”⁴. An interactive example is available at <https://www.desmos.com/calculator/yqafsxccuz> demonstrating these 3 concepts in the 3-bin ($n = 3$) case using a ternary diagram⁵. In these interactive examples, one

⁴If the natural logs in Equations 4.13 to 4.15 are replaced with $\log_2$, then the units become “bit”. If replaced with $\log_{10}$, then the units become “Hart”.

⁵A similar interactive ternary diagram can be found at <https://www.geogebra.org/3d/uucjwfdy> or <https://www.desmos.com/3d/5dfa197ef3>, where the Kullback–Leibler Divergence $D_{KL}(\mathbf{p}, \mathbf{q})$ is plotted as a 2-D surface over all of the allowed ranges of $q_1, q_2, q_3 \geq 0$ and $q_1 + q_2 + q_3 = 1$ for a user-defined distribution of $\mathbf{p}$. However due to the 3-D nature of this simulation, it is computationally more intensive, so the former link is slow to update while the latter is prone to displaying unphysical artefacts.

4.1 Classification of unfolding algorithms

can verify for themselves that when any of the 3 bins of $\mathbf{q}$ approaches 0, the cross-entropy and Kullback–Leibler Divergence approach infinity, and so do their derivatives. This fact will be useful later in the derivation of Section 4.3.2.

These three quantities are related to each other as follows:

$$\text{CrossEntropy}(\mathbf{p}, \mathbf{q}) = \text{Entropy}(\mathbf{p}) + \text{Kullback–LeiblerDivergence}(\mathbf{p}, \mathbf{q}). \quad (4.16)$$

Or, in more compact algebraic notation,

$$H(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}) + D_{KL}(\mathbf{p}, \mathbf{q}). \quad (4.17)$$

When applied to the neutron spectrum, the (self-)entropy $H(\mathbf{p})$ measures how evenly neutrons are distributed across all bins; while the Kullback–Leibler Divergence $D_{KL}(\mathbf{p}, \mathbf{q})$ measures how “surprised” an experimentalist should be if they expected neutron distribution $\mathbf{p}$ but the measurement shows neutron distribution $\mathbf{q}$ instead.

Notice that among the three definitions in Equation 4.13 — 4.15 above, the definition of entropy $H(\mathbf{p})$ (Equation 4.13) is the only one that is independent of distribution $\mathbf{q}$. In other words, when Equation 4.13 was used as the regularising function (such as in the case of the unfolding program MEALU [112]), it would make no reference to any *a priori* spectrum ϕ_0 , and simply try to maximise the entropy (or minimise *negative* entropy) of the normalized flux $\frac{\phi_{sol,i}}{\sum_{i=1}^n \phi_{sol,i}} = p_i$, achieving prior-less unfolding. However, this is also the equivalent of asking the algorithm to find the flattest neutron spectrum that can explain $\mathbf{N}_{\text{meas}}$, which in general is not the most realistic or physically sensible solution.

A better choice of regularising function is the cross-entropy, explained below.

Example regularising algorithm: MAXED MAXED (MAXimum Entropy Deconvolution) is a seminal unfolding program written by Marcel Reginatto in 1998, with numerous subsequent experiments that made use of it to perform unfolding. It is a type of regularising unfolding algorithm (Section 4.1.3) where the regularising function is cross-entropy. Instead of minimising a sum of cross-entropy and χ^2 , it takes a slightly different approach of minimising the cross-entropy while being constrained to a manifold of constant $\chi^2(\mathbf{N}_{\text{meas}}, \mathbf{N}_{\text{sol}}) = \Omega$, where Ω is a user-defined value. By denoting normalization of any neutron spectrum with a prime ('), i.e.

4.1 Classification of unfolding algorithms

$\phi' = \frac{1}{\sum_i \phi_{\text{sol},i}} \phi$ (See Appendix K.4), the objective of MAXED can be defined as

$$\text{MAXED: } \phi_{\text{sol}} = \underset{\phi}{\operatorname{argmin}} \left[- \sum_{i=1}^n \left(\phi'_i \log \left(\phi'_{0,i} \right) \right) \right], \text{ where } \chi^2 = \Omega, \quad (4.18)$$

$$= \underset{\phi}{\operatorname{argmin}} [H(\phi, \phi_0)], \text{ where } \chi^2 = \Omega. \quad (4.19)$$

and χ^2 is defined in Equation 3.19.

The technique of Lagrangian multipliers is used to derive the necessary equations for optimization, by splitting this single expression to be optimised into $n + m + 1$ equations (p.5 of [56]). The Lagrangian multipliers chosen in the original texts are μ and $\lambda_k, 1 \leq k \leq m$. Two clever tricks are then employed: firstly, $n + 1$ equations are eliminated by rewriting some of the terms in the remaining m equations using the relations described by the extraneous $n + 1$ equations. This changes the number of degrees of freedom of the problem from n to m (especially useful in an underdetermined case when $m < n$), and ends up with the following set of m equations:

$$N_{\text{meas},k} + \lambda_k \sigma_k^2 \sqrt{\frac{\Omega}{\sum_{l=1}^m (\lambda_l^2 \sigma_l^2)}} - \sum_{j=1}^n R_{kj} \phi'_{0,j} \exp \left\{ - \sum_{l=1}^m \lambda_l R_{lj} \right\} = 0, \quad \text{for } k = 1, \dots, m \quad (4.20)$$

where σ_k is just the standard deviation of the k^{th} measurement $N_{\text{meas},k}$ (Equation 3.11 of the original document [56]). Using the Lagrangian multipliers λ_k , the normalized solution neutron spectrum can be reconstructed by the following equation:

$$\phi'_{\text{sol}} = \phi'_0 \circ \exp(-\underline{\lambda} \underline{\mathbf{R}}), \quad (4.21)$$

where $\mathbf{a} \circ \mathbf{b}$ denotes the element-wise multiplication of vectors $\mathbf{a}$ and $\mathbf{b}$.

The set of Equations 4.20 is a problem of root-finding in m -dimensions, which is still fairly difficult to solve without plenty of computational power and advanced numerical linear algebra packages, which may not have been available in that era. The second clever trick is a trick best described as “inverse Lagrangian multiplier”: Reginatto observed that the left-hand sides of the set of m Equations 4.20 form a conservative vector field over the domain of $\underline{\lambda}$, thus he expressed this vector field as $\frac{\partial Z}{\partial \underline{\lambda}}$, where Z is given below:

$$Z = - \sum_{i=1}^n \phi'_{0,i} \exp \left\{ - \sum_{k=1}^m \lambda_k R_{ki} \right\} - \sqrt{\Omega \sum_{k=1}^m (\lambda_k^2 \sigma_k^2)} - \sum_{k=1}^m N_k \lambda_k. \quad (4.22)$$

Maximising Z in Equation 4.22 has the same effect as finding the roots of all m equations in 4.20. This recasts the problem back into an optimization problem, but in m -dimensions instead of the original n -dimensions, which is very useful in underdetermined problems where $m < n$. He then applied some off-the-shelf optimization programs to solve the problem of $\operatorname{argmax} Z$. In UMG3.3 (authored by Reginatto), the few-channel version of MAXED uses the simulated annealing algorithm (with default starting temperature=1.0 and cooling factor=0.85 at every iteration) as it is a robust optimization algorithm that is unlikely to be trapped in local maxima if any are present; and for the many-channel version of UMG3.3, L-BFGS-B was implemented for a similar rationale.

In Reginatto’s original paper of 1998, he demonstrated that using cross-entropy as the regularising function is mathematically rigorous as it is the only function that will yield a unique solution invariant under change of coordinates and also satisfy “subset independence” and “system independence” (see original text [56] for more details). Hence in Section 4.3, we will build on this argument, but taking more care when defining cross-entropy.

4.1.4 Purely iterative

The final class of algorithms in the classification tree of Figure 4.1 employs none of the aforementioned regularization techniques, hence they are purely iterative in nature. Starting at a user-provided *a priori* spectrum ϕ_0 , they begin taking steps in the downhill direction of ϕ space (or $\log(\phi)$ space) in the hopes of finding a point with a lower loss value, where the loss function may be χ^2 as defined in Equation 3.19, or a modified version of it (See Section 4.1.4 paragraph **Example purely iterative algorithm 2: GRAVEL**). Eventually, when the loss value is lower than a user-set threshold, the algorithm terminates, and returns the final iteration of ϕ as the ϕ_{sol} . These algorithms simply sidestep the problem of underdetermination by making arbitrary decisions that determine the direction and size of next steps, typically with no consideration of what the appropriate step size is, or what assumptions are made by descending in the chosen direction. Two of the most famous examples of purely iterative codes are SAND-II and GRAVEL. While these algorithms are effective at generating physically plausible spectra, they lack rigorous mathematical arguments over why the ϕ_{sol} must be the point closest to ϕ_{true} when the response matrix is underdetermined. In plain words, this class of algorithms finds **a** solution spectrum, but with no guarantee that it is **the** solution most similar to the original spectrum.

4.1 Classification of unfolding algorithms

Regardless, these types of algorithms had once been very popular due to their memoryless nature, which means each step in its path of descent in the loss function landscape is decided purely on the current values of ϕ and $\mathbf{N}_{\text{meas}}$, with no explicit reference to the *a priori* or any previous steps. This, combined with the simplicity of the formula required to calculate the scale factors at each step, made it a very popular style of unfolding algorithm during the early days of spectra unfolding when computational resources, especially memory, were limited.

One of the disadvantages of purely iterative algorithms is that they do not always take the most efficient route to the point of the smallest χ^2 value - which may not be a problem in the case of fully- and overdetermined unfolding. However, as shown in Figure 3.3, the solutions to underdetermined unfolding problems are not unique, so algorithms that take a large detour from ϕ_0 to the manifold of minimum loss value may end up returning a ϕ_{sol} that is nowhere near ϕ_{true} even though it has a low χ^2 value. Sometimes, these programs may crash due to having runtimes too long to be supported by the computer, terminating before reaching the user-set threshold of χ^2 . Another disadvantage is that due to a lack of proper loss function used when deriving these algorithms, uncertainty propagation can currently only be achieved by a Monte Carlo approach of repeating the optimization with slightly perturbed initial $\mathbf{N}_{\text{meas}}$, which is very computational-resource intensive. In fusion neutron spectrum unfolding, each optimization may take up to several tens of thousands of steps, and error propagation would require repeating this at least several hundred times.

Example purely iterative algorithm 1: SAND-II SAND-II was written by physicists in the U.S. Air Force Weapons Laboratory during the 1960s specifically for neutron spectrum unfolding with activation foils [119]. It applies the following procedure to iteratively update the solution neutron spectrum (starting at $\phi_{\text{sol}} = \phi_0$ initially).

At the first iteration, SAND-II sets the initial guess solution as the *a priori* spectrum $\phi_{\text{sol}} = \phi_0$. Then, $\mathbf{N}_{\text{sol}}$ for this iteration is calculated as $\mathbf{N}_{\text{sol}} = \underline{\mathbf{R}}\phi_{\text{sol}}$ according to Equation 3.8. After that, a weight matrix $\underline{\mathbf{W}}$ is created using $\mathbf{N}_{\text{sol}}$. The detailed list of steps generating $\underline{\mathbf{W}}$ can be found in Appendix E. This step applies a user-chosen degree of smoothing, N_s , to suppress “spurious structures” in the unfolded spectra. As the author has frankly expressed in [54], the choice of N_s is fairly arbitrary. There are no standardised ways of choosing N_s and it would be up to the user of SAND-II to experiment with different values of N_s , leaving much of the unfolding procedure to human interpretation and making it irreproducible.

4.1 Classification of unfolding algorithms

Using this $\underline{\underline{\mathbf{W}}}$ matrix, ϕ_{sol} can be updated as follows:

$$\lambda = \frac{1}{\mathbf{1}^m \cdot \underline{\underline{\mathbf{W}}}}, \quad (4.23)$$

$$\Delta = \left(\log(\mathbf{N}_{\text{meas}}) - \log(\mathbf{N}_{\text{sol}}) \right) \cdot \underline{\underline{\mathbf{W}}} \circ \lambda, \quad (4.24)$$

$$\phi_{\text{sol}} \leftarrow \phi_{\text{sol}} \circ \exp(\Delta), \quad (4.25)$$

where $\mathbf{1}^m$ = an m -dimensional vector of all 1's. We then repeat this set of steps, beginning at the weight matrix calculation ($\underline{\underline{\mathbf{W}}}$), if the $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})$ is still above the threshold that the user defined; otherwise it terminates, returning the last iteration of ϕ_{sol} as the solution spectrum.

Note that the last step involves the multiplication of a positive vector ϕ_{sol} with another positive vector $\exp(\Delta)$. The intention behind this multiplicative operation (as opposed to an additive operation) is to ensure that the next iteration of ϕ_{sol} will remain a positive vector ($\mathbb{R}^+ \times \mathbb{R}^+ \mapsto \mathbb{R}^+$), so this avoids the unphysical results of negative fluxes.

According to [120], repeated iterations of SAND-II have the serendipitous effect of minimising the “lognormal least squares” when the system is overdetermined, giving the same result as minimising this loss function:

$$\text{loss}_{\text{SAND-II}} = \left(\log(\mathbf{N}_{\text{sol}}) - \log(\mathbf{N}) \right) \cdot \underline{\underline{\mathbf{I}}} \left(\log(\mathbf{N}_{\text{sol}}) - \log(\mathbf{N}) \right), \quad (4.26)$$

where $\underline{\underline{\mathbf{I}}}$ is the identity matrix, included to match the form of Equation 3.19 while making it clear that this loss function is not weighted by the uncertainty of the measurements $\mathbf{N}_{\text{meas}}$. This quantity is essentially an unweighted χ^2 for the difference between $\mathbf{N}_{\text{meas}}$ and $\mathbf{N}_{\text{sol}}$ in log space. However, there is no conclusive literature on whether SAND-II will converge onto the minimiser of the same loss function in an underdetermined system, which is more relevant in fusion neutron spectrum unfolding.

Due to the unusual loss function, and its lack of clear mathematical motivation or reasoning in its semi-empirical derivation, SAND-II is not deemed a suitable algorithm for unfolding fusion neutron spectra accurately. It is also not guaranteed to converge, as it has no built-in functionality to reduce the step size when near the optimum point. This means that it is possible to constantly overshoot the optimum point and end up teetering on either side of the optimum without any further decrement in loss value/ χ^2 .

4.1 Classification of unfolding algorithms

Example purely iterative algorithm 2: GRAVEL GRAVEL was written by Manfred Matzke in 1994 as an improved version of SAND-II. It accounts for the uncertainty in $\mathbf{N}_{\text{meas}}$ ($\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}}$) by modifying one of the steps in SAND-II, and was released as part of the unfolding program HEPROW [69].

To aid with the derivation in linear algebra terms, we introduce a new quantity $\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}}$, which denotes the inverse of the covariance matrix associated with $\mathbf{N}_{\text{meas}}$. In the HEPROW implementation of GRAVEL, this covariance matrix is assumed to be diagonal, so that $(\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}})_{jk} = \frac{\delta_{jk}}{\text{variance on } (\mathbf{N}_{\text{meas}})_k}$.

$$\mathbf{N}_{\text{sol}} = \underline{\underline{\mathbf{R}}} \phi_{\text{sol}}, \quad (4.27)$$

$$\underline{\underline{\mathbf{W}}} = \frac{1}{\mathbf{N}_{\text{meas}}} \circ \underline{\underline{\mathbf{R}}} \circ \phi_{\text{sol}}, \quad (4.28)$$

$$\boldsymbol{\lambda} = \frac{1}{\mathbf{N}_{\text{meas}} \circ \mathbf{N}_{\text{meas}} \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \cdot \underline{\underline{\mathbf{W}}}}, \quad (4.29)$$

$$\Delta = \left(\log(\mathbf{N}_{\text{meas}}) - \log(\mathbf{N}_{\text{sol}}) \right) \cdot \left(\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \mathbf{N}_{\text{meas}} \circ \mathbf{N}_{\text{meas}} \circ \underline{\underline{\mathbf{W}}} \circ \boldsymbol{\lambda} \right), \quad (4.30)$$

$$\phi_{\text{sol}} \leftarrow \phi_{\text{sol}} \circ \exp(\Delta), \quad (4.31)$$

Similarly to SAND-II, the multiplication with a positive vector in the last (ϕ_{sol} update) step ensures neutron fluxes in all bins remain positive. It is unclear whether an analogue to Equation 4.26 exists for GRAVEL, see Appendix F for more details. It is a distinct possibility that repeated iteration of GRAVEL does not yield any result equivalent to the minimization of a loss function.

As GRAVEL was not designed for the extreme conditions of fusion neutron spectrum unfolding with activation foils, it also tends to perform poorly when solving unfolding problems with these levels of underdetermination ($m \approx 10$, $n \approx 175$, which is much worse than, say, fission neutron spectra where $E_{\text{max}} \approx 2$ MeV, requiring fewer bins and hence having a smaller n , or scintillator detectors, where $m \approx n$ as number of bins in the response matrix can be arbitrarily increased). The lack of smoothing conditions imposed in $\underline{\underline{\mathbf{W}}}$ to suppress “spurious structures” allows for the growth of unphysical artefacts, such as peaks corresponding to the resonance peaks of the foils. These artefacts get amplified more when the number of iterations increases (due to its memoryless nature). This is especially problematic considering that GRAVEL suffers from the same issue as SAND-II of not being able to change the step size when approaching a point of minimum χ^2 , so it often fails to converge within a small number of iterations. To the

author’s knowledge, there is no literature available on the convergence condition of GRAVEL, meaning we cannot predict when GRAVEL will/will not converge⁶.

Before the modern variant of the QR singular value decomposition algorithm was popularised by Golub et al. in LAPACK [121], an iterative program such as SAND-II that doesn’t involve any matrix inversion, and hence takes up very little memory, is a perfectly reasonable choice of unfolding algorithm, despite its weaker mathematical foundation [122]. And GRAVEL, which is a modified version of SAND-II [83] did not attempt to rebuild the mathematical foundation or make use of the increased computational power to invert the Hessian matrix to find the appropriate step size and step direction.

But given the improvements in computer architectures and programming languages [123], where simpler linear algebra computation syntax reduces the rate of error and amount of work required to implement and test new unfolding algorithms, it is easier than ever to leverage the improved computing capabilities to implement a more mathematically sound unfolding algorithms than these old iterative schemes offer, effectively rendering this class of algorithms obsolete.

4.2 Assessing the correctness and suitability of unfolding algorithms

Due to the ill-posed nature of the general unfolding Equation 3.6, there exists no single uniquely “correct” solution to the unfolding problem when the problem is underdetermined, without a substantial amount of subjective opinion inputted in the form of *a priori* information provided by the experimentalist performing the unfolding. This poses a challenge to anyone interested in answering the question “Which algorithm can most accurately reproduce the correct solution” especially in the field of fusion neutronics, where *a priori* information and understanding of fusion neutron spectra is in scarce supply.

A handful of studies had highlighted the issue that there exists no fair method of comparing the performance of unfolding algorithms against each other over a large range of unfolding problems, and expressed an ambition to develop a general benchmarking framework to answer the question of “which unfolding code is the best”. However, all of them quickly lost momentum as they either did not generate any citations from any parties interested in developing the framework [124], or simply generated no citations after the workshop they were submitted

⁶It is possible that we will never know whether GRAVEL will reliably converge due to the Halting Problem. However, to prove whether this is or is not the case is far beyond the scope of this PhD.

for [125]. Both projects were possibly stopped by the same obstacle, i.e. they could not proceed to develop their framework further without including domain-specific *a priori* information, destroying the generality of their framework in the process.

Regardless of the reason, the result is that today there exists no mature framework of benchmarking neutron spectrum unfolding programs to determine their performances. Consequently, there is also no comprehensive literature review that suggests which of the unfolding algorithms is best for fusion neutron spectra unfolding or similar domains of problems, where the type of *a priori* spectra that will be used as input is still largely unknown, and the degree of underdetermination will be the highest ever seen in all fields of unfolding.

Fortunately, this does not spell the end of this thesis' attempt to find a suitable unfolding algorithm for fusion. We approach the problem of finding a suitable unfolding algorithm in the specific domain of fusion neutron spectrum unfolding problem by considering the following four aspects: reliability, choice of loss function, independence from the *a priori*, and independence from the response matrix.

Not all algorithms will efficiently converge when presented with a difficult unfolding problem, and those that do converge are not guaranteed to produce a realistic spectrum. For example, in a highly constrained problem starting with an *a priori* that deviates from ϕ_{true} a lot and with large (> 100) values of m and n , as is the case for fusion neutron spectra unfolding, simpler algorithms such as gradient-descent or GRAVEL will take up to several tens of thousands of iterations to converge, or fail to converge altogether. Meanwhile, directly taking the inverse (Equation 4.41) will converge to the solution in a single step, but will almost certainly lead to negative fluxes or unphysical dips in the neutron spectrum when the *a priori* spectrum deviates from the true neutron spectrum even just slightly. Therefore care has to be taken to choose unfolding algorithms that converge to in a reasonable number of iterations without causing unphysical features to appear in the solution spectrum.

Not all loss functions are created equal. While all algorithms ultimately reduce the deviation of $\mathbf{N}_{\text{sol}}$ from $\mathbf{N}_{\text{meas}}$, the choice of loss function determines whether the solution produced points to a physically realistic spectrum or not. Some algorithms choose loss functions that minimise the roughness of the neutron spectrum, such as those that use Fisher information as the regularising function. In the domain of fusion neutronics, where resonance absorption or lack of shielding from the 14 MeV neutron source might lead to sharp dips or peaks in the neutron spectrum, such algorithms that promote the smoothness of the spectrum would be a bad choice as it would bulldoze over these features. Substantial consideration has been placed into this and is elaborated further in Section 4.2.2.

At the same time, an unfolding algorithm that does not depend strongly on the *a priori* is preferred. For example, STAY'SL simply incorporates the *a priori* spectrum's information as shown in Equation 3.21. This places a large and un-tuneable weight on the importance of the *a priori*, especially when the degree of underdetermination is high, which is not desirable when the expected neutron spectrum is not well known or well characterised. On the other hand, algorithms that require no *a priori* spectrum input, and instead maximise smoothness in its loss function in an adjustable manner can avoid this pitfall, but with the aforementioned problem of not respecting the limitations of fusion neutron spectra. Therefore we have to strike a balance between the two extremes and find algorithms that place a small (preferably tunable) but non-zero weight on the *a priori* spectrum.

When the condition number of $\underline{\mathbf{R}}$ is high or the *a priori* is not similar to ϕ_{true} , which is usually the case for activation foils [126], algorithms that take in so little *a priori* information will give sub-par results. Namely, features in the response matrix, such as resonance peaks of the reaction cross section, will appear on the unfolded solution. This phenomenon is frequently seen in SAND-II [54] and GRAVEL. Ideally, we wish to be able to repeat our measurements with a different detector or a different set of foils which has different response characteristics, and still be able to obtain the same solutions; but algorithms such as GRAVEL and SAND-II are unable to achieve this. Therefore we would ideally avoid algorithms that amplify artefacts in the response matrix indefinitely with every iteration, suggesting that purely iterative algorithms (which are memoryless) should be avoided.

To this end, several algorithms that sidestep all of the problems mentioned in this section have been developed. They are designed to take in just enough *a priori* information to decouple the solution spectrum from features in the response matrix without relying too heavily on the *a priori* spectrum, and will always converge quickly. These algorithms all use cross-entropy as the foundation of their loss functions and are designed to allow uncertainty to propagate analytically without the use of Monte Carlo methods.

4.2.1 The model of ϕ_{true} , ϕ_{sol} , and ϕ_0

If we reduce neutron spectrum unfolding to its essence, it is concerned with the three spectra as shown in Figure 4.3, whose relationship with each other is explained in Table 4.1.




Figure 4.3: An illustrative diagram to distinguish the relationships between various ϕ from one another. ϕ_{true} is unknown in a real experiment; but may be known in a synthetic dataset (e.g. Section 4.4).

Table 4.1: Table attached to Figure 4.3 explaining the three distance metrics D .

Deviation	Description
$D_{\text{sol},0}$	A good unfolding algorithm shall minimise this.
$D_{\text{sol,true}}$	The end goal of unfolding is to minimise this.
$D_{0,\text{true}}$	A good <i>a priori</i> selector shall minimise this.

When viewed in this way, it is immediately clear what the role of a neutron spectrum unfolding algorithm is: it shall reduce $D_{\text{sol},0}$; meanwhile, it is the responsibility of the unfolding program user to select an appropriate *a priori* spectrum ϕ_0 , either by selecting an existing neutron spectrum and manually modifying it or by using an *a priori* selection program to do so (such as [127]). Without knowing what kind of *a priori* the user will choose and how large $D_{0,\text{true}}$ would typically be, we cannot estimate how good the unfolded spectrum would be (i.e. how small $D_{\text{sol,true}}$ would be).

This model helps us understand why it is difficult to benchmark the performance of these algorithms on fusion neutron spectrum (as of now) — $D_{\text{sol,true}}$ has two components: $D_{\text{sol},0}$ and $D_{0,\text{true}}$. The former component is easy to minimise with the help of algorithms developed in Section 4.3. However, in fusion neutronics, due to the large degree of underdetermination, the latter component dominates, and due to the lack of literature in the list of ϕ_{true} that may arise and the list of possible ϕ_0 that may be used, this is not possible. One can try and generate synthetic data that matches fusion neutron spectra, but that in and of itself is a fairly complex challenge as it involves plasma physics and complex tokamak models.

Therefore, due to the many obstacles in the way of obtaining realistic data for benchmarking, the algorithms produced in this thesis will only be demonstrated on a small set of data, without a comprehensive benchmarking process. To compensate for the weaker empirical evidence, this

thesis will focus on making the mathematical theories rigorous to prove that they are sound, and are fit for unfolding fusion neutron spectra.

This model also highlights the need for establishing what D is — for the moment, D is only a placeholder metric measuring the deviation from one neutron spectrum to another⁷. In Section 4.2.2, we choose what metric D represents, and from that, we can properly develop algorithms to minimise D . These algorithms are then, by definition, the most effective algorithm that exists to minimise $D_{\text{sol},0}$. In other words, if we are only concerned with $D_{\text{sol},0}$, these algorithms are provably the best ones to use.

4.2.2 Choosing the metric

Reginatto put forth a very convincing argument (Section 4.1.3.3 paragraph **MAXED**) that entropy is the most “correct” metric (being that it is self-consistent, subset-independent and system-independent) to measure the deviation of one neutron spectrum from another. Therefore the following Section 4.3 focuses the attention on using cross-entropy as the regularising function.

Stepping back from the loss function of MAXED in Equation 4.18 and re-examining the definition of cross-entropy in Equation 4.14, we can see that it is fundamentally asymmetric, i.e. in general $H(\mathbf{p}, \mathbf{q}) \neq H(\mathbf{q}, \mathbf{p})$. It is then important to consider whether we should set $(\mathbf{p}, \mathbf{q}) = (\phi'_{\text{sol}}, \phi'_0)$ as MAXED did, or as $(\phi'_0, \phi'_{\text{sol}})$, because the solution ϕ_{sol} obtained by minimising the former would be different to the solution obtained by minimising the latter.

Looking at Equation 4.17 and Section 4.1.3.3, to use the MAXED’s definition $\phi_{\text{sol}} = \underset{\phi}{\operatorname{argmin}} H(\phi, \phi_0)$ means that we are simultaneously minimising the deviation of ϕ_{sol} from ϕ_0 while trying to choose the **least flat** neutron spectrum as ϕ_{sol} . Subjectively speaking, this is problematic and has room for improvement, as there aren’t any arguments for why a less flat/more uneven neutron spectrum makes for a more realistic neutron spectrum. If we choose a different quantity to minimise, $\phi_{\text{sol}} = \underset{\phi}{\operatorname{argmin}} H(\phi'_0, \phi')$, we would then be only concerned with minimising the deviation of ϕ_{sol} from ϕ_0 , without applying an undue driving force to make the spectrum either smoother or spikier. This opinion is supported by a careful investigation into the definition of cross-entropy [118], which shows that the first distribution $\mathbf{p}$ is defined as the “prior” distribution, thus also showing that $H(\phi'_0, \phi'_{\text{sol}})$ is a more natural choice of loss function.

Therefore in Sections 4.3.2 and 4.3.3, this function will be applied using the definition cross-entropy $= H(\phi'_0, \phi')$.

⁷Note that in this thesis, when speaking of a “metric”, we do not require such metric to be symmetric, i.e. in general, $\text{metric}(a, b) \neq \text{metric}(b, a)$.

4.3 New algorithms developed

This section forms one of the major contributions to knowledge in this PhD, as new algorithms are created with the following features:

- Quick convergence — Using Newton’s method and sped up by line search, these algorithms all rapidly approach the global optimum in their loss function, in very few iterations⁸. This means the programs will not crash due to non-convergence.
- Guarantee that the converged solution is the optimum — the unfolding algorithms enforce checks that ensure the converged solutions are not local maxima or saddle points, but true minimiser solutions to the minimization problems stated in their premise.
- Protection against negative fluxes — as a corollary from the previous two points, in all iterations of these algorithms, the neutron spectrum will always have positive fluxes, as their implementations (utilising line search) hedge the step size appropriately to avoid negative fluxes.
- Rigorous mathematical derivation — strong logical reasoning at each step, valid for both underdetermined and overdetermined cases.
- Analytical uncertainty propagation — all of these algorithms have explicit formulae for uncertainty propagation, such that computationally expensive and untransparent Monte Carlo methods would not be required. This also allows the user to have a good glimpse of the expected uncertainties on the solution neutron spectrum before the experiment, which would be useful in Chapter 5.
- Versatility — the algorithms can be applied to unfold any particle spectra measured with any instrument with a linear response, as no assumptions were made to restrict them to apply to only fusion neutron spectra with activation foils.

Previous algorithms usually satisfy none or only a few of these requirements, while the list of algorithms developed in the following section satisfies all of the above features. The keen reader can see Appendix B to D for detailed mathematical derivations for each of these new algorithms, including the mathematical proofs mentioned above. All of the algorithms named

⁸All of these algorithms converged within 100 iterations when tested by synthetic data of $n = 175, m = 16$ (Section 4.4.3), and fewer than 200 iterations when tested by experimental data of $n = 315, m = 1306$ (Section 4.5). Of course, this number would change depending on the values of (m, n) and the deviation of ϕ_0 from ϕ_{true} . Regardless, these values are much more promising than existing implementation(s) of MAXED, GRAVEL and SAND-II.

4.3 New algorithms developed

Table 4.2: Comparison of all cross-entropy-based algorithms that have been/will be discussed in this chapter.

Algorithm name	loss function	optimiser	Respects the definition of entropy?	Newly developed?	Allow target χ^2 to be set to 0
MAXED	$H(\phi_{\text{sol}}, \phi_0)$, fixed at $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) = \Omega$	Simulated Annealing/ L-BFGS-B	No ⁹	Yes	Yes
IMAXED	$H(\phi'_{\text{sol}}, \phi'_0)$, fixed at $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) = \Omega$	Newton’s method + Wolfe condition	No	Yes	Yes
AMAXED	$H(\phi'_0, \phi'_{\text{sol}})$, fixed at $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) = \Omega$	Newton’s method + Wolfe condition	Yes	Yes	Yes
AMAXED- Regularization	$(\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) + \tau H(\phi'_0, \phi'_{\text{sol}}))$	Line search + {Newton’s method + Wolfe condition}	Yes	Yes	No ¹⁰

in this section, along with GRAVEL, SAND-II, and MAXED, have been implemented into a Python package named the “unfoldingsuite” so that their performances can be compared fairly across the board. One of the useful features of this unfoldingsuite implementation is that, unlike previous unfolding programs that allow only for variances of $\mathbf{N}_{\text{meas}}$ to be entered [68, 128], the user can enter the full covariance matrix, such that the problem of overestimating the uncertainty mentioned in Section 3.2.1.1 is resolved.

4.3.1 IMAXED

Several possible optimizations were observed in the UMG3.3 MAXED and improved upon when making the IMAXED (Improved MAXED) package in the unfoldingsuite.

The biggest improvement is the use of an optimization algorithm to optimise Z — a simpler Newton’s method, sped up by line search, was used to converge to the location of maximum Z faster. A rigorous derivation and proof of the uniqueness of the global maximum is available in Appendix B. Another improvement is the inclusion of covariance information between different measurements N_k , which was not available in the UMG3.3 implementation of MAXED.

These changes led to more reliable convergence to the solution ϕ_{sol} , as the UMG3.3 version of MAXED sometimes fails to converge to the optimal solution (as the simulated annealing algorithm terminates itself when an arbitrary predetermined number of iterations has passed,

⁹The ϕ_{sol} and ϕ_0 used in cross-entropy H definition in MAXED and IMAXED are not normalized.

¹⁰AMAXED-Regularization cannot converge to exactly $\chi^2 = 0$, but it can achieve a result arbitrarily close to it.

4.3 New algorithms developed

without verifying whether the global maximum has been reached or not). IMAXED will always converge and only terminate when the global maximum of Z is found (verified by the second derivatives of Z w.r.t. ϕ all being negative), as long as $\Omega \geq$ the minimum achievable χ^2 for that problem.

The objective function of IMAXED is identical to MAXED's, shown in Equation 4.19. The only difference from MAXED lies in the implementation of its optimiser, detailed in Appendix B.

4.3.2 AMAXED

As explained in Section 4.2.2, a better usage of cross-entropy is given by $H(\phi'_0, \phi'_{\text{sol}})$. By minimising this quantity while holding χ^2 constant, AMAXED (Alternative definition MAXED) was created:

$$\text{AMAXED: } \phi_{\text{sol}} = \underset{\phi}{\operatorname{argmin}} \left[- \sum_{i=1}^n \left(\phi'_{0,i} \log(\phi'_i) \right) \right], \text{ where } \chi^2 = \Omega, \quad (4.32)$$

$$= \underset{\phi}{\operatorname{argmin}} [H(\phi'_0, \phi')] \quad , \text{ where } \chi^2 = \Omega, \quad (4.33)$$

Using Equation 4.17, this can be converted to

$$= \underset{\phi}{\operatorname{argmin}} [D_{KL}(\phi'_0, \phi')] \quad , \text{ where } \chi^2 = \Omega. \quad (4.34)$$

Between Equation 4.33 and Equation 4.34, the two quantities within the argmin bracket differ by only a constant $\phi'_0 \circ \mathbf{log}(\phi'_0)$, thus mathematically the argmin of both terms should be the same. However in practice, the latter expression has a smaller numerical value, allowing a more faithful representation by floats and therefore a more accurate solution to be produced by alleviating the effects of floating point errors.

Compared to the MAXED definition of loss function in Equation 4.19, the new definition of loss function by AMAXED (Equation 4.34) has several advantages. First is that, as shown in Section 4.1.3.3, $\lim_{\phi_{\text{sol},i} \rightarrow 0^+} D_{KL}(\phi'_0, \phi'_{\text{sol}}) = \infty$ and $\lim_{\phi_{\text{sol},i} \rightarrow 0^+} \frac{\partial D_{KL}(\phi'_0, \phi'_{\text{sol}})}{\partial \phi_{\text{sol},i}} = \infty$, thereby penalizing the optimiser for trying to cross into the negative fluxes region, preventing unphysical negative fluxes from appearing in the solution ϕ_{sol} . This feature is not present in MAXED as $\mathbf{p}$ in $H(\mathbf{p}, \mathbf{q})$ is allowed to become negative without blowing up the loss value. The other advantage is the ease of uncertainty propagation: the curvature (Hessian matrix) of the loss function can be inverted directly to obtain the covariance matrix of ϕ_{sol} . This expression of the Hessian matrix, along with details about AMAXED's derivation, and how we guarantee the convergence

of AMAXED, can be found in Appendix C. Similar to IMAXED, it will always converge to a viable solution when $\Omega \geq$ the minimum achievable χ^2 for that problem.

4.3.3 AMAXED-Regularization

In both AMAXED and IMAXED, the solution ϕ_{sol} is found at a fixed, user-set value of χ^2 . Appendix A.2 argues why this is an unnatural way to weight the importance of χ^2 against the importance of the regularising function, making it unhelpful for first-time users who would not know what χ^2 to choose, as well as failing to provide any intuition as to how far the ϕ_0 is from an acceptable solution ϕ_{sol} .

Instead, a generalised Tikhonov regularization algorithm is proposed (as suggested by Appendix A.3), where the cross-entropy definition is the same as the definition used in AMAXED:

$$\text{AMAXED-Regularization: loss} = (\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) \cdot \underline{\mathbf{S}}_{\mathbf{N}_{\text{meas}}}(\mathbf{N}_{\text{meas}} - \mathbf{N}_{\text{sol}}) + \tau \sum_{i=1}^n \phi'_{0,i} \log \left(\frac{\phi'_i}{\phi_{0,i}} \right) \quad (4.35)$$

$$= \chi^2(\mathbf{N}_{\text{meas}}, \mathbf{N}_{\text{sol}}) + \tau \cdot D_{KL}(\phi'_0, \phi'). \quad (4.36)$$

As mentioned in Equation 4.5, τ is a positive real parameter, i.e. confined to the interval $(0, \infty)$. The user can experiment with different values of τ to see which one gives the most physically realistic solution spectrum ϕ_{sol} . This provides a much more intuitive parameter to tune the ratio of the weight of the regularising function against the weight of χ^2 .

In addition to running it in the `fix_tau` mode, the `unfoldingsuite` implementation of AMAXED-Regularization can also be run in the `fix_chi2` mode, where the τ is varied until the final solution ϕ_{sol} has a $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})$ equal to the user-provided target χ^2 value, for ease of comparing its performance against the other algorithms. Of course, this only works if the target $\chi^2 \geq$ the minimum achievable χ^2 for that problem, otherwise the program will fail to converge as no values of τ would be able to give a solution ϕ_{sol} that satisfies $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) = \text{target } \chi^2$.

This algorithm can also play the role of selecting priors: If multiple *a priori* spectra are available, it can be used to measure which one is the best at fitting to the data at hand. This is done by keeping τ , $\mathbf{N}_{\text{meas}}$ and $\underline{\text{cov}}(\mathbf{N}_{\text{meas}})$ constant across multiple runs, changing only the ϕ_0 , the user can then compare which of those runs gives a lower χ^2 , and hence is a better fit for the data $\mathbf{N}_{\text{meas}}$.

AMAXED-Regularization shares some of the same advantages as AMAXED, being able to propagate uncertainty without using Monte Carlo methods, having no other stationary

points than the global minimum, and the fact that negative fluxes do not appear. The detailed derivation that leads to AMAXED-Regularization, the proofs of the uniqueness of the global minimum as the only stationary point, as well as the uncertainty propagation equations, are all given in Appendix D.

4.3.4 Pseudo-Inverse

This algorithm is very simple and hence does not perform well on datasets of high complexity (i.e. large degree of underdetermination). However, as it is used as part of Chapter 5 and forms a foundational part of the unfoldingsuite [129], it is included in this thesis for completeness.

The Pseudo-Inverse unfolding algorithm is a special purely iterative algorithm, one which takes only a single step to reach the point of minimum χ^2 . When faced with Equation 3.8, the simplest action required to reconstruct the neutron spectrum is

$$\phi_{\text{sol}} = \underline{\underline{\mathbf{R}}}^{-1} \mathbf{N}_{\text{meas}} . \quad (4.37)$$

However, this method does not work for any non-square matrix $\underline{\underline{\mathbf{R}}}$, as $\underline{\underline{\mathbf{R}}}^{-1}$ only exists if $\underline{\underline{\mathbf{R}}}$ is square and non-singular. But an alternative, called a pseudo-inverse $\underline{\underline{\mathbf{R}}}^\dagger$, exists for both underdetermined and overdetermined cases, where the matrices are rectangular ($m < n$ and $m > n$ respectively). The brief explanation of how pseudo-inverses are calculated is that Singular Value Decomposition (SVD) is performed on $\underline{\underline{\mathbf{R}}}$ to find out all its nonsingular directions of transformation, and then a matrix that reverses those transformations is constructed. The resultant matrix has the property that $\underline{\underline{\mathbf{R}}} \underline{\underline{\mathbf{R}}}^\dagger \underline{\underline{\mathbf{R}}} = \underline{\underline{\mathbf{R}}}$. While $\underline{\underline{\mathbf{R}}}$ has shape $m \times n$, $\underline{\underline{\mathbf{R}}}^\dagger$ has shape $n \times m$. Being able to calculate $\underline{\underline{\mathbf{R}}}^\dagger$ allows the unfolding to be completed in a single step.

It is tempting to substitute $\underline{\underline{\mathbf{R}}}^\dagger$ in the place of $\underline{\underline{\mathbf{R}}}^{-1}$ in Equation 4.37. However, this is not advisable, and the reason is that this makes no use of the *a priori* information provided. Instead, to account for *a priori* information and uncertainties on $\mathbf{N}_{\text{meas}}$, the following derivation is required:

$$\phi_{\min \chi^2} = \underset{\phi}{\operatorname{argmin}} \left(\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}} \phi \right) \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \left(\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}} \phi \right) \quad (4.38)$$

$$\begin{aligned} \left. \frac{d\chi^2}{d\phi} \right|_{\phi_{\min \chi^2}} &= \mathbf{0} \\ \mathbf{0} &= -2 \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \left(\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}} \phi_{\min \chi^2} \right) \\ \mathbf{0} &= \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \left(\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}} \phi_{\min \chi^2} \right) . \end{aligned} \quad (4.39)$$

4.3 New algorithms developed

Separating $\phi_{\min \chi^2}$ into two components: in such a way that the *a priori* information ϕ_0 is preserved:

$$\phi_{\min \chi^2} = \phi_0 + \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \right)^\dagger \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \phi_0 \right), \quad (4.40)$$

and using the shorthand $\underline{\underline{\mathbf{S}}}_\phi = \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}}$, we have

$$\phi_{\min \chi^2} = \phi_0 + \underline{\underline{\mathbf{S}}}_\phi^\dagger \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{S}}}_\phi \phi_0 \right). \quad (4.41)$$

Once a point of minimum χ^2 is identified as $\phi_{\min \chi^2}$ on Equation 4.41, then a whole family of solutions ϕ_{sol} can be identified:

$$\phi_{\text{sol}} = (1 - \beta) \phi_{\min \chi^2} + \beta \phi_0 \quad (4.42)$$

The Pseudo-Inverse algorithm can then take a second step to backtrack from this point of minimum χ^2 to any other user-defined values of $\chi^2 = \Omega$ by choosing $\beta > 0$, as long as $\Omega > \chi^2(\underline{\underline{\mathbf{R}}} \phi_{\min \chi^2}, \underline{\underline{\mathbf{R}}})$. Taking advantage of the quadratic nature of the χ^2 landscape, β is given by:

$$\beta = \sqrt{\frac{\Omega - \chi^2(\underline{\underline{\mathbf{R}}} \phi_{\min \chi^2}, \mathbf{N}_{\text{meas}})}{\chi^2(\underline{\underline{\mathbf{R}}} \phi_0, \mathbf{N}_{\text{meas}}) - \chi^2(\underline{\underline{\mathbf{R}}} \phi_{\min \chi^2}, \mathbf{N}_{\text{meas}})}} \quad (4.43)$$

The behaviour of this algorithm is best explained visually in Figure 4.4: Instead of inverting the entire measurement $\mathbf{N}_{\text{meas}}$ to reach the manifold of minimum χ^2 directly, we begin at the *a priori* spectrum, and only invert the difference between $\underline{\underline{\mathbf{R}}} \phi_0$ (the would-have-been reaction rate if $\phi_{\text{true}} = \phi_0$) and the actual response vector $\mathbf{N}_{\text{meas}}$. This allows us to take the smallest step (as measured by the L_2 metric) necessary to bring us onto the minimum χ^2 manifold, minimising our deviation from ϕ_0 while minimising χ^2 at the same time. This step would appear as a step perpendicular to the minimum χ^2 manifold (Figure 4.4).

If we allow ourselves to abuse mathematical terminology briefly and name the directions in ϕ space that don't affect the χ^2 as singular directions of $\underline{\underline{\mathbf{R}}}$, then we can describe how ϕ_{sol} can be broken into two parts: the nonsingular components of ϕ_{sol} are equal to $\underline{\underline{\mathbf{R}}}^\dagger \mathbf{N}_{\text{meas}}$, while its singular components are equal to the singular components of ϕ_0 .

This algorithm has three different behaviours depending on whether $m < n$, $m = n$, or $m > n$. In an underdetermined system, $\underline{\underline{\mathbf{R}}}^{T\dagger} \underline{\underline{\mathbf{R}}}^T = \underline{\underline{\mathbf{I}}}$ and $\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}}^\dagger \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} = \underline{\underline{\mathbf{I}}}$, transforming Equation 4.41




Figure 4.4: Several randomly chosen *a priori* spectra are unfolded by the pseudo-inverse unfolding algorithm in an underdetermined 2-bin system down to $\Omega = \text{target } \chi^2 = 0$. The flux values on the ϕ_1 and ϕ_2 axes are deliberately left ambiguous to emphasize that this behaviour is scale-invariant and translationally-invariant, i.e. the same behaviour is observed regardless of the positivity/negativity of fluxes.

4.3 New algorithms developed

to:

$$\begin{aligned}
\phi_{\min \chi^2} &\equiv \phi_0 + \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \underline{\underline{\mathbf{R}}} \right)^\dagger \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \underline{\underline{\mathbf{R}}} \phi_0 \right) \\
&\equiv \phi_0 + \left(\underline{\underline{\mathbf{R}}}^\dagger \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}^\dagger \underline{\underline{\mathbf{R}}}^{T\dagger} \right) \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \underline{\underline{\mathbf{R}}} \phi_0 \right) \\
&\equiv \phi_0 + \underline{\underline{\mathbf{R}}}^\dagger \left(\mathbf{N}_{\text{meas}} - \underline{\underline{\mathbf{R}}} \phi_0 \right) \quad \text{when } m < n. \quad (4.44)
\end{aligned}$$

In a fully-determined system the last term $-\underline{\underline{\mathbf{S}}}_{\phi=\phi}^\dagger \underline{\underline{\mathbf{S}}}_{\phi=\phi} \phi_0 = -\underline{\underline{\mathbf{I}}} \phi_0 = -\phi_0$ also cancels out with the first term ϕ_0 , and the previous relations $\underline{\underline{\mathbf{R}}}^{T\dagger} \underline{\underline{\mathbf{R}}}^T = \underline{\underline{\mathbf{I}}}$, $\underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}^\dagger \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} = \underline{\underline{\mathbf{I}}}$ still applies, making the Equation 4.41 equivalent to:

$$\begin{aligned}
\phi_{\min \chi^2} &\equiv \underline{\underline{\mathbf{S}}}_{\phi}^\dagger \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \mathbf{N}_{\text{meas}} \\
&\equiv \underline{\underline{\mathbf{R}}}^\dagger \mathbf{N}_{\text{meas}} \quad \text{when } m = n; \quad (4.45)
\end{aligned}$$

And lastly, applying the same equation onto an overdetermined system is equivalent to:

$$\phi_{\min \chi^2} \equiv \underline{\underline{\mathbf{S}}}_{\phi}^\dagger \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \mathbf{N}_{\text{meas}} \quad \text{when } m > n. \quad (4.46)$$

In general, in an overdetermined system, no ϕ can simultaneously satisfy all m constraints because there are at most $m - n$ conflicting constraints. Equation 4.46 solves these conflicts by choosing the best compromise ϕ that violates each of these constraints the least as weighted by the covariance matrix $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}})$. In both the fully- and overdetermined cases, no *a priori* information would be accounted for evident from the lack of ϕ_0 in Equations 4.45 and 4.46, as the information provided by $\mathbf{N}_{\text{meas}}$ alone is sufficient to uniquely determine ϕ_{sol} .

In an underdetermined system, Pseudo-Inverse has the same effect as finding a solution ϕ_{sol} on the minimum χ^2 manifold that also minimises the L_2 (i.e. Euclidean) distance from the *a priori* spectrum ϕ_0 . Unfortunately, there is no guarantee that this new ϕ_{sol} will have nonnegative flux in all its bins. Therefore it is possible (and in fact, virtually unavoidable when n is large) to get unphysical negative fluxes in this algorithm.

Uncertainty propagation in the pseudo-inverse algorithm is achieved by taking the inverse of the curvature of the χ^2 function. By differentiating the slope of the χ^2 function shown in Equation 4.39, we obtain the curvature (Hessian) matrix:

$$\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}} = \underline{\underline{\mathbf{S}}}_{\phi} = \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \underline{\underline{\mathbf{R}}}, \quad (4.47)$$

4.4 Demonstration with synthetic data

where $\underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} = \left(\underline{\underline{\mathbf{cov}}}(\mathbf{N}_{\text{meas}})\right)^{-1}$ as previously mentioned in Section 4.1.4. However, in underdetermined cases, $\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}$ is not a full-rank matrix and therefore cannot be inverted normally. Thus to obtain $\underline{\underline{\mathbf{cov}}}(\phi_{\text{sol}})$, we must use the pseudo-inverse again:

$$\underline{\underline{\mathbf{cov}}}(\phi_{\text{sol}}) = \underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}^{\dagger} = \left(\underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}}\right)^{\dagger}. \quad (4.48)$$

It is useful to bear in mind that the resultant $\underline{\underline{\mathbf{cov}}}(\phi_{\text{sol}})$ would still be a singular matrix even after this operation.

4.4 Demonstration with synthetic data

To demonstrate these algorithms in a controlled setting, several synthetic activation foil neutron spectrum unfolding datasets of increasing complexity are created, such that unfolding algorithms can be applied to these datasets to obtain the solution spectra ϕ_{sol} 's.

4.4.1 2-bin example

The simplest non-trivial unfolding problem requires at least 2 neutron bins. Having a number of bins = the number of dimensions on a page allows the vector space ϕ to be visualized intuitively on the page in Figures 4.5 to 4.8. Colour bars were occasionally used to convey further information in extra dimensions in these plots, and their colour schemes had been kept consistent across all figures in this subsection, where a **cyan** colour was used to indicate higher χ^2 , and a **magenta** colour was used to indicate higher values of the regularising function (cross-entropy or Kullback–Leibler Divergence).

To keep the problem complex enough to be interesting while being numerically easy to compute, the following variables were used: $\underline{\underline{\mathbf{R}}} = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix}$, $\phi_{\text{true}} = \begin{pmatrix} 3 \\ 3 \end{pmatrix}$, and to keep the problem simple, noise is assumed to be zero $\epsilon = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ such that Equation 3.7 gives $\mathbf{N}_{\text{meas}} = \begin{pmatrix} 9 \\ 9 \end{pmatrix}$. Its covariance matrix is assumed to be diagonal with zero epistemic error, i.e. $\underline{\underline{\mathbf{cov}}}(\mathbf{N}_{\text{meas}}) = \underline{\underline{\mathbf{Diag}}}(\mathbf{N}_{\text{meas}}) = \begin{pmatrix} 9 & 0 \\ 0 & 9 \end{pmatrix}$. The *a priori* spectrum used was $\phi_0 = \begin{pmatrix} 2.0 \\ 0.8 \end{pmatrix}$, giving a value of initial $\chi^2(\underline{\underline{\mathbf{R}}}\phi_0, \mathbf{N}_{\text{meas}}) = 5.840$.

Since AMAXED and AMAXED-Regularization both use $H(\phi'_0, \phi') = D_{KL}(\phi'_0, \phi') + H(\phi'_0)$ as the regularising functions, the loss function landscapes that both algorithms face are very

4.4 Demonstration with synthetic data




Figure 4.5: Loss function landscape for the problem set out in Section 4.4.1, where the regularising function is cross-entropy, defined as $H(\phi'_0, \phi') = D_{KL}(\phi'_0, \phi') + H(\phi'_0)$, and $H(\phi'_0)$ is a constant that can be omitted.




Figure 4.6: Loci of solution spectra ϕ_{sol} created by AMAXED-Regularization and AMAXED with all values of target $\chi^2 \leq 5.840$.

4.4 Demonstration with synthetic data

similar. This landscape is plotted in Figure 4.5, where the left plot shows the χ^2 component of the loss function, the middle plot shows the regularising component of the loss function, and the right plot shows the sum of these two (which is only applicable to AMAXED-Regularization).

As the target $\chi^2 = \Omega$ is varied, different values of ϕ_{sol} are obtained for each algorithm. Since ϕ_{sol} varies smoothly w.r.t. Ω , it traces out a continuous path as Ω is decreased. These paths are plotted as the loci in Figure 4.6. Comparing the middle against the right plot, the trajectories of ϕ_{sol} on both plots match up perfectly between $\Omega = 0$ to 2.793. However, the trajectories of ϕ_{sol} for AMAXED-Regularization and AMAXED bifurcates when Ω is increased beyond 2.793, up to 5.840. The bifurcation point is marked with an orange circle $\circ$ on both plots. This bifurcation occurs because there are two equally valid solutions to the problem of “ $\underset{\phi}{\operatorname{argmin}} D_{KL}(\phi'_0, \phi')$ confined to the manifold $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) = \Omega$ ”, evident from the fact that the line $D_{KL} = 0$ intersects with two points of the $\chi^2 = \Omega$ circle when $\Omega > 2.793$. Due to the way that the optimization engines are set up, AMAXED-Regularization always converges on the nearer intersection point while AMAXED always converges on the further one, thus the bifurcation when the target χ^2 is set to > 2.793 . Due to the monotonically increasing nature of χ^2 and D_{KL} , it is possible to deduce that the loci of ϕ_{sol} at $\Omega = 5.840$ for both algorithms should be bounded by the **red** box for both algorithms even before executing either algorithm, by examination of the loss function landscape alone. The bounded area includes all points ϕ in the closed set

$$\{\chi^2(\underline{\mathbf{R}}\phi, \mathbf{N}_{\text{meas}}) \leq 5.840\} \cap \left\{ D_{KL}(\phi'_0, \phi') \leq D_{KL} \left(\phi'_0, \left(\underset{\phi}{\operatorname{argmin}} \left(\chi^2(\underline{\mathbf{R}}\phi, \mathbf{N}_{\text{meas}}) \right) \right)' \right) \right\}. \quad (4.49)$$

As for MAXED and IMAXED, a different regularising function is used, leading to a different loss function landscape (left of Figure 4.7). As both algorithms have the same loss function (differing only in the implementation of the optimization engine), for every value of Ω , the solution spectra ϕ_{sol} arrived by both algorithms are the same, producing overlapping loci.

Similar to Figure 4.6, it is possible to deduce the bounding box for the solution ϕ_{sol} of MAXED and IMAXED at $\Omega = 5.840$ from examining the loss function landscape alone. This is drawn in **red** in Figure 4.7, enclosing the closed set

$$\{\chi^2(\underline{\mathbf{R}}\phi, \mathbf{N}_{\text{meas}}) \leq 5.840\} \cap \left\{ H(\phi', \phi'_0) \leq H \left(\left(\underset{\phi}{\operatorname{argmin}} \left(\chi^2(\underline{\mathbf{R}}\phi, \mathbf{N}_{\text{meas}}) \right) \right)', \phi'_0 \right) \right\}. \quad (4.50)$$

For completeness, Figure 4.8 shows how SAND-II, GRAVEL, and Pseudo-Inverse solves this problem. The smoothing conditions of SAND-II are not applied here due to the low number of




Figure 4.7: Loci of solution spectra ϕ_{sol} created by IMAXED and MAXED with all values of target $\chi^2 \leq 5.840$.




Figure 4.8: The optimization path taken by three purely iterative unfolding algorithms, plotted in log scale since both SAND-II and GRAVEL are algorithms that attempt to take steps in $\log(\phi_{\text{sol}})$ space. All three paths are overlapping as they all took the same path for this specific problem.

bins (see the end of Appendix E). All three algorithms are indisputably better than the rest of the algorithms in solving this specific unfolding problem, evident from the fact that both algorithms took only one step to arrive at the correct solution $\phi_{\text{sol}} = \phi_{\text{true}}$. This is because of the fully-determined nature of the system allowing for quick convergence by these three algorithms. However it is important to note that these purely iterative algorithms perform poorly with large underdetermined systems, when compared to other algorithms, and thus cannot be relied upon for fusion neutronics problems, as shall be shown in Section 4.4.3.

The reader should practise caution when attempting to generalise the intuitions learned from this subsection. Some concepts remain applicable in higher dimensions; for example, the bounding boxes of possible ϕ_{sol} stated in Equations 4.49 and 4.50 remain true in higher dimensions. They are not only helpful for understanding the unfolding problem but also useful for checking whether the unfolding program is implemented correctly or not¹¹. Other features stated in this subsection disappear in more complex problems; for example, the bifurcation of the loci of ϕ_{sol} in Figure 4.6 is not observed when the problem is underdetermined.

¹¹These bounding boxes were implemented as criteria that the computed solution ϕ_{sol} must pass in the tests of the unfoldingsuite.

4.4.2 5-bin example

Increasing the complexity of the problem slightly to a larger number of bins and degree of determination, an underdetermined system with a response matrix of shape 2 rows $\times$ 5 columns is used as an example for the following subsection. The underlying “true” neutron spectrum was arbitrarily chosen as $\phi_{\text{true}}^T = (12.5 \ 5.0 \ 20 \ 20 \ 10)$, and the response matrix is similarly created by choosing random integers between 0 to 5, giving $\underline{\underline{\mathbf{R}}} = \begin{pmatrix} 2 & 0 & 5 & 4 & 0 \\ 0 & 4 & 0 & 0 & 1 \end{pmatrix}$. Using Equation 3.7, and again applying the simplest assumptions that $\epsilon^T = \begin{pmatrix} 0 & 0 \end{pmatrix}$ and $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}}) = \underline{\underline{\text{Diag}}}(\mathbf{N}_{\text{meas}})$, the measured response vector and its associated uncertainties are calculated as $\mathbf{N}_{\text{meas}}^T = \begin{pmatrix} 205 & 30 \end{pmatrix}$, $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}}) = \begin{pmatrix} 205 & 0 \\ 0 & 30 \end{pmatrix}$. An almost-flat *a priori* spectrum $\phi_0^T = (10 \ 12 \ 12 \ 11 \ 10)$ is also arbitrarily selected. The algorithms are all supplied with the same $\underline{\underline{\mathbf{R}}}$, $\mathbf{N}_{\text{meas}}$, $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}})$ and ϕ_0 , and are instructed to find a ϕ_{sol} that produces the minimum χ^2 . Some algorithms, namely MAXED, IMAXED, and AMAXED, are written in a way such that $\Omega = \text{target } \chi^2$ can equal to zero (see Appendix A.1 for explanation), so their target χ^2 are set to 0. Other algorithms that do not allow $\chi^2 = 0$, namely AMAXED-Regularization, GRAVEL and SAND-II, had their target χ^2 set to 0.01 instead. The results are shown in Figure 4.9. Note that purely iterative algorithms are implemented in such a way that it will keep iterating until the solution spectrum’s χ^2 drops below the target χ^2 . This explains why for the two purely iterative algorithms in Figure 4.9 (GRAVEL and SAND-II), their solution spectrum’s χ^2 is much smaller than the target χ^2 value of 0.01. The rest of the algorithms tend to have a small mismatch between the solution spectrum’s χ^2 and the target χ^2 as well due to the tolerance written into the termination conditions.

While it is tempting to make the conclusion that “AMAXED, AMAXED-Regularization, and Pseudo-Inverse performed the best” from looking at the plots alone, we cannot conclude this in a quantifiable way without choosing a metric to measure the distance $D_{\text{sol,true}}$, as explained in Section 4.2.1. But if we allow ourselves to make this choice, and choose Kullback–Leibler Divergence $D_{KL}(\phi_{\text{true}}, \phi_{\text{sol}})$ and cross-entropy $H_{KL}(\phi_{\text{true}}, \phi_{\text{sol}})$ as the metrics to measure this distance, then we can measure this deviation, and these measurement results are presented in Table 4.3. These results, showing that AMAXED, AMAXED-Regularization, and Pseudo-Inverse performed the best, are not surprising at all: as long as the *a priori* is correctly chosen, the deviation of ϕ_0 from ϕ_{true} should be small, meaning that $D_{KL}(\phi_{\text{true}}, \phi_0)$ and $\sqrt{\sum_i^n |\phi_{\text{true},i} - \phi_{0,i}|^2}$ are both small. Since the solution of AMAXED and AMAXED-Regularization are both

4.4 Demonstration with synthetic data




Figure 4.9: Result of the unfolding problem in Section 4.4.2 as solved by various algorithms.

4.4 Demonstration with synthetic data

Table 4.3: How far the unfolded solutions deviated from the “true” spectrum in Figure 4.9, measured by Kullback–Leibler Divergence and cross-entropy.

Spectrum ϕ	$D_{KL}(\phi'_{\text{true}}, \phi')$ (unit: nat)	$H(\phi'_{\text{true}}, \phi')$ (unit: nat)
ϕ_0 , <i>a priori</i> spectrum	0.100	1.609
ϕ_{sol} of SAND-II	0.037	1.546
ϕ_{sol} of SAND-II (smoothing condition applied)	0.053	1.562
ϕ_{sol} of GRAVEL	0.037	1.546
ϕ_{sol} of Pseudo-Inverse	4.734×10^{-3}	1.514
ϕ_{sol} of MAXED	5.881×10^{-3}	1.515
ϕ_{sol} of IMAXED	5.881×10^{-3}	1.515
ϕ_{sol} of AMAXED	5.459×10^{-3}	1.514
ϕ_{sol} of AMAXED-Regularization	5.389×10^{-3}	1.514

minimizers of $D_{KL}(\phi_0, \phi_{\text{sol}})$, and Pseudo-Inverse’s solution is the minimiser of the L_2 distance between $(\phi_0, \phi_{\text{sol}})$, it is not hard to imagine that they likely are also minimisers of $D_{KL}(\phi_{\text{true}}, \phi_{\text{sol}})$ (deviation from ϕ_{true} to ϕ_{sol}) by referring to the diagram in Figure 4.3.

Similar to Section 4.4.1, this result does not allow us to make any general statement about the quality of one algorithm against another, because from a mathematical perspective, there yet uncountably many other *a priori* spectra ϕ_0 that would result in completely different ranking of quality of solutions than what is shown in Table 4.3. Due to reasons explained in Section 4.2.1, the frequency at which these more pathological $(\phi_0, \phi_{\text{true}})$ pairs may occur in practice is not known, and therefore we cannot use Table 4.3 to make the generalised conclusion of whether AMAXED, AMAXED-Regularization and Pseudo-Inverse will always yield the results most similar to ϕ_{true} .

4.4.3 172-bin example

To demonstrate these algorithms on a more realistic level of complexity, the following set of synthetic data mimicking fusion neutron spectrum unfolding with activation foils is used. The underlying “true” neutron spectrum ϕ_{true} is set as the ITER-DT reference input spectrum from the FISPACT wiki [85]. This is provided in the VITAMIN-J group structure, which has 175 bins; but the three lowest energy bins were removed as they recorded zero neutron fluxes, which would otherwise make computing the Kullback–Leibler Divergence in Table 4.5 not possible. To create a response matrix, a subset of 16 radionuclides from foil-set C in reference [130] was selected, listed in Table 4.4. Three foils were assumed to be used: erbium, molybdenum, and tin. Each foil was assumed to be 100% pure with exactly 1 mole of the element, with fractions

4.4 Demonstration with synthetic data

Table 4.4: Response matrix for the problem in Section 4.4.3.

Foil element	detectable radionuclide	dominant reaction path	number of reactants per mole of foil element
Er	^{163}Er	$^{162}\text{Er}(\text{n},\gamma)$	8.371×10^{20}
	^{165}Er	$^{164}\text{Er}(\text{n},\gamma)$	9.641×10^{21}
	^{171}Er	$^{170}\text{Er}(\text{n},\gamma)$	8.979×10^{22}
Mo	^{101}Mo	$^{100}\text{Mo}(\text{n},\gamma)$	5.868×10^{22}
	^{91}Mo	$^{92}\text{Mo}(\text{n},2\text{n})$	8.822×10^{22}
	^{96}Nb	$^{96}\text{Mo}(\text{n},\text{p})$	1.004×10^{23}
	^{97}Nb	$^{97}\text{Mo}(\text{n},\text{p})$	5.770×10^{22}
	$^{98\text{m}}\text{Nb}$	$^{98}\text{Mo}(\text{n},\text{p}_1)$	1.463×10^{23}
Sn	$^{113\text{m}}\text{Sn}$	$^{114}\text{Sn}(\text{n},2\text{n}_1)$ and $^{112}\text{Sn}(\text{n},\gamma_1)$	9.816×10^{21}
	$^{125\text{m}}\text{Sn}$	$^{124}\text{Sn}(\text{n},\gamma_1)$	3.487×10^{22}
	$^{123\text{m}}\text{Sn}$	$^{124}\text{Sn}(\text{n},\text{p}_1)$ and $^{122}\text{Sn}(\text{n},\gamma_1)$	6.275×10^{22}
	^{111}Sn	$^{112}\text{Sn}(\text{n},2\text{n})$	5.841×10^{21}
	$^{112\text{m}}\text{In}$	$^{112}\text{Sn}(\text{n},\text{p}_1)$	5.841×10^{21}
	$^{116\text{m}}\text{In}$	$^{116}\text{Sn}(\text{n},\text{p}_1)$	8.756×10^{22}
	$^{117\text{m}}\text{In}$	$^{117}\text{Sn}(\text{n},\text{p}_1)$	4.625×10^{22}
	^{117}In	$^{117}\text{Sn}(\text{n},\text{p})$	4.625×10^{22}

of isotopes distributed according to its natural abundance on Earth. The response matrix for the creation of the k^{th} radionuclide is then given by the following sum over all reactants r that can produce radionuclide k upon neutron bombardment:

$$R_{ki} = \sum_{\text{All reactant isotopes } X} (\text{number of isotope } X) \overline{\sigma_{k,X}(E_i)} \quad (4.51)$$

Where $\overline{\sigma_{k,X}(E_i)}$ is the microscopic cross section of the production of radionuclide k upon neutron bombardment of isotope X , averaged over the neutron energy group i . This gives a response matrix of 16 rows $\times$ 172 columns. Folding the response matrix $\underline{\mathbf{R}}$ with the “true” neutron spectrum ϕ_{true} gives a response vector $\mathbf{N}_{\text{meas}}$ that represents the number of each radionuclide created during 1 second of irradiation by the ITER–DT neutron spectrum (see Equation 3.7). As per the two previous subsections, we assume $\boldsymbol{\varepsilon} = \mathbf{0}$ and $\underline{\text{cov}}(\mathbf{N}_{\text{meas}}) = \underline{\text{Diag}}(\mathbf{N}_{\text{meas}})$. The *a priori* neutron spectrum is created by parametric approximation of ϕ_{true} : a background component (of the form $\log(y) = m \log(x) + c$) and a Gaussian peak component (initialized near the DT peak) is fitted onto the original ITER–DT spectrum to obtain a crude approximation of this DT spectrum. This is shown as the blue line in Figure 4.10. The target χ^2 is chosen as an arbitrary small value of 1000th of the initial χ^2 , i.e. target $\chi^2 = \frac{\chi^2(\underline{\mathbf{R}}\phi_0, \mathbf{N}_{\text{meas}})}{1000} = \frac{548.611}{1000} = 0.549$. The resulting ϕ_{sol} of each algorithm is then presented in Figure 4.10, and their deviations from ϕ_{true} are quantified as numerical values D_{KL} and H in Table 4.5. When reading Table 4.5, the

4.4 Demonstration with synthetic data

Table 4.5: How far do the unfolded solutions deviate from the “true” spectrum in Figure 4.10, measured by Kullback–Leibler Divergence and cross-entropy.

Spectrum ϕ	$D_{KL}(\phi'_{\text{true}}, \phi')$ (unit: nat)	$H(\phi'_{\text{true}}, \phi')$ (unit: nat)
ϕ_0 , <i>a priori</i> spectrum	0.086	13.532
ϕ_{sol} of AMAXED-Regularization	0.081	13.527
ϕ_{sol} of AMAXED	0.081	13.527
ϕ_{sol} of IMAXED	0.074	13.520
ϕ_{sol} of MAXED	0.074	13.520
ϕ_{sol} of SAND-II	0.203	13.649
(smoothing applied: $N_s = 5$)		
ϕ_{sol} of SAND-II	0.071	13.517
ϕ_{sol} of GRAVEL	0.091	13.537
ϕ_{sol} of Pseudo-Inverse	N/A (uncomputable due to negative flux)	N/A (uncomputable due to negative flux)

same caveat as given in Section 4.4.2 applies: one should not attempt to generalise this ranking of algorithms to other problems, as this ranking is undecidable for underdetermined unfolding problems until ϕ_0 and ϕ_{true} are provided.

All algorithms’ ϕ_{sol} reaches the target χ^2 , except for GRAVEL and SAND-II, which overshoot slightly to a lower χ^2 than the target. AMAXED and AMAXED-Regularization reached the same ϕ_{sol} as expected since both use the same regularising function. Similarly, MAXED and IMAXED reached the same ϕ_{sol} , with the difference being that MAXED used more iterations than IMAXED as expected as its optimization engine is less efficient, but uses less memory.

In the range of neutron bins < 100 eV (thermal, epithermal, and some intermediate neutrons), AMAXED and MAXED managed to closely replicate ϕ_{true} ; MAXED and IMAXED also managed to replicate ϕ_{true} , but to a lesser degree. This is because the microscopic cross sections used in the response matrix are generally large (and typically increase with lethargy) in this region due to thermal capture, increasing the system’s sensitivity to deviations of ϕ_0 from ϕ_{true} in this energy range.

In the range between 100 eV to 10 MeV, spurious structures such as single bins that rose above or dipped below neighbouring neutron bins in ways that were not observed in the original neutron spectrum, started to appear. This is due to strong peaks or troughs in the response matrix due to resonance reactions: where the $\mathbf{N}_{\text{sol}}$ does not match $\mathbf{N}_{\text{meas}}$ sufficiently, unfolding algorithms will increase or decrease the offending rows’ $N_{\text{sol},k}$ until they match the data $N_{\text{meas},k}$, and in the process leave behind positive or negative imprints of the offending row’s response matrix in the neutron spectrum; these imprints when there are sharp resonance peaks in the relevant rows of the response matrix. These effects are particularly pronounced in the purely

4.4 Demonstration with synthetic data




Figure 4.10: Attempts at unfolding a reference ITER spectrum using different unfolding algorithms.

iterative algorithms due to their lack of regularising function and memoryless nature. SAND-II similarly suffers from spurious structures, but it has included a smoothing parameter N_s in its later releases to combat this issue of spurious structures. However, when comparing the **orange** line with the **green** line in the bottom plot of Figure 4.10, it is obvious that tuning N_s to even a relatively small value of 5 is enough to cause a massive peak to appear in the range of 2–8 MeV, a feature that is not present in the original neutron spectrum, while doing very little to suppress these spurious structures. Despite these spurious structures, none of the algorithms managed to replicate the peak between 20 keV to 2 MeV. This is because none of the reactions is sensitive to neutron energies in this range, so unfolding algorithms have the tendency to leave the ϕ_0 unchanged in these “dark” regions unilluminated by the response matrix. This highlights the need for a robust method of constructing a suitable response matrix by selecting the correct foils.

In the fusion-relevant energy range of around 14 MeV, all of the algorithms managed to unfold a ϕ_{sol} that fits the original ϕ_{true} ’s DT peak well, simply by virtue of leaving *a priori* unperturbed, as ϕ_0 was already well-fitted to ϕ_{true} in this energy range by construction.

4.4.3.1 Theoretical resolution

Section 4.4.3 demonstrated that for a given neutron energy range, the resolution of ϕ_{sol} is proportional to the number of rows of the response matrix that are sensitive to this energy range of interest. It is easy to see this rule holds for $m \geq n$: in an overdetermined or fully-determined system in the absence of noise, the neutron spectrum can be perfectly reconstructed using any of the above unfolding algorithms by setting the target χ^2 to as low as possible (i.e. 0). However evident from Figure 4.10, when $m < n$, different algorithms achieve different levels of fidelity when reconstructing the neutron spectrum, and peaks have the tendency to be incorrectly reconstructed as imprints of resonance peaks elsewhere.

Therefore, to more formally quantify the resolution of activation foil neutron spectrum unfolding, we can make the following claim: when attempting to reconstruct the neutron spectrum in the range E_{min} to E_{max} , the neutron spectrum can be divided up to m bins where m = number of linearly independent components of the response matrix, to make the system fully-determined. In the ideal, noiseless scenario, this allows the neutron spectrum to be unfolded perfectly, establishing the best-case scenario’s resolution to be $\frac{E_{\text{max}} - E_{\text{min}}}{m}$, i.e. this would be the upper bound to its resolution. In practice, for a fusion neutron spectrum with neutrons between 0.025 eV to 17 MeV, and the number of rows in the response matrix rarely exceeds 30




Figure 4.11: Response matrix used, normalized to allow for optimal visualization. This is obtained using the simulation software SCINFUL. This is divided into 1306 pulse-height channels (rows) $\times$ 315 neutron groups (columns), making it an overdetermined problem.

(Chapter 6). If pressed to fill in the “resolution” column of the activation foil neutron spectrum unfolding entry in Table 2.1 with a numerical value, then the response “ > 0.5 MeV” would be a great candidate, with a very large asterisk attached next to it stating the caveat that “this number is the upper bound for a noiseless fully-determined system: if an underdetermined system is used instead, then some regions of the spectrum may get better resolution at the cost of worse resolution elsewhere”, exemplified by AMAXED-Regularization’s and AMAXED’s improved resolution in at < 100 eV at the cost of reduced resolution at > 0.1 MeV in Figure 4.10.

4.5 Demonstration with experimental data

To prove that the unfoldingsuite can be used to solve real-world unfolding problems, it was deployed against a set of scintillator measurements for a p+Be neutron generator in the Nuclear Physics Institute of the Czech Academy of Sciences (NPI CAS). The details of this experiment were published in the article [131]; this section will only concern the content that was not already published in the article. All of the data (except the neutron spectrum) used to create the plots in this section are obtained from collaborators from NPI who authored the article.




Figure 4.12: How the measured response vector $\mathbf{N}_{\text{meas}}$ compares to the would-have-been response vector $\mathbf{N}_{\text{sol}}$ constructed from $\underline{\underline{\mathbf{R}}}\phi_{\text{sol}}$.

A stilbene detector (described in Section 2.1) operates by measuring the scintillation light brightness due to a neutron scattering event. The total amount of light is proportional to the incident neutron's energy, forming a response matrix $\underline{\underline{\mathbf{R}}}$ shown in Figure 4.11. Element R_{ki} of this response matrix describes the intensity of pulses in the k^{th} pulse-height channel elicited by every incident neutron from neutron group i . The colour scale in Figure 4.11 is normalized by column, such that the maximum intensity of pulses created by neutrons of a given energy = 1.0.

The histogram of the height of pulses measured by the scintillator forms a response vector $\mathbf{N}_{\text{meas}}$ that resembles a bottom left triangle (red line in Figure 4.12). As with previous sections, the covariance matrix of the measured response vector is assumed to be diagonal, i.e. the errors between any two channels of the pulse-height histogram are assumed to be uncorrelated, and each channel is assumed to follow the Poisson distribution such that $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}}) = \underline{\underline{\text{Diag}}}(\mathbf{N}_{\text{meas}})$.

An *a priori* ϕ_0 was constructed from a series of functions that are piecewise linear in log- ϕ space (i.e. piecewise exponential decay functions w.r.t. increasing neutron energy) that, together formed a single continuous function (see the blue line in Figure 4.13). Then the unfolding can begin: The $\underline{\underline{\mathbf{R}}}$, $\mathbf{N}_{\text{meas}}$, $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}})$, and ϕ_0 were given to AMAXED-Regularization, which unfolded the solution shown in Figure 4.13. The uncertainty envelope of $\pm 1\sigma$ was obtained by taking the square root of the main diagonal of $\underline{\underline{\text{cov}}}(\phi_{\text{sol}}) = \text{the inverse of the Hessian matrix}$ as described by Appendix D.3. The target χ^2 of 5000 was chosen via adjusting the target χ^2




Figure 4.13: Neutron spectrum of the neutron generator at 0° away from the source [131], reproduced here for easier referencing.

as low as possible while keeping the solution neutron spectrum ϕ_{sol} monotonically decreasing w.r.t. increasing incident neutron energy, as it was reasoned that that neutron should only downscatter without much absorption, such that the scalar flux should only increase w.r.t. increasing lethargy. This non-zero final value of $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})$ was expected because this was an overdetermined problem. Other algorithms of the unfoldingsuite (AMAXED, IMAXED, GRAVEL, SAND-II, and Pseudo-Inverse) were tested as well, but AMAXED-Regularization remained the best algorithm as its result could achieve the lowest χ^2 without causing unphysical features (i.e. non-negative values in the 1st derivative of ϕ_{sol}) to appear. Therefore, the result of AMAXED-Regularization was accepted and published.

Figure 4.12 shows how closely $\mathbf{N}_{\text{sol}}$ matches the $\mathbf{N}_{\text{meas}}$. The hypothetical response vector histogram $\mathbf{N}_{\text{sol}}$ (what the response vector “would have been” if ϕ_{sol} was indeed the neutron spectrum) was constructed by

$$\mathbf{N}_{\text{sol}} = \underline{\mathbf{R}}\phi_{\text{sol}} = \sum_i^n \phi_{\text{sol},i} \cdot (\underline{\mathbf{R}}^T)_i,$$

4.5 Demonstration with experimental data

each colour of the colour bar in Figure 4.12 shows the contribution to the response vector histogram from the i^{th} neutron group. In the original study [131], time-of-flight (ToF) measurement of the neutron spectrum was also available at energies > 8.225 MeV; this ToF-measured spectrum was shown to agree with the result of the unfolded solution, residing completely within the uncertainty envelope of ϕ_{sol} .

This scintillator example problem is an overdetermined problem ($m = 1306$ channels in the detector, $\times n = 315$ neutron bins) with well-understood physics, thus it is not as difficult as activation foil fusion neutron spectrum unfolding problems. However, it consolidates AMAXED-Regularization as one of the top-performing algorithm, showing that it is able to perform well even with a large number of neutron bins (315), instilling confidence in its capability.

Chapter 5

Activation foil unfolding experiment design optimization

The previous chapters discussed how the choice of unfolding algorithm affects the quality of the unfolded solution when $\mathbf{N}_{\text{meas}}$ is already measured. This chapter explores how to design the unfolding experiment such that $\mathbf{N}_{\text{meas}}$ can be unfolded to the most accurate and precise solution. This mainly involves making changes that affect the response matrix $\underline{\mathbf{R}}$ in Equation 3.6, and the uncertainties associated with the measured response vector, $\underline{\text{cov}}(\mathbf{N}_{\text{meas}})$.

When planning an unfolding experiment, the experimentalist is faced with the problem of choosing a set of experiment parameters that maximise the effectiveness of the unfolding experiment, or expressed as a text equation:

$$\left[\underset{\text{experiment design parameters}}{\text{argmax}} \quad (\text{accuracy and precision of unfolding result}) \right] \\ \{\text{experiment design parameters} \in \text{feasible parameter ranges}\}. \quad (5.1)$$

To perform the constrained optimization in Equation 5.1, the **foil selector framework** was developed, named as such because the primary design parameter it optimises is the foil selection. It is then implemented as a Python module known as the **foil selector program**. There is a subtle difference between the framework and the program it is based upon: the former is a collection of abstract concepts of how foil selection should theoretically be done; while the latter is a repository of codes implementing these ideas. Sections 5.1 to 5.5 of this chapter focuses on deriving this framework; while Section 5.6 focuses on how this program implements the framework. Both the framework and the program are open-source: the framework is described

5.1 Formulation and assumptions




Figure 5.1: The definition of “schedule” represented graphically. A schedule is characterised by 3 lengths of time: irradiation duration ($a - 0$), transport duration ($b - a$), and acquisition duration ($c - b$).

in its entirety in this thesis, which is then published to the public domain; while the program repository is released with the MIT open-source license [132].

Since no previous work had carried out foil selection to similar levels of rigour before, both the foil selector framework and the foil selector program constitute original contributions to knowledge of this PhD. The first record of consideration given to foil selection was found in 1970 [53], where the author listed a series of criteria that a foil used should satisfy to safely and effectively measure the neutron spectra via unfolding. However, there were no further published advancements in foil selection; foil selection for unfolding experiments often proceeded in an *ad hoc* manner [50, 58, 60, 49], choosing whatever foil that the experimentalist may be most familiar with and is easily available in the laboratory. Hence such a procedural foil selection method is a breakthrough.

To begin developing this framework, we must first formalise the problem in mathematical terms (Section 5.1), list all relevant assumptions, outline what can (parameters, Section 5.2) and cannot (constraints, Sections 5.3 and 5.4) be changed, and then decide on how to optimise the changeable parts (Section 5.5).

5.1 Formulation and assumptions

To solve Equation 5.1 for an activation foil unfolding experiment, it is necessary to explain how such an experiment proceeds.

Figure 5.1 breaks down the irradiation and measurement schedule of each foil into 3 stages: the activation foil is irradiated by an unchanging neutron spectrum between time $t = 0$ — a , during this time, radionuclides are both produced and decayed away. Between time $t = a$ — c , only decays occur in the foil as it is removed from the neutron field. The experimentalist places the foil into a gamma-ray detection set-up at time $t = b$ to record gamma-rays exiting

5.1 Formulation and assumptions

the foil until $t = c$. The number of gamma-rays recorded in this time can be used to deduce N_{meas} , the number of radionuclides created during time $t = 0 - a$, using the procedure stated in Section 5.1.1.

The response vector $\mathbf{N}$ for activation foil neutron spectrum unfolding is usually defined as either the production rate of detectable isotopes or the total number of such isotopes produced. This thesis will be using the latter definition. To calculate the gross number of the k^{th} root¹ radionuclides N_k produced by, the k^{th} row of the response matrix $(\underline{\mathbf{R}})_k$ is required. In Section 3.1.2, the precise expression for this response matrix was dismissed as a “proportionality constant” multiplied by the microscopic cross section (Equation 3.5). This proportionality constant will be revealed in Equation 5.8 below.

Activation foil irradiation begins with a foil made of stable isotopes, inserted into a neutron field with a time-invariant neutron spectrum. A fraction of neutrons that are incident on the foil will interact with the nuclei in the foil, creating new nuclei. We are only interested in the production of radioactive nuclides, as any stable nuclides produced cannot be detected by gamma-ray spectroscopy. The production rate of the k^{th} root radionuclide at any given moment directly due to irradiation by the neutron spectrum ϕ is given by

$$(\text{Production rate})_k = A \sum_{i=1}^n \left(1 - e^{-(N_d \sigma_{ki})T}\right) \phi_i, \quad (5.2)$$

where A is the area of the foil (unit: cm^2), T is the thickness of the foil (unit: cm), the definitions of group-wise microscopic cross section σ_{ki} and flux ϕ_i are defined in Equation 3.4 and Equation 3.2 respectively, and N_d is the number density of the reactant isotope (unit: number of atoms cm^{-3}). Equation 5.2 only calculates the production rate of k due to one reactant isotope. If multiple reactant isotopes can produce k , the exponent in Equation 5.2 must be summed over all of the relevant reactant isotopes X , to obtain the macroscopic cross section Σ_{ki} ²:

$$\Sigma_{ki} = \sum_{\substack{\text{All isotopes} \\ \text{in the foil} \\ \text{at } t=0}}^X \sigma_{ki,X} N_{d,X}, \quad (5.3)$$

¹The term “root” (radio)nuclide is not a widely adopted term, but it is clear from looking at Equation 5.9 that nuclide A is the “root” of the decay chain. Therefore there is a need of coining this term to refer to nuclide such as A , which is produced as the first nuclide in the decay chain spawned by the k^{th} reaction.

²The macroscopic cross section Σ_{ki} can be distinguished from the summation sign $\sum_n$ by the location of subscripts: (k, i) are placed to the right of the Σ , rather than directly below. For further details please refer to the entry for Σ in Appendix K.2.

5.1 Formulation and assumptions

which represents the probability of reaction per unit thickness (unit: cm^{-1}). This means Equation 5.2 can then be modified into

$$(\text{Production rate})_k = A \sum_{i=1}^n \left(1 - e^{-\Sigma_{ki}T}\right) \phi_i. \quad (5.4)$$

Using the assumption applied in Section 5.1.5 to take the first-order approximation of the term in the parentheses, this equation can be simplified into

$$(\text{Production rate})_k = A \sum_{i=1}^n \Sigma_{ki} T \phi_i \quad (5.5)$$

$$= AT \sum_{i=1}^n \Sigma_{ki} \phi_i. \quad (5.6)$$

Summing over the entire irradiation period, the gross number of nuclei of root radionuclide produced is N_k (ignoring their decays),

$$N_k = ATa \sum_{i=1}^n \Sigma_{ki} \phi_i, \quad (5.7)$$


where a = the duration of irradiation, as defined in Figure 5.1. Comparing this against Equation 3.6, we can show that

$$R_{ki} = ATa \Sigma_{ki}. \quad (5.8)$$

5.1.1 Measuring the response vector $\mathbf{N}$ for activation foils

How then, can we infer the quantity $\mathbf{N}$ above (gross number of nuclei of the root radionuclide produced) from the gamma-ray spectrum? We rely on counting the number of γ -rays emitted by the root radionuclide and its descendants detected during the acquisition phase, and then rewinding back in time using the relevant decay information to deduce the $\mathbf{N}$ that must have been present to cause these γ -ray counts.

Let the root nuclide be named A , so that the production and decay of A follow the paths:






Figure 5.2: Example time-evolution of the population of a hypothetical daughter nuclide produced by neutron irradiation, with decay constant $\lambda_A = \frac{1}{3}$ per minute.

If a number of nuclei of $A = P_{A,0}$ were spontaneously created at time $t = 0$, its population would follow a simple exponential decay curve

$$P_A(t) = \begin{cases} 0 & \text{when } t < 0 \\ P_{A,0}e^{-\lambda_A t} & \text{when } t \geq 0 \end{cases}. \quad (5.10)$$

$\mathcal{B}_A(t) = \frac{P_A(t)}{P_{A,0}}$ is named the 1st generational Bateman equation, for notational convenience.

However, if these $P_{A,0}$ nuclei of A were produced continuously (at a constant rate) over the period from $t = 0$ — a rather than spontaneously produced all at once at time $t = 0$, then the radionuclide population would have to be obtained by convolution of a top-hat function with Equation 5.10 $\overline{P}_A(t) = (\Pi_a * P_A)(t)$, where the top-hat function is defined as

$$\Pi_a(t) = \begin{cases} \frac{1}{a} & \text{when } 0 \leq t \leq a \\ 0 & \text{otherwise} \end{cases}. \quad (5.11)$$

An example of this convolved curve is shown in Figure 5.2. The instantaneous activity (i.e. decay rate) of A equals $\lambda_A \times$ (the current population of A). Therefore, the total number of decays of A that occurred between time $t = b$ to $t = c$ can be calculated by integrating the population of A w.r.t. time over that period, and then multiplying it by the decay constant λ_A ,

5.1 Formulation and assumptions

as shown in Figure 5.2, obtaining the expression:

$$\text{Number of decays of } A = \lambda_A \int_{t=b}^{t=c} \overline{P_A}(t) dt. \quad (5.12)$$

This method of calculating the activity can be generalised to isotopes further down the decay chain. For example, if the decay chain has started with $P_{A,0}$ nuclei of A and 0 nuclei of B spontaneously created at time $t = 0$, then the population of radionuclide B can be described by the Bateman equation for the 2nd generation of radionuclides in the decay chain:

$$P_B(t) = \begin{cases} 0 & \text{when } t < 0. \\ P_{A,0} \frac{\lambda_A}{\lambda_B - \lambda_A} (e^{-\lambda_A t} - e^{-\lambda_B t}) & \text{when } t \geq 0, \end{cases} \quad (5.13)$$

$\mathcal{B}_{AB}(t) = \frac{P_B(t)}{P_{A,0}}$ is named the 2nd generational Bateman equation for notational convenience as well.

Similarly, if these $P_{A,0}$ nuclei of A were built up slowly rather than created spontaneously, then the radionuclide population of B can be obtained by convolving $\overline{P_B}(t) = (\Pi_a * P_B)(t)$, and the number of decays of B can then be obtained by multiplying the integral with λ_B as shown in Figure 5.3.

$$\text{Number of decays of } B = \lambda_B \int_{t=b}^{t=c} \overline{P_B}(t) dt. \quad (5.14)$$

This pattern of finding the population and total number of decays using $\overline{P_X}(t) = (\Pi_a * P_X)(t)$ for radionuclide X can be extended to any arbitrary n^{th} generation of radionuclide down the decay chain $A \rightarrow B \rightarrow \dots \rightarrow X$, giving the general equation of:

$$\text{Number of decays of } X = \lambda_X \int_b^c (\Pi_a * P_X)(t) dt. \quad (5.15)$$

where

$$P_X(t) = P_{A,0} \mathcal{B}_{AB\dots X}(t) = P_{A,0} \times \text{the } X^{\text{th}} \text{ generational Bateman equation.} \quad (5.16)$$

Equation 5.16 is only valid when the decay chain of X is linear, i.e. a decay chain $A \rightarrow B \rightarrow \dots \rightarrow X$ with no branches. To generalise this result to decay chains with branches, we can multiply Equation 5.15 with all of the branching fractions, then sum over all of the pathways that reach X , in a procedure detailed in Appendix G, such that P_X is instead given by Equation G.3.




Figure 5.3: Example time-evolution of the population of a hypothetical daughter nuclide B, produced as a secondary product of neutron irradiation in the decay chain [reactant] $\xrightarrow{n}$ A $\xrightarrow{\lambda_A}$ B $\xrightarrow{\lambda_B}$... with decay constants $\lambda_A = 1.77$ per minute, $\lambda_B = 0.36$ per minute.

Then, to calculate the total number of γ -rays counted due to the decay of X , we simply multiply the number of decays of X by the γ -rays counted per decay:

(Number of γ -rays measured at the peak with energy E_l due to the decay of X)

$$= (\text{Number of decays of } X) \times \sum_l^L \epsilon_{\text{abs}}(E_l) I_{X,l}, \quad (5.17)$$

where the $\epsilon_{\text{abs}}(E_l)$ is the absolute detection efficiency of the gamma-ray detector set-up at the energy of the l^{th} gamma-line of that nuclide, L is the total number of gamma-lines emitted by X , and $I_{X,l}$ is the intensity of the l^{th} gamma line denoting how many of such γ -rays would have been emitted per decaying nucleus of nuclide X .

If we create a vector $\mathcal{N}$ where $\mathcal{N}_l$ represents the number of counts in the l^{th} gamma-ray peak in the gamma-ray spectra (with units of gamma-ray counts), then Equation 5.17 can be re-written into:

$$\mathcal{N}_l = \sum_k^m \left(\epsilon_{\text{abs}}(E_l) \sum_{\substack{\text{all radionuclides} \\ \text{which are} \\ \text{descendants of} \\ \text{root nuclide } k}}^X I_{X,l} \left(\lambda_X \int_b^c \left(\Pi_a * \frac{P_X}{P_{A,0}} \right) (t) dt \right) N_k \right). \quad (5.18)$$

5.1 Formulation and assumptions

Or, to account for more complex decay chains containing branches, we use the notation in Equation G.4 where the population of X at time t due to the decay of root nuclide k during the irradiation is $P_{k...X}(t)$, giving:

$$\mathcal{N}_l = \sum_k^m \left(\epsilon_{\text{abs}}(E_l) \sum_{\substack{\text{all radionuclides} \\ \text{which are} \\ \text{descendants of} \\ \text{root nuclide } k}}^X I_{X,l} \left(\lambda_X \int_b^c \frac{\overline{P_{k...X}(t)}}{P_{k,0}} dt \right) N_k \right). \quad (5.19)$$

Equation 5.19 can be re-written into a matrix multiplication by simplifying the complicated summation expression inside the bracket as $\underline{\underline{\mathbf{G}}}$:

$$\mathcal{N} = \underline{\underline{\mathbf{G}}} \mathbf{N}, \quad (5.20)$$

where

$$G_{lk} = \epsilon_{\text{abs}}(E_l) \sum_{\substack{\text{all radionuclides} \\ \text{which are} \\ \text{descendants of} \\ \text{root nuclide } k}}^X I_{X,l} \left(\lambda_X \int_b^c \frac{\overline{P_{k...X}(t)}}{P_{k,0}} dt \right). \quad (5.21)$$

Combining Equation 5.20 with Equation 3.6, we have:

$$\mathcal{N} = \underline{\underline{\mathbf{G}}} \underline{\underline{\mathbf{R}}} \phi \quad (5.22)$$

$$\mathcal{N} = \underline{\underline{\mathbf{R}}} \phi \quad (5.23)$$

Where $\mathcal{N}$ is the “effective response vector” describing the number of counts in each peak in the gamma-ray spectra, and $\underline{\underline{\mathbf{R}}}$ is the “effective response matrix”, with units of gamma-ray counts cm^2s . The relationships between the operators ($\underline{\underline{\mathbf{R}}}$, $\underline{\underline{\mathbf{G}}}$, and $\underline{\underline{\mathbf{R}}}$) and the measured quantities (ϕ , $\mathbf{N}$, and $\mathcal{N}$) are best expressed in Figure 5.4. The upper bound of the rank of $\underline{\underline{\mathbf{G}}}$

$$\text{rank}(\underline{\underline{\mathbf{G}}}) \leq \min(\text{rank}(\underline{\underline{\mathbf{R}}}), L);$$

and by the definition stated in Section 3.1.2, every row of the response matrix $\underline{\underline{\mathbf{R}}}$ must produce a root radionuclide that produces *measurable* gamma-ray output, so $L \geq m$. Since $\text{rank}(\underline{\underline{\mathbf{R}}}) = \min(m, n)$, the rank of $\underline{\underline{\mathbf{G}}}$ must be bounded from above by

$$\text{rank}(\underline{\underline{\mathbf{G}}}) \leq \text{rank}(\underline{\underline{\mathbf{R}}}) = \min(m, n).$$




Figure 5.4: Diagram showing how the neutron spectrum (1) is converted into gross number of root radionuclides produced (2) using Equation 5.7, which is then converted to the number of decays of radionuclides (3) using Equation G.4, then finally converted to gamma-ray counts at the photopeaks (4) using Equation 5.19

The only way for $\underline{\underline{\mathbf{G}}}$ to have a rank fewer than this upper bound is if two of the root radionuclides have decay chains with the exact same gamma-ray spectra, which can be assumed to have zero

likelihood of happening in practice³. Therefore,

$$\text{rank}(\underline{\mathbf{G}}) = \min(m, n), \quad (5.24)$$

$$\text{rank}(\underline{\mathbf{R}}) = \min(m, n) = \text{rank}(\underline{\mathbf{R}}). \quad (5.25)$$

Apart from the assumption that “no two decay chains shall have identical gamma-ray spectra” mentioned above, several other important assumptions were made when creating the foil selector. The necessity and validity of the rest of these assumptions are explained in the following Sections 5.1.2 to 5.1.6.

5.1.2 Reactant quantity time-invariance

We assume that the number of reactants remains effectively unchanged throughout the irradiation step so that the linear relationship between flux and reaction rate holds. This means the consumption of existing reactants does not significantly impact the number of nuclei produced in the end, and the addition of new reactants (including the production of new nuclides that open up new reaction channels) does not significantly impact the effective response vector $\mathcal{N}$. This can similarly be validated with a back-of-the-envelope calculation.

Consider the former half of the claim: consumption of existing reactants does not significantly impact N_k . This is verified by extracting the probability of reaction term P_{reaction} from Equation 5.7.

$$\begin{aligned} N_k &= a(AT) \sum_i \Sigma_{ki} \phi_i \\ &= a(\text{volume of foil}) \sum_i N_d \sigma_{ki} \phi_i \\ &= a(\text{number of reactants in foil}) \sum_i \sigma_{ki} \phi_i \\ \frac{N_k}{\text{number of reactants in foil}} &= a \sum_i \sigma_{ki} \phi_i \\ P_{\text{reaction}} &= a \sum_i \sigma_{ki} \phi_i. \end{aligned}$$

Taking conservative (i.e. leaning to the side of overestimation) estimates of standard flux $\phi = 10^{14} \text{ cm}^{-2}\text{s}^{-1}$ ([20]) and $\sigma(E) = 10^4 \forall 0 \leq E \leq 20 \text{ MeV}$, $a = 2 \text{ hours} = 7200\text{s}$. Then the

³In practice, this only happens in cases where both decay chains contains only nuclides that only emit positron-annihilation γ -rays at 511 keV. This has been dealt with in the foil selector framework by excluding elements that only produces 511 keV γ -rays. In the unlikely case that this still arises, or if we do not wish to exclude reactions that emits only 511 keV as their gamma-ray spectra, then more details about how to preserve the rank of $\underline{\mathbf{R}}$ is given in Appendix H.

fraction of reactant depleted is then given by

$$P_{\text{reaction}} \approx (7200)(10^4 \times 10^{-24})(10^{14}) = 7.2 \times 10^{-3},$$

meaning the amount of reactant used up would still not reach 1%.

The latter half of the claim is harder to defend: The number of counts found in the measured gamma-ray spectra $\mathcal{N}$ can be significantly altered in two ways when reactant x reacts with neutrons to produce reactant y : One is when x is abundant while y is scarce, where the amount of x transmuted into y overwhelms the amount of y originally present, causing an inaccurate estimation of y ; and the other is when the counts of γ -rays measurable per decay of radionuclide product of y vastly outweighs the counts of γ -rays measurable per decay of radionuclide product of x , causing the number of counts of gamma-rays of some peaks $\mathcal{N}_i$ to increase beyond expectation, or be drowned out by background (electronic noise, Compton scattered gamma-ray counts, pile-up events, or external radiation sources). It is difficult to come up with a definitive rule/criterion for when this happens due to the chain of if-conditions required to catch these cases. Therefore, the responsibility of checking this particular assumption has been delegated to the user of the foil selector framework, such that foils violating this assumption may be rejected or have their response matrix modified so that the reactions of interest are not affected by the aforementioned perversion in the measurement of $\mathcal{N}$ due to the production of unwanted isotopes.

5.1.3 Neutron spectrum time-invariance

Another assumption of invariance is that the neutron spectrum is expected to be time-independent for Equation 5.8 to be valid. This allows the radioisotope production rate to be modelled by Equation 5.15 as the convolution of the normalized population curve with the top-hat function (Equation 5.1.1). And it should be obvious that the neutron spectra experienced by all foils are assumed to be the same, since all of these measurements will be used to reconstruct a single neutron spectrum by the end of the experimental analysis.

To guarantee the validity of this assumption would, again, be down to the responsibility of individual users, who should check that their neutron source has a time-invariant neutron spectrum.

5.1.4 Neutron spectrum position-invariance

Another assumption related to the invariance of the neutron spectrum is that it is position-independent: we assume that the reactions occurring at the front of the foil do not scatter a significant number of neutrons into other (typically lower) energy bins, increasing reaction rate at the back of the foil significantly. This allows us to assume that the reaction rate induced by neutrons in the i^{th} bin does not interfere with the reaction rate induced by neutrons in the j^{th} bin (where $j \neq i$). This assumption is necessary for Equation 5.2 to be valid.

This is a valid assumption as the microscopic cross sections for (down-)scattering are typically smaller than the microscopic cross sections at the resonances of more energetically favourable reactions, such as neutron capture; and the latter has been sufficiently defended against in the assumption in Section 5.1.5, where the probability of the latter has been limited to $\leq P_{\text{threshold}}$, so the probability of the former must be $\ll P_{\text{threshold}}$. The typical maximum elastic neutron scattering cross sections are around 10^2 barns, while the typical neutron capture cross sections at resonances routinely reach $> 10^5$ barns, so it is reasonable to assume the maximum probability of the neutron scattering is less than $10^{-3} \times P_{\text{threshold}}$. Using neutron scattering theory, we know that any neutrons with initial energy E elastically scattering in a medium with atomic mass A will have final energy Ξ uniformly distributed in the energy space $\left(\left(\frac{A-1}{A+1}\right)^2 E\right) \leq \Xi < E$. Assume a uniformly spaced energy group structure were used, and let the energy range $\left(\left(\frac{A-1}{A+1}\right)^2 E\right) - E$ covers g energy bins, where g is an integer > 1 . Then the maximum number of neutrons elastically scattered from bin i to bin j is $10^{-3} \times P_{\text{threshold}} \left(\frac{\phi_i}{g}\right)$. If $\phi_j \geq \phi_i$, then an addition of $10^{-3} \times P_{\text{threshold}} \left(\frac{\phi_i}{g}\right)$ (where, for the sake of argument, $P_{\text{threshold}}$ is chosen as 20%) would be an insignificant fractional change ($= 2 \times 10^{-4}$) to the flux in the j^{th} bin, satisfying the assumption.

5.1.5 Limited self-shielding from neutrons

The neutron spectrum is not only assumed to be unchanging in shape across the foil's front to back, but also unchanging in intensity, i.e. removal of neutrons from the i^{th} neutron group via either scattering or absorption of the neutrons at the front of the foil does not significantly decrease the number of neutrons in the same group available for reaction in the back of the foil. The nonlinear expression in Equation 5.4 is simplified into a linear expression in Equation 5.7 with the help of this assumption. Without this assumption unfolding algorithms would be even more complicated and computationally intensive as they cannot draw on the existing knowledge and packages on linear algebra.




Figure 5.5: The first-order approximation of the probability of having the k^{th} reaction breaks down when $\Sigma_{ki}T = N_d\sigma_{ki}T$ is too large.

This allows us to take only the first-order approximation of the reaction rate w.r.t. thickness, i.e.

$$(1 - e^{-\Sigma_{ki}T}) = 1 - 1 + \Sigma_{ki}T - \frac{(\Sigma_{ki}T)^2}{2!} + \frac{(\Sigma_{ki}T)^3}{3!} - \dots \quad (5.26)$$

$$\approx \Sigma_{ki}T. \quad (5.27)$$

(The macroscopic cross section Σ was defined in Equation 5.3.)

However, this assumption may be violated when foil contains a reaction which has a large resonance peak. Consider Figure 5.5, where N_d and σ_{ki} are chosen to resemble the resonance capture cross section of Cadmium, whose maximum microscopic cross section reaches $> 10^6$ barns. A foil with a modest thickness of 0.01 mm is already enough to make the first-order approximation of Equation 5.26 deviate from the true probability by 53.3%. This can lead to an underestimation of the flux in the relevant resonance regions of the neutron spectrum during unfolding.




Figure 5.6: The difference in path length variability between irradiation by a unidirectional source and an extended source is illustrated. The red lines denotes the path taken by neutron across the activation foil.

The most elegant solution, often employed in neutron activation studies [133–135], is to multiply the response matrix with a term known as the “self-shielding correction factor”:

$$R_{ki,\text{corrected}} = R_{ki} \times \frac{(1 - e^{-\Sigma_{ki}T})}{\Sigma_{ki}T}. \quad (5.28)$$

While at first sight, this is the most straightforward workaround that requires minimal effort, we must also consider that there are uncertainties associated with the nuclear data σ_{ki} , and uncertainties associated with T due to manufacturing tolerances, as well as the angle at which neutron crosses the foil. Figure 5.6 shows the effect of using an extended neutron source, rather than a point-like neutron source at infinity that can be approximated as a unidirectional beam: some neutrons will pass through the foil at an angle, rather than perfectly perpendicularly. Then the path length taken by that neutron to pass through the foil would differ from the measured thickness T , so Equation 5.28 is no longer valid. Furthermore, the uncertainties on σ_{ki} and T will bring extra uncertainty into the response matrix, making uncertainty propagation even more complicated. Therefore, multiplying by a “self-shielding correction factor” would only serve to complicate the analysis.

Instead, we can defend against this assumption violation by requiring the foil to be thin enough to meet the following criterion:

$$\begin{aligned}
 \max_E(P_{\text{reaction}}(E)) &\leq P_{\text{threshold}} \\
 (1 - e^{-(\Sigma_{k,\text{max}})T}) &\leq P_{\text{threshold}} \\
 1 - P_{\text{threshold}} &\leq e^{-\Sigma_{ki}T} \\
 T &\leq \frac{-\log(1 - P_{\text{threshold}})}{\Sigma_{k,\text{max}}}, \tag{5.29}
 \end{aligned}$$

where the probability of reaction is P_{reaction} , the threshold probability $P_{\text{threshold}}$ is a user-chosen parameter, and $\Sigma_{k,\text{max}}$ is the maximum macroscopic cross section of that reaction. More details can be found in Appendix I.

In practice, choosing a foil thinner than criterion 5.29 is not always possible due to concerns about the price or fragility of such a foil, and possibly due to the lower bound to the thickness established in Section 5.1.6 below, and a thicker foil must be used instead. In those cases, the experiment analysis procedure can simply disregard the reactions that violate that criterion during the unfolding experiment, and proceed to analyse the rest of the reactions as normal (with the additional assumption that the consumption of neutrons by the disregarded reaction does not significantly affect the reaction rates of the rest).

5.1.6 Complete retention of radionuclides produced

The foil selector framework assumes that 100% of the radionuclides produced by all neutron-induced reactions and nuclear decays will remain in the foil. However, due to conservation of momentum, part or all of the momentum of the neutron will be transferred to the nucleus, so a small fraction of them, particularly those produced on the far side of the foil opposite to the side facing away from the neutron source, will escape the foil into the irradiation environment before they lose all of their kinetic energy due to Coulomb interactions with the bulk of the foil. This causes the framework to overestimate the number of radioactive nuclei that should have stayed in the foil to produce γ -rays during the acquisition phase. If this loss of radionuclides were not accounted for during the unfolding analysis, then the total number of reactions N_{meas} inferred from the gamma-ray spectrum analysis would be underestimated, leading to an underestimation of neutron flux in the energy bins that produced these nuclides. Nuclear decays do not release enough energy to cause nuclei to recoil energetically enough to exit the foil, so the only legitimate

5.1 Formulation and assumptions

concern is with neutron-induced reactions that occur during the irradiation phase producing fast ions.

Fortunately, this assumption often holds at fusion-relevant energy ranges. Consider the highest energy neutron source available in fusion: the DT peak at 14.1 MeV. In the most extreme case where all of that energy was transferred to the nucleus (achievable by neutron capture, which is very unlikely at high energy but nonetheless has a non-zero cross section of reaction), a fast ion of <14.1 MeV will be produced⁴. Slowing down such an ion typically requires < 10 μ m of material. This is exemplified by calculating the stopping range of a few materials commonly used in unfolding with SRIM [136]:

- 14.1 MeV Al ion stopping in aluminium foil: 6.25 μ m $\pm$ 1487Å
- 14.1 MeV Au ion stopping in gold foil: 1.08 μ m $\pm$ 1601Å
- 14.1 MeV Sn ion stopping in tin foil: 3.15 μ m $\pm$ 3149Å

For ease of handling, foils used in neutron spectrum unfolding experiments typically have thicknesses > 0.1mm, so this means in the absolute worst-case scenario < 10% of the nuclei produced will be ejected, and > 90% of them would be retained. And for other reactions where more than one projectile (e.g. an excited nucleus and an inelastically scattered neutron, or a proton from an (n,p) reaction) is produced, the nucleus is the heaviest fragment, so it will carry much less energy than the lighter fragment, carrying an energy $\ll$ 14.1 MeV, likely in the range of 1 — 10 keV. Therefore the risk of these radionuclides not being retained in the foil is even lower.

This requirement can be formalised to the criterion:

$$\frac{\text{maximum stopping range for the radionuclide produced in that foil}}{\text{thickness of foil}(T)} \leq \text{user-chosen fraction} . \quad (5.30)$$

Criterion 5.30 ensures that the amount of radionuclides retained is always more than or equal to the user-chosen fraction.

Users who wish to conduct unfolding experiments at higher energies, typically at beamlines that mimic high altitude atmospheric neutron spectrum, should beware of violations to criterion 5.30, as the stopping range required to slow and entrap the fast ion produced increases with energy. If due to practical limitations such as foil availability, criterion 5.30 cannot be adhered to, then a sample holder or backing material made of neutron transparent material (such as a

⁴The outgoing ion will typically have less than 14.1 MeV as the release of γ -ray in the (n, γ) neutron capture reaction will carry away a portion of the energy.

5.2 Experiment design parameters

graphite capsule like the one used in [58]) may be used instead to capture the escaping ions. N.B. If used, these holders or backing materials will also have to be placed into the gamma-ray detector during the acquisition phase and included in any MCNP modelling of the gamma-ray detector, so the γ -rays emitted by the captured radionuclide are still detected.

5.1.7 Accurate *A priori*

The purpose of the foil selector framework is to select the best response matrix for measuring the neutron spectrum. When using the foil selector framework, we are determining what response matrix is best for measuring the neutron spectrum. The optimal response matrix to minimise uncertainty (i.e. covariance matrix $\text{cov}(\phi_{\text{sol},i}, \phi_{\text{sol},j})$) differs depending on ϕ_{sol} .

For example, consider two hypothetical response matrices with $m = 2, n = 2$: $\underline{\underline{\mathbf{R}}}_1 = \begin{pmatrix} 1, & 3 \\ 1, & 2 \end{pmatrix}$ and $\underline{\underline{\mathbf{R}}}_2 = \begin{pmatrix} 3, & 1 \\ 2, & 1 \end{pmatrix}$. Given two different neutron spectra $\phi_1 = \begin{pmatrix} 3 \\ 5 \end{pmatrix}$ and $\phi_2 = \begin{pmatrix} 5 \\ 3 \end{pmatrix}$, if we allow ourselves to assume the distribution of $\mathcal{N}_k$ is Poissonian and independent of each other (i.e. standard deviation of the k^{th} reaction $\sigma(N_k) = \sqrt{N_k}$, and $\text{cov}(N_1, N_2) = 0$), then it is clear that $\underline{\underline{\mathbf{R}}}_1$ is preferred over $\underline{\underline{\mathbf{R}}}_2$ for measuring ϕ_1 , as $\underline{\underline{\mathbf{R}}}_1 \phi_1$ gives a higher count in both reactions and hence a lower fractional uncertainty than $\underline{\underline{\mathbf{R}}}_2 \phi_1$. The converse is true, i.e. $\underline{\underline{\mathbf{R}}}_2$ is preferred over $\underline{\underline{\mathbf{R}}}_1$ for measuring ϕ_2 as it gives more counts in both reactions.

To conclude this demonstration, the optimal choice of response matrix is dependent on what ϕ_{true} is present. Unfortunately, ϕ_{true} is unknown to the experimentalist at the time of planning the unfolding experiment, as the very goal of the unfolding experiment is to obtain ϕ_{true} (or a close approximation of it) via unfolding. Therefore, the user of the unfolding program must provide a guessed *a priori* spectrum ϕ_0 instead. We can only assume that true neutron spectrum ϕ_{true} closely resembles the ϕ_0 i.e. the point representing ϕ_{true} in ϕ space is very close to the point representing ϕ_0 . While there is no way to justify this assumption as generally correct in all use cases of the foil selector framework, it is an unavoidable assumption. Therefore the responsibility of checking $\phi_0 \approx \phi_{\text{true}}$ can only be delegated to the user instead of handled by the foil selector framework.

5.2 Experiment design parameters

The following two sections discuss the factors that affect unfolding effectiveness as outlined in Figure 5.7. It attempts to build a fairly comprehensive but not exhaustive planning checklist




Figure 5.7: Context diagram showing how practical limitations may constrain the feasible range of parameters, and how may affect the accuracy and precision of the unfolding.

for research groups that may need to perform their unfolding experiments, as it is impossible to capture the needs of all types of unfolding experiments, past and future.

The factors directly controlling the effectiveness of unfolding (rightmost section of Figure 5.7) will be discussed further in Section 5.5, while the experiment design parameters (middle section of the figure) that in turn control these are discussed in this section. These are the “experiment design parameters” shown in Equation 5.1 at the beginning of this chapter when outlining the optimization problem.

1. **Number of foil samples used in the unfolding experiment:** Using more foils means a larger variety of materials can be irradiated and measured (see point 7). Alternatively, repeated irradiations and measurements of foils of identical composition can increase the count of γ -rays detected for those reactions as well.

2. **Foil volume:** a larger volume of foil (i.e. a larger number of reactant nuclei) means more radionuclides produced (and hence more decay counted) per unit flux incident on the foil. More reactions may have sufficient counts of γ -rays to be used in unfolding analyses as well, increasing the degree of determination.
3. **Foil thickness:** two considerations compel the choice of thicknesses in opposing directions. A thinner foil experiences less issue with self-shielding (Section 5.1.5), but a thicker foil experiences less issue with radionuclide escapement (Section 5.1.6). To quell both issues, the thickness has to be chosen from a range of values between the lower bound T_{lower} established by criterion 5.30 and the upper bound T_{upper} established by criterion 5.29. The upper and lower bounds for each reaction are different. If the final chosen thickness T_{final} lies outside of this range for some reactions (i.e. $T_{\text{final}} > T_{\text{upper}}$ or $T_{\text{final}} < T_{\text{lower}}$), then those reactions cannot be used in the unfolding analysis and will have to be discarded, reducing m and decreasing the degree of determination.

Another consideration, not reflected in Figure 5.7, is that the self-shielding of γ -rays (instead of self-shielding of neutrons as discussed in Section 5.1.5) by the foil adds to the complexity of the gamma-ray spectroscopy, and reduces the absolute detection efficiency of γ rays as discussed in point 8, so having a thinner foil is also desirable in that respect.

4. **Duration of irradiation:** longer irradiation means more radionuclides are produced, and a larger population of radionuclides is present at the end of the irradiation for the gamma-ray detector to detect, leading to a higher count of γ -rays. This is particularly useful for radionuclides further down the decay chain (far away from the root nuclide) or long half-lives, as they take longer to reach their saturation population. More of such radionuclides will be eligible for use in unfolding if the irradiation duration is increased, increasing the number of available reactions m and increasing the degree of determination as a result.
5. **Duration of transport:** a shorter transport duration reduces the fraction of radionuclides that have decay away before measurement, so that a larger fraction of radionuclides may have a sufficient count rate to be detectable. Radionuclides with short half-lives and close to the root of the decay chain are more sensitive to these changes, as a large fraction of them would have decayed away during the transfer from the neutron field to the gamma-ray detector. More of these radionuclides will be eligible for use in unfolding when transport

duration is shortened, increasing the number of available reactions m and increasing the degree of determination as a result.

6. **Duration of gamma-acquisition:** a longer acquisition period means a larger fraction of the radionuclides decay on the gamma-ray detector, increasing the counts on the gamma-ray spectrum and decreasing the fractional error on the measured number of counts. Similar to increasing the duration of irradiation, having a longer duration of gamma-acquisition is particularly beneficial for nuclides with long half-lives or further down the decay chain, as it takes longer to deplete the storage of these nuclides, making more of these radionuclides eligible for use in unfolding experiments and increasing the degree of determination.
7. **Range of materials used:** the type of radionuclide produced is entirely dependent on what reactant nuclides are present in the foils. These include control of:
 - Chemical composition (i.e. what elements are present) of the foil dictates what neutron-induced reactions may occur, and hence what radionuclides may be measured, or what noise may be present in the spectrum due to contaminants.
 - Isotope composition control of the foil (i.e. foil enriched with a certain isotope) may be available to well-funded research groups in rare cases.

Having a large variety of radionuclides increases m of the response matrix, improving the degree of determination. A correct choice of foil material (e.g. using foils free of impurities that can be activated to pollute the background of the gamma-ray spectrum) can also reduce the uncertainty on the γ -ray counts extracted by analysing the gamma-ray spectrum, improving the unfolding precision. The theoretical benefits of improving the range of materials used are demonstrated in [130].

8. **γ -ray spectroscopy capability of the current set-up:** A few examples of how improved γ -ray spectroscopy capability can raise the effectiveness of unfolding are given below. By choosing a gamma-ray detection set-up with higher absolute detection efficiency (e.g. by placing the foil closer to the detector or using a different detector), a larger number of counts of γ -rays can be detected (at the risk of increasing the rate of pile-up events in the gamma-ray detector). The detector electronics may also play a role: when different pre-amplifier pulse-shaping parameters are used, e.g. by reducing the decay time of the signal pulse (N.B. unrelated to nuclear decay), the saturation threshold count rate may be

increased at the expense of decreased energy resolution, increasing the maximum achievable γ -ray count rate while decreasing the capacity differentiating radionuclides with similar gamma-ray spectra from each other (hence reducing the range of radionuclides detectable), showing the trade-off between accuracy and precision. Using noise-reduction techniques such as Compton suppression can also reduce the error in the number of γ -rays counted, improving the unfolding precision. As the population of all radionuclides falls during the acquisition phase, the background level due to Compton scattered gamma-rays also reduces; if the gamma-ray spectrum is recorded in a time-resolved manner (e.g. recorded under list mode), then it is possible to take advantage of this reduction in background level to improve the uncertainty on the number of counts of the gamma-ray spectrum. Additionally, the time-resolved nature of list-mode acquisition allows better separation of counts due to different root nuclei, reducing the uncertainty on N_{meas} , as shown in Appendix H, increasing the precision.

To conclude the above section, seven design parameters may be optimised for each foil sample. This implies that the number of dimensions in the optimization space increases by seven for every additional foil, increasing the space of possible solutions that need to be checked exponentially, making the problem intractable as no solution may be generated in a reasonable time frame. To hedge against this curse of dimensionality, several of these design parameters can be fixed by the user, leaving only a handful of design parameters to optimise by the framework. These decisions will be explained in Sections 5.4 and 5.5. The user will have to decide on some design parameters manually (as opposed to allowing the framework to make these decisions automatically) before the framework begins. The user may take into consideration some practical concerns outlined in Section 5.3 to make these decisions.

5.3 Limitations on the feasible range of the experiment design parameters

This section concerns the leftmost part of the diagram in Figure 5.7, which lists factors that limit the range of design parameters that can be picked in Section 5.2.

1. **Experiment time:** This limits the length of the experiment schedule, and hence the length and number of irradiation-to-gamma-acquisition cycles that can be completed.

2. **Budget:** Budget constraints can limit the number of foils and types of materials that can be purchased. Making thinner foils is also technically challenging, so budget limitations may reduce the range of thicknesses available for purchase from suppliers.
3. **Nuclear data library used:** Not every nuclear data library is created equal. For example, FENDL [137] evaluates the most accurate nuclear data available for a smaller set of materials to fusion-relevant energies, while TENDL [138] tries to include nuclear data for as many materials as possible, filling theoretical values in places where there are no experimental values. Since microscopic cross section values for some reactions often differ wildly across libraries, showing more than a magnitude of difference, a different choice of nuclear data library will result in the framework concluding that a completely different set of foils is best for unfolding. If an incorrect choice of nuclear data library (one where the nuclear data quality is poor and the uncertainties are underestimated) is used, the unfolding will also proceed poorly to result in an inaccurate spectrum, because the response matrix will be inaccurately calculated.

Each library has its inadequacies, such as missing decay data, incomplete microscopic cross sections records (evident from the fact that the sum of all reactions' microscopic cross sections adds up to less than the total microscopic cross section, or having no data about microscopic cross sections above 20 MeV). If, in one foil, there are too many missing nuclides when building the decay chain, or very incomplete records of reaction cross sections of the reactants, then the experimentalist may choose to avoid using that material completely, reducing the range of usable materials.

4. **Neutron source characteristics:** A more unstable neutron spectrum would require a shorter irradiation time, to adhere to the assumption of having a constant (time-invariant) spectrum; and if a more intense neutron spectrum (i.e. neutron spectrum with higher flux) is used, then a shorter irradiation duration may be needed in order to not exceed the saturation threshold count rate of the γ -ray detector.

If the neutron source is an extended source rather than a unidirectional beam (see Figure 5.6), the maximum permissible thickness for the foil may be reduced as the user may wish to choose a lower $P_{\text{threshold}}$ in criterion 5.29 due to the increased path length variability. A neutron source with higher neutron energy will require thicker foils as the longer stopping distance of the more energetic nuclides produced requires a thicker medium

to stop in according to criterion 5.30, limiting the minimum permissible thickness for the foil.

5. **Materials handling capabilities:** foils that are extremely thin may be too fragile to be handled by the laboratory, rendering them unusable in unfolding experiments. This limits the range of usable thicknesses of foils.

Additionally, some materials may pose hazards to the handling personnel or the environment. Therefore the use of these materials may be restricted to only laboratories where expertise for handling the relevant hazards is available. An example list of these concerns is listed below:

- Chemical reactivity of that material, which affects its:
 - flammability,
 - and likelihood to react corrode/be corroded by the environment/container in which it is irradiated, including reaction with oxygen which may cause the foil to disintegrate after oxidation by exposure to air.
- Biological hazards, such as:
 - carcinogenicity,
 - and toxicity.
- Radiological hazards, due to:
 - fissile isotopes that occur after neutron irradiation of fertile/fissionable materials,
 - and any long-lived radioactive sources (e.g. Cobalt-60) produced by the irradiation.
- Whether the experiment set-up can tolerate *in situ* melting of the activation foil. If not, then the following constraint has to be taken into account as well:
 - the irradiation environment's temperature must remain below the solidus temperature of the material.

6. **Available human resources and experimental equipment:** Sometimes, whether a certain radionuclide can be detected, and how many of them can be detected, depends on the quantity and quality of the gamma-ray detectors, and the amount of human intervention required.

Access to more gamma-ray detectors allows for parallel irradiation and measurement schemes, such as those shown in Figure 5.8, which means a larger number of foils can be simultaneously irradiated and acquired.

Access to higher quality gamma-ray detectors can increase the detection efficiency, energy resolution, saturation threshold count rate, reduce the Compton background level, or even encode timing information into the gamma-ray spectrum. The impact of these effects is discussed in detail in Section 5.2 point 8.

If human intervention is required to extract the foil, place it into the gamma-ray detector and begin acquisition, the transfer time will be longer (and more inconsistent) than if this process were done in an automated fashion (such as the “rabbit” pneumatic transfer system [139, 140]). The workload required in these manual foil transfer cases is also higher, thus the complexity of the experiment schedule may be limited by the number of personnel available to assist with the experiment. For example, having two experiment personnel alone cannot effectively utilise both gamma-ray detectors simultaneously if frequent changes of foils are required. Therefore the available human resource limits the number of parallel irradiation-and-measurement cycles that can be carried out.

The feasible ranges of parameters will differ as some different constraints apply to different research groups. For example, some research groups have the capability to handle melting and re-solidification of low melting point foils (e.g. Pb, which has melting point 327°C, lower than the 800°C rated for the first walls of fusion reactors [22]) during the irradiation [58], as well as handle carcinogenic (e.g. Be) foils [6], expanding their range of available foils; while other groups may have access to time-resolved HPGe software [47] which allows radionuclides with short half-lives to be analysed more accurately.

5.4 Parameters fixed by the user

Some of the parameters shown in Figure 5.7 cannot be optimised programmatically and have to be manually decided by the experimentalist, to restrict the scope of the optimization problem and make it tractable. This section explains why the values taken by most of the parameters should be set manually; these reasons typically revolve around the need for *a priori* knowledge possessed only by the human operator, or having extremely large solution space that does not justify the programming time and effort required to find the optimal solution in it, especially when the *a priori* knowledge of the human operator can get a close-enough-to-optimal solution.




Figure 5.8: Consider a neutron source that can irradiate one foil at a time. If multiple detectors are available, a parallel schedule may be utilised to allow the gamma-ray spectra of multiple foils to be acquired concurrently (see top plot for example). Otherwise, if only a single detector is available, then the experimentalist is limited to a series schedule, acquiring the gamma-ray spectra of the previous foil while a new foil is being irradiated (see bottom plot for example).

1. **Experiment schedule**, namely the irradiation duration, transport duration and acquisition duration: There are many ways of arranging irradiation and measurement schedules. For example, two of the many possible ways of irradiating and acquiring the spectra for 10 foil samples are demonstrated in Figure 5.8. This goes to show that the solution space of possible schedules is incredibly large even for a small number of foil samples, and may look wildly different, so there are no trivial methods of converting these schedules into a parametric representation, or plugging them into a scalar-valued function that outputs an “unfolding efficiency”. As a result, there are no known ways of optimising the experiment schedule. Instead, we have to look towards computer science for help.

As each gamma-ray detector can only be used to acquire the spectrum of one foil sample at a time, arranging when to irradiate the foil, and for how long, becomes what is known as a “job-shop scheduling” problem in computer science, which is already a very complex field in its own right [141]. If experiment schedules need to be optimised, then existing job-shop scheduling algorithms must be modified to also account for the number of γ -ray counts detected. This will be a herculean undertaking, and since no suitable expertise is available in the physics community to provide the relevant training or assistance, a decision was made to sidestep this issue completely. Instead, the user of the foil selector framework will make up their own schedule, unaided by the computer program to decide the length of irradiation, transport and acquisition duration, thereby fixing these parameters in place and excluding them from the optimization process.

To minimise the effort required by the user when inputting the experiment schedule, and the complexity of the problem faced by the foil selector framework, all foil samples are assumed to have identical irradiation durations, transport durations, and acquisition durations. Therefore the user does not have to input multiple values of time a , b and c for each foil, and the framework assumes foils are interchangeable, so the order in which foil samples are irradiated and measured does not have to be optimised.

Obviously, users will have to account for the amount of time available for performing the experiments when choosing their values of irradiation durations, transport durations and acquisition durations to be fed to the foil selector framework. Additionally, they have to consider the minimum transport durations that are realistically and consistently achievable under their experiment set-up, and any relevant characteristics of their neutron source (such as the stability of the neutron spectrum, or maximum irradiation time of the

beamline) that may impose extra restrictions on the irradiation duration. The acquisition duration will be limited by the number of gamma-ray detectors available as well.

2. **Choice of the nuclear data library:** As explained in Section 5.3 point 3, the choice of which nuclear data library to use is entirely up to the user of the foil selector framework, as this is inputted as one of the data sources of the foil selector framework.

The foil selector program implemented in Section 5.6 can only check the completeness of the nuclear data library (i.e. how many isotopes are missing from the decay data library or whether the sum of cross sections is consistent with the total cross sections), but this is not the sole reason for why one nuclear data library should be preferred over another. Another reason is that the accuracy of nuclear data libraries cannot be evaluated without *a priori*/external knowledge, i.e. the foil selector program cannot claim whether one library is more accurate than another without consulting the scientific literature or the documentation behind the nuclear data libraries themselves. This task is more suited for human operators, i.e. the experimentalist to do. Therefore, it is preferable to let the user manually select what nuclear data library they would like to use, depending on what foils they are planning to use and what range of energy their expected neutron spectrum covers.

3. **Group structure:** This consideration was not included in Figure 5.7 as it does not strongly change the accuracy and precision of unfolding in the same manner; however, it still has to be specified from the start as the entire foil selection and unfolding process requires a consistent group structure.

Different group structures may lead to slightly different unfolding results due to intra-bin flux variation [78], and hence a different set of foils selected by the foil selector framework. Moreover, due to the assumption of having piecewise constant flux within each bin (Section 3.1.1), a finer group structure is required to capture the neutron spectrum variation in places where the scalar flux $\phi(\mathbf{r}, E; t)$ changes quickly w.r.t. energy. This requires *a priori* knowledge of where the steep features are expected to appear on the neutron spectrum. Therefore this is best chosen manually by the user at the beginning of the foil selector framework.

4. **Number of foil samples used:** The accuracy and precision of unfolding result can always be improved by adding one more foil into the foil set: if the additional foil has identical composition, irradiation duration, transport duration, and acquisition duration

of one of the existing foils, then it can help reduce the uncertainty by essentially repeating the measurement and confirming existing values of $\mathbf{N}_{\text{meas}}$, increasing the precision; else if the additional foil has a composition dissimilar to all of the foils seen in the existing foil set or have a much longer irradiation/acquisition durations, then new rows are added into the response matrix to represent the reactions that are found in the additional foil, hence increasing the degree of determination and increasing the accuracy of unfolding.

From this example, we can conclude that the efficiency of unfolding strictly increases with the number of foils used. If left alone to its own devices, the foil selector framework will keep trying to increase the number of foils indefinitely until it runs out of memory; and if a user sets a hard limit of “maximum number of foil samples” = ϑ that the outputted foil set may contain, the framework will simply output foil sets that is exactly equal to this limit every time, using the reasoning above. Considering that the framework will always output foil sets whose number of foils = the user-chosen integer ϑ , then there is no point for the framework to even begin evaluating the unfolding effectiveness of foil sets whose number of foils $< \vartheta$. Therefore, the framework asks the user to choose the required number of foils samples ϑ required in the output, and only evaluates the effectiveness for foil sets with exactly ϑ foils, reducing the space of solution that it has to search through as it now no longer has to consider the solutions with $< \vartheta$ foils.

5. **Foil compositions** (partially fixed only): The composition of foil can be parametrised in two ways: either as a list of isotopes and their corresponding atomic fraction (at.%) or as a list of chemical elements and their corresponding atomic fraction, where each element is assumed to have isotope composition equal to its natural abundances.

Each foil can be made of at least one isotope, but there are no upper limits to how many isotopes or elements it can consist of. This means that finding a foil with the optimal composition becomes an optimization problem with a variable number of dimensions, where the value of each dimension represents the atomic fraction of an isotope. Similar to the scheduling problem in point 1, exploring the problem of variable-dimension optimization would require mathematical expertise not readily available to the author and will be a sufficiently complex project in its own right that should not be included in the scope of this investigation.

Also, if the foil selector were granted absolute authority over the composition of the foil, the optimization result may produce a foil composition that is not manufacturable. For example, if the atomic fraction of isotope compositions were optimised, the foil selector

framework may output a solution foil composition that requires a large fraction of a rare isotope, which would be expensive, if not impossible, to manufacture via isotope separation. Even if this problem is avoided by choosing to optimise the elemental composition instead, the framework may still output foil compositions that are hard to manufacture, such as alloys that are unstable, or high entropy alloys that require rapid cooling rates to form. Theoretically, these problems may be curbed by adding extra constraints or additional terms to the objective function into the variable-dimension optimization algorithm, such as a function that penalises the use of dissimilar metals at similar proportions. However, this only serves to increase the complexity of this hypothetical variable-dimension optimization algorithm required to solve the foil composition optimization problem.

In light of these issues, the foil selector framework is not given free rein to optimise the foil composition. Instead, it works collaboratively with the user to choose the best set of foil that is realistically attainable by the user. The user provides a long list of materials and their compositions, which act as the **candidate** foils. The foil selector framework is only allowed to choose the materials for the finalised foil set from this list of candidate foils. Users can curate such a list of available materials from the catalogues of laboratory supply providers, or simply by listing existing materials in the laboratory storage cupboard.

An important benefit of this approach is that the user can remove all of the materials that cannot be used in their experiment set-up or handled by their research group, resolving any concerns regarding hazards mentioned in Section 5.3 point 5.

6. **Gamma-ray spectrum acquisition set-up:** The absolute detection efficiency, background noise level, energy resolution, dead-time following each pulse, and timing information are all controlled by the gamma-ray spectrum acquisition set-up. The experimentalist can adjust these quantities by changing the detection set-up: the absolute detection efficiency can be increased by reducing the separation distance between the foil and the detector crystal, at the cost of increasing the background level due to Compton and pile-up events, or by using a different type of detector; calibrating the voltage supplied to the pre-amplifier (if applicable) and amplifier, and the pulse-processing parameters in the analogue-digital converter (ADC), can affect the energy resolution and dead-time following each pulse; and using different gamma-ray spectrum acquisition software can allow incorporation of timing information into the spectrum.

These changes impact the quantity and quality of the information that can be extracted by unfolding: the more gamma-rays detected ($\mathcal{N}$) at the peaks with predetermined energies,

the smaller its fractional uncertainty is, and as a result a smaller fractional uncertainty can be expected of $\mathbf{N}_{\text{meas}}$, leading to higher precision. However, if the energy resolution is too low to separate neighbouring peaks from each other, or if the background level is high, then the uncertainty on the $\mathcal{N}$ may increase, or peaks may even have to be removed from consideration altogether, possibly reducing the rank of $\underline{\mathcal{R}}$, reducing the accuracy.

Much of this information has to be manually inputted by the user, as it is infeasible to incorporate a working model of a gamma-ray detector into the already complex foil selector framework. Therefore it cannot advise the user of the optimal detector type, settings for the ADC, voltage for the pre-amplifier and amplifier, and acquisition software that will produce the best results. Instead, the responsibility of finding the optimal gamma-ray spectrum acquisition set-up is delegated to the framework's user. The framework shall only prompt the user for inputs of relevant information about the gamma-ray detection set-up where necessary. These include a parameter denoting whether timing information is available or not, noise threshold (discussed in more details in greater detail in Section 5.6.4), a preliminary efficiency curve discussed in more detail in the paragraph below, as well as the saturation threshold count rate given by Equation 5.32.

$$\text{user-chosen threshold} \geq \frac{\text{total dead time}}{\text{acquisition duration}} \quad (5.31)$$

$$\text{user-chosen threshold} \geq \frac{(\text{dead time per pulse}) \times (\# \text{ of counts on the } \gamma\text{-spec})}{\text{acquisition duration}}$$

$$\begin{aligned} \frac{\text{acquisition duration}}{\# \text{ of counts on the } \gamma\text{-spec}} &\geq \frac{\text{dead time per pulse}}{\text{user-chosen threshold}} \\ \text{max. count rate} &= \frac{\text{user-chosen threshold}}{\text{dead time per pulse}} \geq \frac{\# \text{ of counts on the } \gamma\text{-spec}}{\text{acquisition duration}}. \end{aligned} \quad (5.32)$$

The preliminary efficiency curve functions as a trial solution to solve the interdependency issue (colloquially known as a “chicken-and-egg” problem), where the area and thickness of the foil cannot be decided (see the end of Section 5.6.3) until the γ -ray count rate per unit volume is calculated, which is dependent on the efficiency; but the efficiency curve cannot be precisely calculated unless the area and thickness of the foil are known. Using the preliminary efficiency curve in place of the precisely calculated efficiency curve, the cycle of interdependence is broken and the approximate required dimensions of the foil can be calculated.

Some users may be dissatisfied with the rigidity of the framework as it does not advise the user on the optimal experiment schedule or gamma acquisition set-up. To overcome these inflexibilities, users are encouraged to re-run the foil selector with changed parameters for any of the above six points, iteratively approaching the optimal set of parameters, effectively performing a kind of manual optimization. In particular regarding point 2, to patch the issues of missing isotopes in some nuclear data libraries, they can also splice together their own custom nuclear data library by copying the ENDF files (Evaluated Nuclear Data Files, [142]) of all isotopes they intend to use into a directory, and then providing that directory path to the foil selector program.

5.5 Parameters optimised by the framework

The only three factors left to be controlled by the foil selector framework are

- Choice of material
- thickness of foil used
- Foil volume

With practical concerns in mind, we aim to first tackle the optimization problem of these three factors in two stages.

The first stage is to choose the dimension of the candidate foils, in a semi-automated fashion, where a computer recommends the optimal ranges of values, while the user (the experimentalist) finalises the values to be chosen based on practical limitations.

1. The framework informs the user of the maximum permissible volume and range of permissible thickness of each candidate foil to avoid reaching the saturation threshold count rate and or violating criteria 5.29 or 5.30;
2. The user then determines the largest and thickest foil samples that may be purchased from a supplier or manufactured in-house, while staying below these limits, and adds the foil dimension (area A and thickness T) into the input JSON file. In the process, the user is effectively choosing and fixing the foil volume and thickness parameters, rather than letting the foil selector choose an arbitrary thickness and volume which may be hard to order or manufacture.

- One of the simplest ways the foil selector framework user could make this decision is to first choose the thickness T by rounding down from the maximum permissible thickness $T_{\max}$ to the next thickest foil thickness available for order. Then the foil area A may be decided by rounding down from the maximum permissible area $A_{\max} = \frac{V_{\max}}{T}$ to the next widest foil available for order.
3. The framework then simulates the irradiation and the subsequent gamma-ray spectroscopy measurement of these candidate foils in their user-determined dimension, evaluating this foil set's effectiveness at unfolding the neutron spectrum.

The second stage is for the framework to select a subset of best foils from the set of all candidate foils whose dimensions were specified by the user. Since the number of foils has been constrained in Section 5.4, it is a very well-defined problem of “choosing ϑ foils out of Θ candidate foils such that the resulting unfolded spectrum has maximum accuracy and precision”.

The choice of foil material is the most difficult, and hence least considered experiment design parameter, as it requires a tediously long computer program that loads the nuclear data, as well as a complete understanding of how to quantify the effectiveness of unfolding, which is not possible to present on short journal paper submissions. Hence this is likely the most important contribution to knowledge of this PhD.

5.5.1 Optimising for accuracy and precision

The precision and accuracy of the unfolded solution can be captured by the covariance matrix $\underline{\underline{\text{cov}}}(\phi_{\text{sol}})$, which can be calculated by borrowing the same procedure used in Equation 4.48⁵. Physically, this covariance matrix predicts the curvature of the χ^2 function (given by Equation 3.19) around any potential ϕ_{sol} that will be reached by unfolding the result of the activation foil gamma-ray spectra measurements.

In an ideal world, we would be able to convert the $\underline{\underline{\text{cov}}}(\phi)$ to a single scalar value for each foil set representing its effectiveness when applied to unfold a neutron spectrum, in a manner that captures both accuracy and precision. One natural choice would be the hypervolume of the $\chi^2 = 1$ region, which is a number proportional to the determinant $||\underline{\underline{\text{cov}}}(\phi)||$. For a 2-bin neutron spectrum, this “hypervolume” = the area encircled by the $\chi^2 = 1$ line (exemplified by Figure 3.1); and in a 3-bin neutron spectrum, this “hypervolume” is the volume enclosed by

⁵The Pseudo-Inverse algorithm is chosen because it makes the fewest assumptions about the neutron spectrum, as well as the simplicity of $\underline{\underline{\text{cov}}}(\phi)$.

5.5 Parameters optimised by the framework

the $\chi^2 = 1$ surface, etc. Using this measure, the best set of foils is then supposed to yield the smallest possible $\chi^2 = 1$ region hypervolume.

However, such a measure would only work in an overdetermined or fully-determined system. This is because the $\chi^2 = 1$ region's hypervolume is infinite (exemplified by the left plot of Figure 3.3) when it has at least one singular direction (i.e. it is underdetermined); and there is no easy ways of establishing whether the determinant of $\underline{\underline{\text{cov}}}(\phi_1)$ is larger than the determinant of $\underline{\underline{\text{cov}}}(\phi_2)$ or vice versa, if both determinants are infinite.

Yet, this does not spell the end of our attempt to find a suitable measure of the effectiveness of a set of foils. Consider two singular covariance matrices in ϕ space: $\underline{\underline{\text{cov}}}(\phi_1) = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ and $\underline{\underline{\text{cov}}}(\phi_2) = \begin{pmatrix} 4 & 2 \\ 2 & 1 \end{pmatrix}$. By inspection, it is clear that the curvature of the χ^2 function associated with the former system is stronger, as the latter represents a wider spread distribution of ϕ , since the standard deviation of ϕ_2 is twice that of ϕ_1 . In other words, there must exist some measure capable of quantifying the effectiveness of a set of foil at unfolding a neutron spectrum $\phi_{\text{true}} = \phi_0$ from examining $\underline{\underline{\text{cov}}}(\phi)$ alone, unaffected by whether or not it is singular.

Instead of trying to capture both “accuracy” and “precision” in a single scalar value, it was decided that one scalar measure be developed to represent each of these two quantities.

To quantify the accuracy, the following metric is used:

$$\mathbf{Accuracy} = \text{number of nonsingular directions of the } \chi^2 \text{ function} \quad (5.33)$$

$$= \text{rank}(\underline{\underline{\text{cov}}}(\phi)) \quad (5.34)$$

$$= \text{rank}(\underline{\underline{\text{cov}}}(\phi)^\dagger) \quad (5.35)$$

$$= \text{rank}(\underline{\underline{\mathbf{S}}}_\phi) .$$

If the ϕ_{true} were to deviate from ϕ_0 in any of the singular directions of $\underline{\underline{\text{cov}}}(\phi)$, the measured data $\mathbf{N}_{\text{meas}}$ would not be able to capture such deviation, and the unfolding algorithms will not be able to reproduce this deviation in the unfolded solution ϕ_{sol} , leading to an inaccurate solution. If $\underline{\underline{\text{cov}}}(\phi)$ is full-rank, then there are no directions where the foil set is “blind” to the deviation of ϕ_{true} from ϕ_0 . Therefore, the more nonsingular directions $\underline{\underline{\text{cov}}}(\phi)$ has, the closer to a full-rank matrix it is, and the less inaccurate the solution will be, lending itself nicely to be the proxy measure for “accuracy” for the foil set (Equation 5.35).

As for quantifying the precision, we turn to the ± 1 standard deviation envelope shown in Figure 4.13 (orange envelope) for inspiration. The smaller this envelope, the more confident

5.5 Parameters optimised by the framework

the unfolding program is with the neutron spectrum. The quantity shown in Equation 5.36 is roughly inversely proportional to the square of this area:

$$\int_{E_0}^{E_n} \frac{dE}{\text{var}(\phi(E))}, \quad (5.36)$$

where $\phi(E)$ is the scalar flux. The choice of $\frac{1}{\text{var}(\phi_{\text{sol}})}$ is a deliberate one, as it has a property of linearity that will be beneficial to the optimization below.

Since $\phi(E)$ is taken as a piecewise constant function:

$$\phi(E) = \frac{\phi_i}{E_i - E_{i-1}} \quad \forall \quad E_{i-1} < E \leq E_i, \quad (5.37)$$

Equation 5.36 can be written as:

$$\sum_{i=1}^n \int_{E_{i-1}}^{E_i} \frac{dE}{\text{var}\left(\frac{\phi_i}{E_i - E_{i-1}}\right)} \quad (5.38)$$

$$= \sum_{i=1}^n \int_{E_{i-1}}^{E_i} \frac{(E_i - E_{i-1})^2}{\text{var}(\phi_i)} dE \quad (5.39)$$

$$= \sum_{i=1}^n \frac{(E_i - E_{i-1})^3}{\text{var}(\phi_i)}. \quad (5.40)$$

This allows the precision to be quantified as:

$$\mathbf{precision} = \text{scalar curvature of the } \chi^2 \text{ function at } (\phi_0) \quad (5.41)$$

$$= \mathbf{w} \cdot \mathbf{Diag}(\underline{\underline{\mathbf{S}}}_{\phi}) \quad (5.42)$$

$$= \sum_{i=1}^n \left(w_i \left(\underline{\underline{\mathbf{cov}}}(\phi)^\dagger \right)_{ii} \right) \dots \quad (5.43)$$

The scalar curvature described by Equation 5.42 predicts how resilient the system is against noise: if the experimentalist were to repeat the unfolding experiment many times using the same set-up to obtain a list of unfolded solution ϕ_{sol} , this distribution of ϕ_{sol} 's would be spread out into some distribution; the spread of this distribution is then quantified by this number in Equation 5.42, lending itself to be the proxy measure for “precision” for the foil set excellently.

$\mathbf{w}$ is the weight vector, its function is to assign an importance value to the curvature in the i^{th} bin. Under the definition of piecewise constant $\phi(E)$ given by Equation 5.40, the weight vector $\mathbf{w}$ is defined as

$$w_i = (E_i - E_{i-1})^3. \quad (5.44)$$

5.5 Parameters optimised by the framework

However, if a logarithmic weighting scheme were to be used, where the user is more concerned with the accuracy of the flux scheme at low flux rather than high fluxes, then

$$w_i = (\ln(E_i) - \ln(E_{i-1}))(E_i - E_{i-1})^2.$$

Or if the flux were assumed to be piecewise-constant in lethargy space rather than linear ϕ space, then

$$w_i = (\ln(E_i) - \ln(E_{i-1}))^2(E_i - E_{i-1}).$$

Finally, if both modifiers apply (logarithmic weighting scheme and fluxes are assumed piecewise-constant in lethargy space), then

$$w_i = (\ln(E_i) - \ln(E_{i-1}))^3. \quad (5.45)$$

In general, the user can supply any positive vector $\mathbf{w}$ as the weight vector of their choice as long as this same $\mathbf{w}$ is reused for assessing the effectiveness of all candidate foil sets.

As explained in point 4 of Section 5.4, this problem now concerns finding the foil set made of ϑ foils, out of Θ candidate foils, making it a discrete optimization problem. The solution space would have a size $\binom{\Theta}{\vartheta}$, which is finite; however this number may grow to very large if the number of candidate foils considered is significant. There are at least 75 elements that are solid at room temperature with at least one stable isotope: if all of these elemental foils are considered, without allowing duplications or alloys, Θ is already = 75, and if the user aims to choose only $\vartheta = 6$ foils, then there are $\binom{75}{6} = 201359550$ combinations in the solution space. Therefore a strategy must be devised to search through this space to find the best solutions. This is where the linearity of both the **accuracy** metric and the **precision** metric becomes handy.

In general, the rank of the response matrix $\underline{\mathbf{R}}$ equals the number of types of root nuclides produced m , as microscopic cross sections are rarely linear combinations of each other; so when a new foil is added to a foil set, the rank of the response matrix of the foil set = rank(foil set) increases by the rank of the additional foil's response matrix = rank(new foil).

$$\mathbf{accuracy}(\text{foil superset with (foil set A + foil set B)})$$

$$\approx \mathbf{accuracy}(\text{foil set A}) + \mathbf{accuracy}(\text{foil set B}), \quad (5.46)$$

when $A \neq B$. This rule almost always holds when $m \ll n$, where the response matrix of foil set A has plenty of “unoccupied” (singular) directions for the singular vectors⁶ of the foil set B to inhabit. This is a useful heuristic to bear in mind, even if it isn’t a rigorous rule.

As for the **precision**, it is easier to prove its linearity. Consider two sets of foils A and B, with respective response matrices $\underline{\mathbf{R}}_A$ and $\underline{\mathbf{R}}_B$. If a new foil set (A+B) is constructed by combining these two sets of foils, the χ^2 of the new system is simply χ^2 of A + χ^2 of B. Consequently, the curvature of its χ^2 function, given by $\underline{\mathbf{S}}_\phi$ in Equation 5.42, can also be obtained by the sum of the $\underline{\mathbf{S}}_\phi$ of foil set A and the $\underline{\mathbf{S}}_\phi$ of foil set B.

To find the solution to the single-objective optimization problem of “Choose ϑ foils from the Θ candidate foils that give the highest accuracy/precision”, all that is needed is:

1. Create the response matrix for each of the Θ foils
2. Take the **accuracy/precision** of these Θ response matrices, forming a list of length = Θ .
3. Sort this list from left (smallest) to right (largest).
4. Take the ϑ rightmost elements of this list. Their corresponding foils form the optimal ϑ -foils set.

However, this is a multi-objective optimization problem, where both **accuracy** and **precision** need to be optimised. Therefore, the following algorithm is devised:

1. Create the response matrix for each of the Θ foils
2. Take the **accuracy** and **precision** of these Θ response matrices, forming two lists (referred to as **sorted_accuracy** and **sorted_precision** below), each of length = Θ .
3. Sort both lists from left (smallest) to right (largest).
4. Create two more empty lists, **most_accurate_combinations** and **most_precise_combinations**. Put the ϑ rightmost elements of **sorted_accuracy** as the first foil set in **most_accurate_combinations**; Put the ϑ rightmost elements of **sorted_precise** as the first foil set in **most_precise_combinations**.
5. Check if any foil sets are simultaneously present in **most_accurate_combinations** and **most_precise_combinations**. If so, terminate and return this match; otherwise, proceed to the next step.

⁶A singular vector is the generalization of an eigenvector in a full-rank square matrix.

6. Add the next-highest **accuracy** foil set into `most_accurate_combinations`; and add the next highest **precision** foil set into `most_precise_combinations`. Then return to the previous step.

This algorithm is crude and can be made more robust against any deviations from the approximation made by Equation 5.46 by ensuring at least several matches are found before the algorithm terminates. Having found more than one match across these two lists before terminating also gives the experimentalist more choices in what foils may be purchased under their budget constraint.

Due to the sorted-list nature of `sorted_accuracy` and `sorted_precision`, it is very straightforward to check for the next-highest **accuracy** foil and the next-highest **precision** foil: the algorithm only has to check at most $\min(\vartheta, \Theta - \vartheta)$ foil sets to determine which one is the next-highest accuracy/precision. The proof for this is trivial and will be left to the keen reader to do in their own time⁷. This turns the optimization problem into a tractable one: rather than blindly and randomly scanning the solution space, this algorithm scans the solution space in an ordered, top-down manner, beginning from the highest-performing combinations.

This method of optimization has not been proven to always arrive at the optimal solution; however, it is a useful heuristic to find a solution that is close to optimal, and is much better than random guesses, or having to search through the entire solution space.

5.6 Implementation as the foil selector program

To implement the foil selector framework described in the above procedurally, a Python program was created. Python was chosen due to the community support in nuclear data packages such as `openmc` [143], and the popularity of the language in general, which encourage open-source sharing and and contributing. The relative readability (compared to lower level languages such as C/C++) was also essential to the rapid development and testing of the program.

This foil selector program is currently written in the form of an interactive dialogue with the user, where the user is only required to have minimal background knowledge in unfolding, with the intention of lowering the barrier to entry for researchers. While the program is currently only available in a personal repository [132], it will be published as a software package with a documentation website and tutorial in the future, to ensure good research software and data quality.

⁷Those who wishes to understand this algorithm and prove that it works, may choose to represent the sorted list as a row of $\Theta - \vartheta$ 1's and Θ 0's.

5.6 Implementation as the foil selector program

The program is split into five steps, each performed by a single script, corresponding to the five subsections below. The user can plan their unfolding experiment by running the scripts on the same working directory sequentially, following the instructions below in Section 5.6.1 to Section 5.6.5.

5.6.1 Formatting *a priori* neutron spectrum and group structure

Table 5.1: Step 1 inputs and outputs.

Input	Output
• group structure (in any format)	• group structure (formatted): $E_i \forall 0 \leq i \leq n$
• <i>a priori</i> neutron spectrum (in any format)	• <i>a priori</i> neutron spectrum (formatted): ϕ_0

The program interactively prompts the user to input the file path containing the numerical values of the *a priori* neutron spectrum and its group structure and asks the user for further information about the inputted files, including:

- What format is the group structure provided as (mean-energy of each group, or the energy boundaries of each group), and with what unit (eV, keV, or MeV);
- What is the unit for the y-axis in the *a priori* neutron spectrum provided (# of neutrons per eV, per keV, per MeV, per unit lethargy, or integrated flux = total number of neutrons per bin);
- Whether the neutron spectrum provided is assumed to be piecewise constant or piecewise linear;
 - And if it is piecewise linear, what interpolation scheme is used (x-axis: linear or log; y-axis: linear or log).
- Any uncertainty associated with the neutron spectrum is also recorded by the program if available.

The program then re-bins the neutron spectrum into a final, user-chosen group structure. This user-chosen group structure would be reused throughout the rest of the steps, such that the response matrices and outputted neutron spectra are all given in the format of this group structure.

An option is also offered to the user to rescale the neutron spectrum, i.e. shifting the neutron spectrum curve horizontally or vertically on a log-log plot. This is particularly useful for

5.6 Implementation as the foil selector program

extrapolating the known neutron spectra of existing neutron sources to higher intensity/energy regimes to be used as the new *a priori* spectrum.

Finally, the following data are saved into the working directory for later steps to use:

- The neutron spectrum, as the “integrated” flux, i.e. recording the number of neutrons in each bin, with unit $\text{cm}^{-2}\text{s}^{-1}$. This would become ϕ_0 ultimately used in the unfolding process.
- The uncertainty (standard deviation of flux in each bin) associated with this *a priori* neutron spectrum ϕ_0 , if available.
- The group structure, recording the bin-boundaries in eV.
- The same neutron spectrum formulated as a function (taking in an input energy E in eV and outputting the neutron flux per unit in $\text{eV}^{-1}\text{cm}^{-2}\text{s}^{-1}$) using the `openmc.data.Tabulated1D` object.

5.6.2 Extracting from nuclear data libraries

Table 5.2: Step 2 inputs and outputs.

Input	Output
• composition.json (mixed format) • cross section data library: $\sigma_k(E)$ • decay data library: I_l • absolute detection efficiency curve: $\epsilon_{\text{abs}}(E)$ • minimum and maximum detectable gamma-ray energies (two scalar values) • group structure $E_i \forall 0 \leq i \leq n$	• .atomic_composition.json (formatted) • microscopic cross section of the relevant root nuclides (extracted to the correct group structure) σ_{ki} • decay data of all nuclides including their: –total number of gamma-rays detectable by the gamma-ray detector, per disintegration of that nuclide –gamma-ray “fingerprint”: $\epsilon_{\text{abs}}(E_l)I_l$ • maximum microscopic cross section: $\max(\sigma_k(E))$

The next step of the foil selector program is to establish the limitations on the feasible ranges of parameters (Section 5.3). The user provides the file path of a JSON file denoting the elemental/isotopic compositions of all candidate foils. The program reads and then formats

this JSON file to a standardised format that reflects the fraction of isotopes in each foil, such that the sum of all fractions adds up to 1. The user is also required to provide the directory or file path to the nuclear cross section library (MF=3 files) and radioactive decay libraries (MF=8 files). Using this information, the program takes the microscopic cross sections of only the relevant reactant isotopes necessary to build the macroscopic cross sections Σ_{ki} values for the production of root nuclides (which will later form the response matrices $\underline{\mathbf{R}}$ according to Equation 5.3) and rebins them into the user-chosen group structure set in Section 5.6.1. This selective extraction of nuclear cross section data is necessary as the nuclear data libraries are often quite voluminous, with e.g. TENDL-2019 sitting at 2.9 GB and ENDF/B-VIII.0 sitting at 438 MB [144]. Hence loading the cross sections of every single nuclide in the nuclear data library would overburden the RAM and swap space typically available on a modern desktop or laptop, leading to a memory error. By only reading the nuclear data library file by file, and only saving the relevant nuclide’s cross sections and decay data, the load time can be vastly reduced and the memory usage will drop to a manageable level. For users familiar with FISPACT-II, this sub-step is analogous to the “COLLAPSE” (of the cross sections) step.

Before the rebinning of the microscopic cross section, the maximum microscopic-cross section $\max_E(\sigma_k(E)) \forall E_0 < E \leq E_n$ is also recorded for each of the reactions. This is used to aid the calculation of the maximum permissible thickness using criterion 5.29, and the detailed procedure is explained in Appendix I.

The decay data is extracted from the provided library file/directory paths as well. The FISPACT-II analogue of this sub-step is the “CONDENSE” (of the decay information) step. While it would be memory efficient to extract the decay information (decay constants, decay products, and gamma-ray spectra) of only nuclides that may appear during the experiment, it would require the program to identify all relevant nuclides by building the decay chains from each root nuclide, which is an intensive computation reserved for the next step (Section 5.6.3). Therefore, the program simply extracts from the decay library the decay information of all nuclides with mass number $M \leq \max(M)$, where $\max(M)$ is the mass number of the heaviest root nuclide. This scrapes more data than is strictly necessary, as most of those scraped nuclides will not occur after the irradiation of the candidate foils. This is an acceptable amount of memory wasted as the decay data libraries are often quite small compared to the cross section libraries, where e.g. decay_2012 is 55 MB in size [144]. By discarding any non-gamma-ray related decay information, and converting the plain text numerical data into floats, the decay data can be condensed to a much smaller size as well, circumventing any concerns about memory error.

Assuming that an identical set-up of gamma-ray spectroscopy would be used for all foils and that all foil sample sizes are roughly the same, then the absolute detection efficiency curve for each foil would be the same. Since the foil selector program only needs to know the total number of counts for each reaction (justified in Section 5.6.3), we only need to store the total number of counts detectable per decay of each nucleus, rather than the full list of gamma-ray intensities. This is obtained by the dot product of the gamma-ray intensities with the efficiency curve:

$$\text{counts per decay of nucleus} = \sum_l^L \epsilon_{\text{abs}}(E_l) I_l. \quad (5.47)$$

This reduces the computation needed in the following step (Section 5.6.3), as otherwise this dot product in Equation 5.47 would have to be recomputed each time the γ -ray count rate is needed.

Since an efficiency curve is required, the user can either input an efficiency curve in one of the accepted formats (MCNP output, Genie’s ISOCS output, or a simple CSV file with energy as the 1st column and absolute efficiency as the 2nd column.) in this step, or choose from one of the two default NaI/ HPGe efficiency curves. The program also accepts two numbers as to the argument `--gamma-energy-limits` which is used as the cut-off minimum and maximum gamma-ray energies for the efficiency curve in keV. The user is expected to use the calibration curve of the gamma-ray detector planned for the experiment to calculate the gamma-ray energies corresponding to the lowest and highest bin of the gamma-ray spectrum, and enter them into these two argument slots. The efficiency curve is effectively set to zero outside of this range.

The rebinned collapsed microscopic cross sections, the condensed decay information, and the maximum σ_{ki} are saved for use in the next step. The gamma-ray “fingerprint” (vector of element-wise product between the list of gamma-ray intensities and the efficiency curve, given by Equation 5.48) of each nuclide is saved as well, for use in Section 5.6.5:

gamma-ray fingerprint of the disintegration of a nuclide of X

$$= \epsilon_{\text{abs}}(E_l) I_l \quad \forall \quad \text{gamma-lines } l \text{ of nuclide } X. \quad (5.48)$$

One technical detail about nuclear data is that the cross section for each channel of reaction is stored in a section of a file that is assigned an identifier called an “MT number” [145]. Usually, each MT number’s reaction produces one and only one nuclide. But in some cases such as MT= 5 (which corresponds to a (n,misc.) reaction), no products are specified; these reactions can be ignored. In other cases where multiple nuclides may be produced, the ratio of how many

5.6 Implementation as the foil selector program




Figure 5.9: Recovering the microscopic cross section of production of two radionuclides from an MT= 16 ENDF entry using the MF=9 file using Equation 5.49.

of each nuclide is produced is given by an energy-dependent fractional yield value. The “partial” microscopic cross section of production of each nuclide can be recovered by the equation

$$\sigma_k(E) = \text{total cross section value stored at the relevant MT file at } (E) \times \text{fractional yield for the } k^{\text{th}} \text{ root nuclide at } (E). \quad (5.49)$$

An example of this reconstruction process for the production of ^{26}Al and its metastable state is shown in Figure 5.9: only the black curve is stored in the MT= 16 file. But by multiplying the black curve to the fractional yield of ^{26}Al , the **blue** part of the curve is obtained; and similarly multiplying by the fractional yield of ^{26m}Al gives the **orange** part of the curve.

The sum of fractional yields at the energy E over all products should always equal unity. The program logs all cases where the sum does not equal unity and prints a warning to the screen.⁸

⁸Such warnings only occurred when the total cross section values are zero and the fractional yields all drop to zero, either as a result of missing nuclear data or having zero available nuclear data. In either case, the lack of normalization of the fractional yield will not cause any problems as $\sigma_k(E)$ in Equation 5.49 will be zero.

5.6.3 Calculating number of counts

Table 5.3: Step 3 inputs and outputs.

Input	Output
• irradiation duration a (a scalar value) • transport duration $b - a$ (a scalar value) • acquisition duration $c - b$ (a scalar value) • <code>--max-gamma-count-rate</code> (a scalar value) • <i>a priori</i> spectrum: ϕ_0 • microscopic cross section σ_{ki} • decay data of all nuclides, i.e.: –total number of gamma-rays detectable by the gamma-ray detector per root nuclide produced, in a CSV file • <code>..atomic_composition.json</code>	• experiment schedule (a, b, c) • total number of decays, activity and count rate of all descendant • <code>number_of_atoms_upper_limits.json</code>

The third step is to establish the parameters to be fixed by the user (Section 5.4).

For the production of every root nuclide, the total number of decays found during the acquisition duration is obtained by integrating the activities from the beginning of the acquisition ($t = b$) to the end of the acquisition ($t = c$) using Equation 5.19, by using the previously condensed decay data. This yields several quantities:

- activity of all radionuclides at $t = b$ per root nuclide produced
- activity of all radionuclides at $t = c$ per root nuclide produced
- activity of all radionuclides at $t = c + 1$ hour per root nuclide produced (used for radiological clearance and disposal purposes)
- detector count rate due to all radionuclides in the foil at $t = b$ per root nuclide produced
- detector count rate due to all radionuclides in the foil at $t = c$ per root nuclide produced

These quantities are then saved as a CSV file which can be opened in any generic spreadsheet software (e.g. Excel). They are useful references to have on hand when performing the actual experiment, where unexpected changes may be required, and the irradiation and acquisition durations can be shortened or lengthened using the information in this CSV file.

5.6 Implementation as the foil selector program

This step of the program also requires a `--max-gamma-count-rate` (also referred to as the saturation threshold in this thesis, see Appendix L.9) argument input from the user. This limits the number of reactants that may be present in the foil at time $t = 0$, as a foil with too many reactants will have a count rate too high that leads to the gamma-ray detector's deadtime to be too high and causing its gamma-ray spectrum to be degraded and the detection efficiency to decrease in an uncontrolled manner.

The program calculates the maximum number of reactants that may be used using the following equation:

$$\begin{aligned} & \text{max. number of atoms permitted to be present in the foil at } t=0 \\ &= \frac{\text{max. gamma-ray count rate}}{\sum_{\substack{\text{all nuclides produced} \\ \text{in that candidate foil}}} (\text{number of } X) \times (\text{max detector count rate per nuclide } X \text{ in the foil})} \end{aligned} \quad (5.50)$$

$$= \frac{\text{max. gamma-ray count rate}}{\sum_k \substack{\text{all nuclides } k \text{ in} \\ \text{that candidate foil}} N_k \times (\text{detector count rate at } t = b \text{ per root nuclide } k \text{ produced in the foil})}, \quad (5.51)$$

The denominator can be calculated using the information obtained from Equation 5.47 from the previous section. Here, the program assumes that the count rate is highest during the beginning of the acquisition period ($t = b$). The number of root radionuclides N_k produced during irradiation is approximated by folding $\underline{\mathbf{R}}$ with ϕ_0 . The response matrix $\underline{\mathbf{R}}$ is obtained as the product of the microscopic cross section with the `atomic_composition.json` as stated in Equation 5.3, and then multiplied by a , as shown in Equation 5.8.

The information of the maximum number of atoms that can be used for each candidate foil (Equation 5.50) is saved in a file named “`number_of_atoms_upper_limits.json`”. The user is expected to manually open `number_of_atoms_upper_limits.json` and calculate the maximum permissible volume and maximum thickness, using

$$V_{\max} = \frac{\text{max. number of atoms}}{\text{number density of atoms}}, \quad (5.52)$$

to calculate the maximum volume, and criterion 5.29 to calculate the maximum thickness (using the maximum microscopic cross section saved at the end of the previous step Section 5.6.2). These calculations of maximum volumes are not automated because the number density of atoms N_d is variable depending on the method of manufacturing.

The user is expected to search the laboratory supplier’s catalogue for the largest possible foil satisfying these limits that can be purchased or fabricated under the experiment budget and time constraint, and then enter the total number of reactant atoms in that maximum-sized foil into a JSON file, to be used in the next section. The foil selector program assumes each candidate foil entry in the JSON file may only be used once. If the user would like to allow for duplicate foils to be used in the outputted foil set, then the corresponding entries in the JSON file must be repeated.

During the decay-chain-building step outlined in Section 5.1.1, the incompleteness of nuclear data poses some challenges. One of the challenges is that sometimes, a nuclide with no valid decay data is supposedly produced after the irradiation. Typically, these are radionuclides in excited states. The current solution to this challenge is to assume that this isomer of the radionuclide is too fast decaying to have a record in the nuclear data library, and would decay to the ground state immediately with no gamma-ray emitted in the detectable energy range. Thus the program defaults to replacing that isomer with the ground state isomer. If the ground state of that radionuclide also has its decay record missing, then the program replaces that nuclide with a placeholder NULL nuclide, which is presumed to be stable and emits no radiation, and prints a warning to screen to prompt the user to manually check that this nuclide is indeed one that is not expected to emit any detectable radiation.

5.6.4 Finding the optimal foil set

Table 5.4: Step 4 inputs and outputs.

Input	Output
• final_candidate_foils_dimensions.json: • target number of foils required in the foil set: ϑ • <i>a priori</i> neutron spectrum • microscopic-cross section σ_{ki} • experiment schedule (a, b, c) • decay data of all nuclides, i.e.: –their gamma-ray “fingerprints”: $\epsilon_{\text{abs}}(E_l)I_l$.	• solution foil set (first match between both lists) • list of foil sets with highest precision • list of foil sets with highest accuracy • Effective response matrix: $\underline{\mathcal{R}}$

5.6 Implementation as the foil selector program

The fourth step is the most important one: finding the foil set that optimises for both objective functions stated in Section 5.5.1.

Reading in the `selected_candidate_foils_dimensions.json` consisting of the masses of the foils that the user had decided upon, the finalised number of reactants in each candidate foil can be calculated. Using the microscopic cross section and the irradiation duration a from the experiment schedule, the response matrix can be calculated using Equation 5.3 in the same manner as in the previous step. The matrix $\underline{\underline{\mathbf{G}}}$, as defined by Equation 5.21 is also calculated using the decay data previously saved from step 2 (the gamma-ray “fingerprint” from Equation 5.48). This effectively assumes that the same absolute detection efficiency curve is applied on all foils as the efficiency curve inputted in step 2, independent of the actual dimensions of the foil as provided by the user. This assumption is necessary to reduce the amount of inputs the user has to make, as they would otherwise have to input Θ efficiency curves, one for each foil used. The effective response matrix $\underline{\underline{\mathbf{R}}} = \underline{\underline{\mathbf{G}}}\underline{\underline{\mathbf{R}}}$ is then constructed for each foil. The effective response matrices are then folded with the *a priori* neutron spectrum ϕ_0 to obtain the effective response vectors $\mathcal{N}$, the number of gamma-ray expected to be counted at each peak of the gamma-ray spectrum of each foil.

Not all gamma-ray peaks l are detectable. Some peaks have counts $\mathcal{N}_l$ which are too low to be detectable, as they are drowned out by the background noise. To determine how big must $\mathcal{N}_l$ be for the peak to be detectable requires knowledge of the background noise level of the gamma-ray spectrum at its photopeak energy $= E_l$. It is difficult to obtain this background level without a working model of the detector, as they are a sum of the Compton continuum and other complicated artefacts (pile-up events and electronic noise) that are dependent on the number of gamma-rays from other peaks. Since it would be unreasonable to expect the foil selector program’s user to have a response matrix of the gamma-ray detector⁹ planned to be used in their experiment on hand, the foil selector instead makes a simpler assumption to model the background noise in a gamma-ray spectrum: An energy-independent Compton background level proportional to the total number of photopeak counts is expected. Any peaks with counts fewer than a predetermined fraction of this background will be removed from consideration. Expressed mathematically, for gamma-ray peaks to survive this selection process,

$$\mathcal{N}_l > 2\% \times \sum_{\zeta}^L \mathcal{N}_{\zeta}, \quad (5.53)$$

⁹Not to be confused with to the response matrix of the foils, this response matrix converts the incident gamma-rays into detected number of pulses in the gamma-ray detector.

where the arbitrarily chosen factor of “2%” is an admittedly inelegant way to summarise the predetermined fraction and peak-to-Compton ratio to form a single multiplicative constant, which is then multiplied on the total number of counts to form the threshold. Obviously this assumption is inaccurate, thus it shall be improved further as explained in Section 5.7.

All peaks not meeting criterion 5.53 are assumed to be undetectable, and the corresponding rows on the $\underline{\underline{\mathbf{G}}}$ matrix are deleted, reducing the number of rows (and potentially the rank) of $\underline{\underline{\mathbf{R}}}$. All the rest of the peaks are assumed to take a simple, uncorrelated Poissonian distribution, i.e. $\text{var}(\mathcal{N}_i) = \mathcal{N}_i$. This allows $\underline{\underline{\text{cov}}}(\mathcal{N}) = (\underline{\underline{\mathbf{S}}}_{\mathcal{N}})^{-1}$ to be approximated as $\underline{\underline{\text{Diag}}}(\mathcal{N})$, and following Equation 4.47, this yields $\underline{\underline{\mathbf{S}}}_{\phi} = \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathcal{N}} \underline{\underline{\mathbf{R}}}$.

Then, armed with the $\underline{\underline{\mathbf{S}}}_{\phi}$ matrix from each foil, the multi-objective optimization algorithm set out in Section 5.5.1 can be carried out. This yields two lists: `most_accurate_combinations` and `most_precise_combinations`. Both lists are returned, as well as the first match between them, which is the solution foil set. Should the user wish to deviate from the solution foil set suggested by the foil selector program, then these two lists are invaluable references to help them decide what alternatives are useful.

The effective response matrix $\underline{\underline{\mathbf{R}}}$ for each foil is then saved for use in the next step.

5.6.5 Prediction of the gamma-ray spectrum (optional)

Table 5.5: Step 5 inputs and outputs.

Input	Output
• Selected foils (JSON file) • Effective response matrix: $\underline{\underline{\mathbf{R}}}$ • resolution curve of the gamma-ray detector: $R(E)$ • <i>a priori</i> neutron spectrum: ϕ_0	• expected gamma-ray spectra of each candidate foil

This step is an add-on tool to the foil selector. While its purpose extends beyond the original scope of “selecting the optimal set of foils”, it performs tasks in the experiment planning stage that can benefit from reusing the data previously loaded by the foil selector.

During the gamma-ray spectra acquisition phase(s) of the experiment, it is useful to have a copy of the “expected gamma-ray spectra” so experimentalists can confirm they have set up their experiment correctly, and understand how the measured gamma-ray spectra $\mathcal{N}_{\text{meas}}$ differ from the expected gamma-ray spectra $\mathcal{N}_0 = \underline{\underline{\mathbf{R}}}\phi_0$. They may also use this gamma-ray spectrum

to understand where the important peaks are expected to appear, such that the amplifier/ADC settings may be adjusted to maximise the resolution in the gamma-ray energy range of interest. Combined with the spreadsheet of gamma-ray count rates produced at step 3 (Section 5.6.3), they can also adjust the position of the activation foil in the gamma-ray detector to maximise the efficiency without reaching the saturation threshold of the gamma-ray detector.

5.7 Future development plan for the foil selector program

The development of the foil selector program had thus far been the endeavour of a sole author with little external aid for software testing and version management. While some assistance from gamma-spectroscopy experts (Chantal Shand and Alison Bruce, named in the supervision team of this thesis) helped formulate steps 4 and 5 of the foil selector program, and experiments opportunities (Chapter 6 of this thesis and [146]) helped refine the program such that it is more suited to real world use cases, more work is needed to improve the code quality to a level suitable for other research groups to use, such that it may be downloaded and used with no prior introduction.

The most important plan for the foil selector program is to improve the user experience, which will affect its usership. In its current state, it is user-unfriendly as the five steps are stored as five separate scripts, all of which require manual input to the command line, as its development (in parallel with the development of the foil selector framework) was an immense undertaking. The need for frequent user inputs suggests that it may benefit from having a graphical user interface (GUI). With the appropriate resources and manpower (beyond what is available to this PhD), it would be desirable to give the foil selector program a GUI, as well as distribute them as installable .exe files on Windows as well, which is the operating system of choice for many research groups. This would lower the barrier to entry for junior researchers, encouraging better practice when selecting foils for unfolding experiments.

Another benefit of re-writing the foil selector program as a GUI-based program is that data can persist across steps. Currently, the collapsed and condensed nuclear data, as well as various miscellaneous user-inputted information, persist in the same directory in the form of CSV files, Python pickled objects and JSON files even after the run, which takes up memory unnecessarily. If the program can be launched as a GUI, then these data can simply be held in RAM/Swap-space memory (or stored as temporary files if they are too big) and deleted when the program deems them unnecessary/ upon closing the program, reducing the amount of junk files created, and increasing the ease of maintenance.

With this increased ease of maintenance as a GUI-based program, it would be easier to add extra functionalities to the foil selector program. Between step 3 (Section 5.6.3) and step 4 (Section 5.6.4), the manual conversion from number densities (an unintuitive quantity) to volumes V and thicknesses T (more tangible quantities, ones which would be inputted into the purchase orders) has proven to be a very tedious task for the user. In the future, the program can optionally prompt the user to enter a “density” (in g/cm^{-3}) value for each candidate foil via the GUI (currently inputted from `composition.json`), such that the outputted upper limits of dimensions (currently outputted in `number_of_atoms_upper_limits.json`) would also include the upper limits on the foil volumes V and thicknesses T .

To further improve the maintainability and reliability of the foil selector program, some restructuring will be required as well. Due to its scientific and exploratory nature (being developed in parallel with the foil selector framework), it was not possible to develop the software in a test-driven manner. Therefore the foil selector program can benefit from having unit tests and regression tests retrospectively added. It would also benefit from removing its dependency on OpenMC, as that was not the correct package for the task at hand: it was a fairly voluminous Python package (100 MB) with a main focus on neutron transport; its nuclear data reading capability was not well maintained, such that some nuclear data formats (e.g. `.asc` files) could not be read by the version of OpenMC used in the foil selector program development environment, and some properties of the nuclear data libraries cannot be accessed (MF=33 files, i.e. reaction cross section covariance files). This limited the scope of what could be done by the foil selector program, as it cannot select specifically for foils with less uncertainty in their nuclear data. Instead of OpenMC, a new package known as `endf-parserpy` [147] is identified as a package that is potentially more suited to the current task and will take up much less space (6.7 MB), allowing a much smoother installation of the foil selector program. Therefore, one of the next tasks is to look into whether the OpenMC dependency can be replaced by `endf-parserpy`.

In Section 5.6.4 Equation 5.53, an energy-independent count threshold was used as the cutoff value for deciding whether to retain or discard peak from for further analysis. This has proven to be ineffective, as in Chapter 6, the background level due to Compton scattering at > 1000 keV is typically on the orders of single-digit counts per channel, while the background level due to Compton scattering at < 50 keV is on the orders of 4 digits per channel; therefore a different threshold would have to be used to weed out usable peaks from unusable peaks. Therefore, the following improvement is proposed:

The foil selector shall be modified to make a crude simulation of the gamma-ray detector to create a more realistic Compton background spectrum Γ . The user shall be prompted to

5.7 Future development plan for the foil selector program

input two curves: a reference peak-to-Compton ratio curve $\kappa(E_\zeta)$, and the resolution curve $R(E_l)$. Both curves can each be specified by two to three points. The peak-to-Compton ratio is expected to be a first- or second-order polynomial (with positive derivatives) with respect to E_l that outputs the amount of noise that would be created on the gamma-ray spectrum for every γ -ray detected at energy E_l , a dimensionless quantity. Meanwhile, the resolution curve $R(E_l)$ is expected to be a curve of the form $c_1 + c_2\sqrt{E_l}$ or $\sqrt{c_1 + c_2E_l}$, outputting the FWHM of a photopeak at energy E_l , with the unit keV.

The background level approximation function Γ can then be constructed by assuming that these noise generated by each photopeak at gamma-ray energy $= E_\zeta$ is smeared out into a flat distribution to the left of the Compton edge $(E_\zeta \left(1 - \frac{1}{1 + \frac{2E_\zeta}{511\text{keV}}}\right))$, forming the Compton continuum:

$$\Gamma(E_l) = \sum_{\substack{\text{All values of } \zeta \text{ such that} \\ E_\zeta \left(1 - \frac{1}{1 + \frac{2E_\zeta}{511\text{keV}}}\right) > E_l}} \frac{\kappa(E_\zeta)}{E_\zeta \left(1 - \frac{1}{1 + \frac{2E_\zeta}{511\text{keV}}}\right)} \mathcal{N}_\zeta. \quad (5.54)$$

Γ has a unit = counts keV⁻¹. Intuitively, what this function is doing is to account for all of the gamma-rays “scattered” as noise from higher energies to the energy of interest.

A peak at E_l must protrude significantly above the Compton background $\Gamma(E_l)$ to be detectable. For example in Figure 5.10, the **green** peak (peak 3 at 230 keV), which is sitting on the Compton continua of the two higher energy peaks (sum of the **blue** continuum and the **orange** continuum), may be hidden away by the fluctuations of these continua, therefore an experimentalist looking at the spectrum will not be able to confidently identify the blip at $E_l = 230$ keV as a photopeak.

To more formally establish the criterion for when a peak can/can't be detected, we need to compare the number of counts it contains against the number of counts over the Compton background noise that it straddles. For our purpose, only counts within the peak's FWHM are considered.

$$\begin{aligned} \text{total noise under peak at } (E_l) &= \int_{E_l - \frac{\text{FWHM}(E_l)}{2}}^{E_l + \frac{\text{FWHM}(E_l)}{2}} \Gamma(E_l) dE \\ &= \int_{E_l \left(1 - \frac{R(E_l)}{2}\right)}^{E_l \left(1 + \frac{R(E_l)}{2}\right)} \Gamma(E_l) dE; \end{aligned} \quad (5.55)$$




Figure 5.10: A hypothetical (simulated) gamma-ray spectrum. The sum of the Compton continua on the left forms the Compton background noise level $\Gamma(E)$ defined by Equation 5.54. The locations of the peaks are denoted with Gaussian peak shapes.

5.7 Future development plan for the foil selector program

Meanwhile, assuming the photopeak follows a Gaussian distribution, the total number of counts recorded by bins within the range of the FWHM is approximately equal to:

$$\text{total counts within FWHM} = \int_{E_l - \frac{\text{FWHM}(E_l)}{2}}^{E_l + \frac{\text{FWHM}(E_l)}{2}} (\gamma\text{-rays per keV due to the } l^{\text{th}} \text{ photopeak}) dE \quad (5.56)$$

$$= \left(\int_{-\sqrt{\ln(4)}}^{\sqrt{\ln(4)}} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \right) \mathcal{N}_l = \text{erf}(\sqrt{\ln(2)}) \mathcal{N}_l \quad (5.57)$$

$$= \left(\int_{-\frac{2.35}{2}}^{\frac{2.35}{2}} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \right) \mathcal{N}_l \quad (5.58)$$

$$= 0.76097 \mathcal{N}_l \quad (5.59)$$

Assuming the distribution of the Compton background noise level is Poissonian in nature (i.e. Fano factor = 1), and aiming for a 3-sigma confidence level (99.7% confidence, i.e. 3 in 1000 chance of mistaking random fluctuations in the Compton background as a peak), the condition for accepting peak l is given by

$$0.76097 \mathcal{N}_l \geq 3 \sqrt{\int_{E_l \left(1 - \frac{R(E_l)}{2}\right)}^{E_l \left(1 + \frac{R(E_l)}{2}\right)} \Gamma(E_l) dE}. \quad (5.60)$$

Any peaks l with $\mathcal{N}_l$ that does not satisfy this shall be removed from the effective response matrix $\underline{\mathcal{R}}$ in step 4 (replacing Equation 5.53), removing the l^{th} row of the effective response matrix $\mathcal{R}_l$ such that the $\mathcal{N}$ more accurately represents the list of peaks that can be measured experimentally. As for peaks that pass criterion 5.60, the variance on their expected number of

counts $\mathcal{N}_l$ shall increase by $\int_{E_l \left(1 - \frac{R(E_l)}{2}\right)}^{E_l \left(1 + \frac{R(E_l)}{2}\right)} \Gamma(E_l) dE$, which shall impact the **precision** as evaluated by Equation 5.42.

There are some assumptions about this new peak-acceptance scheme that deserve further attention:

1. The background noise level Γ is assumed to be a flat Compton continuum in Equation 5.54.

While this is a better approximation than the flat background noise distribution in Equation 5.53, it is still incorrect, as the Compton continuum is a shape that varies with

the geometry of the detector crystal. It also overlooks any effect of coincidences (which may cause noise to appear in higher energy bins than $E_\zeta \left(1 - \frac{1}{1 + \frac{2E_\zeta}{511keV}}\right)$), backscattered γ -rays from the measurement environment, single-escape, and double-escape peaks. However, this is a necessary assumption because this is the closest approximation of the gamma-ray detector's full spectrum without asking for the full response matrix of the specific gamma-ray detector that the experimentalist plans to use (which is something the average foil selector program user would not have access to). And this assumption is not very far from the truth, as the Compton scattered γ -rays account for the majority of the background noise, forming an approximately flat continuum.

2. The “fingerprint” in step 2 (Equation 5.48) extracted also does not account for cascade summing coincidence in its hypothetical gamma-ray spectrum $\mathcal{N}$. This may lead to peaks being rejected as their respective peaks do not reach the threshold required in criterion 5.60 despite the overall cascade-sum peak passing that criterion, leading to an underestimation of the effectiveness of the relevant foils.
3. The gamma-ray spectrum is discrete in nature, so the integrals in Equations 5.55 and 5.56 becomes summations over all bins that span the FWHM of the peak with energy at its centroid $= E_l$. If the FWHM is too small compared to the bin widths, then there is a risk that more noise would be summed into each bin than suggested by Equation 5.55, such that it may still be undetectable even if it has passed criterion 5.60.

Finally, there are still some open questions posed at the end of Appendix H.1 and I.1 that require further investigation; answering these would make the foil selector framework more robust. Answering the latter question will also allow the foil selector framework to be simplified; and this change in the framework will propagate down to the foil selector program.

Chapter 6


Experimental validation

To demonstrate the use of the foil selector framework created in Chapter 5, an experiment was devised; then the `unfoldingsuite` created in Chapter 4 was used to unfold its result to reconstruct the neutron spectrum.

The experiment was carried out at the ChipIr beamline [148] of the ISIS spallation neutron source, where protons are accelerated up to high energies by a linear accelerator and injected into a synchrotron, in which they are accelerated to even higher energies (up to several GeVs) and then released in pulses to hit a tungsten target (Figure 6.1a). The protons interact with the tungsten nuclei in spallation reactions to release many neutrons whose energies are dependent on the angle relative to the proton beam. ChipIr’s beamline is angled in such a way (Figure 6.2) that its neutron spectrum resembles that of an atmospheric cosmic ray neutron source, having neutron energies up to 700 MeV. Its atmospheric neutron spectrum has applications in aerospace electronics, where components undergo accelerated testing for single-event effects (SEE) induced by high-energy neutrons like those encountered in high altitudes. The high neutron energies render most measurement techniques ineffective (e.g. In this facility, time-of-flight can only measure up to 10 MeV [149] as explained by reference [49], while Bonner Spheres are limited in resolution [150]), thus this presents an opportunity to showcase the unique capabilities of neutron spectrum unfolding with activation foils: it was hypothesized that with the correct selection of foils, the neutron spectrum of ChipIr can be measured precisely and accurately.

6.1 Experiment design

The foil selection was done using a prototype of the foil selector program (Section 5.6), as this experiment was planned concurrently with the development of the foil selector framework. There are some minor discrepancies between the input and output of this prototype version



(a) Front view of the replica. The tungsten target itself is the cylinder with the shiny end to the right and above the finger in the picture; the rest of the stainless steel around the target are water tanks to hold coolant water. The beamline strikes the target in the direction parallel to the direction pointed at.



(b) The aperture that guides fast neutrons to the ChipIr irradiation chamber. This photo is taken at approximately 120° (anticlockwise) offset to the direction of the photo in Figure 6.1a

Figure 6.1: Photos of the replica of the spallation target used in target station 2.

of the program compared to the input and outputs of its contemporary version described in Section 5.6; where these discrepancies make a significant difference in the result obtained by this foil selection process, it would be highlighted explicitly in the text below.

The first step (Section 5.6.1) is to select a group structure and *a priori* spectrum. The CCFE-709 group structure [152] was chosen due to its well-established status, despite the possibility that it may be suboptimal [78]. This is because group-structure optimization and foil selection optimization are interdependent problems (a fine binning would be required to accurately capture parts of the neutron spectrum where $\left| \frac{d\phi(E)}{dE} \right|$ or $\left| \frac{d\sigma(E)}{dE} \right|$ are large, thus requiring knowledge of the neutron spectrum and the set of foils that will be used in the experiment, which are unknowns before the experiment planning begins); further effort placed into group-structure optimization would likely only yield marginal gains, and is not the focus of this thesis.

The neutron spectrum provided in [49] was used as the *a priori* spectrum ϕ_0 . This spectrum only spans the energy range 0.02 eV — 700 MeV. The neutron bins not covered by this energy range are considered to have zero flux and are removed from consideration, in the same vein as Section 4.4.3. This reduces the number of groups down to 537. This *a priori* is plotted in blue in Figure 6.5.

The second step is to extract the appropriate nuclear data from the nuclear data library. To cast the widest net in search of reactions that may be useful for high-energy neutron detection,

6.1 Experiment design




Figure 6.2: The general layout of the two target stations in ISIS. The synchrotron releases proton pulses at 50 Hz, but only every fifth pulse is sent towards target station 2, thus protons arrive in target station 2 at 10 Hz. ChipIr is the **brown** instrument at target station 2. This diagram is shared under the Creative Commons License by [151].

TENDL-2017 [138] was selected as the library used in step 2 as it has the largest collection of reaction cross sections among all nuclear data libraries. The decay-2020 library was used as well as it was the most up-to-date decay library available at the time.

All naturally occurring elements in the periodic table (except those that produced fertile, fissile or fissionable isotopes) were considered. Due to practical constraints, some elements had to be modified or rejected according to the following criteria:

- The element must be available as a foil, tape, or crystal. Otherwise, if it is only available in other forms (powder, liquid, etc.) then it is rejected.
- The element must not be flammable, explosive, or highly toxic.

If an element does not satisfy either of the above criteria, then it falls back to finding a compound consisting of at least 50% of that element (e.g. instead of iodine, caesium iodide was used in the list). If none of the compounds available in the laboratory supplies provider Goodfellow's catalogue satisfies both conditions either, then the element is finally rejected. To reduce the cost of purchase, not only pure elements were considered, but also compounds which are cheaper to manufacture. This yielded a total of 99 materials, giving $\Theta = 99$ according to the notation in Section 5.5. These 99 materials were used as the candidate foils¹ for the next step.

The relative efficiency curve of the gamma-ray detector was obtained via an MCNP simulation of the HPGe detector of a point-source of γ -rays situated at the lowest height that the foil sample may be placed at. This curve was later found to match the experimentally obtained efficiency curve but with significant deviations under 59.5 keV; regardless, this deviation should not impact the foil selection process significantly, as the deviation of the inputted efficiency curve from reality is one of the accepted inaccuracies accounted for in the later steps of the foil selector (Section 5.6.4).

The lower and upper limits of the gamma-ray detector were set to 20 keV and 2800 keV respectively, as these were considered the practical limits of the HPGe detector used in the experiment.

For the third step, the experiment schedule had to be decided. Since there is only one neutron source and one gamma-ray detector, it was decided that foils shall be sequentially irradiated and then measured on the gamma-ray detector. 20 foils were intended to be irradiated in the four days of experiment time available. Therefore it was decided that one hour of irradiation

¹Some of these materials, such as Germanium, were thin crystals rather than flexible foils; while others (copper) were thin film/tapes that could be cut into foils. But for the purpose of this thesis, they will all be called foils from here on.

6.1 Experiment design

time followed by one hour of acquisition for each foil (where the acquisition period of the next foil overlaps with the irradiation period of the previous foil) was a manageable timescale. It was determined that the time it takes from closing the shutter to shut out the neutrons, to opening the irradiation chamber doors to retrieve the foil sample, to placing the foil on the gamma-ray detector to begin gamma-ray spectrum acquisition, shall be no longer than five minutes. This gives the following irradiation schedule, applicable for all foils:

$$a = 1 \text{ hour} \tag{6.1}$$

$$b = 1 \text{ hour and 5 minutes} \tag{6.2}$$

$$c = 2 \text{ hour and 5 minutes} \tag{6.3}$$

The gamma-ray detector's dead time was assumed to be $\approx 50 \mu\text{s}$ per pulse, and aiming for a deadtime $\leq 5\%$, the maximum count rate is inputted as $5\%(\frac{1}{50 \mu\text{s}}) = 1,000 \text{ s}^{-1}$.

Between the third and fourth steps, the prototype foil selector's procedure deviated from what was described in Sections 5.6.3 and 5.6.4: instead of requiring the user to manually select the suitable dimensions for each candidate foil after the third step, it simply set all candidate foil's hypothetical volumes to equal to the maxima calculated by criterion 5.32, and set their thicknesses to any value smaller than Equation 5.29 (as the thickness has no material impact on reaction rate as long as it is under this limit), making the bold assumption that all foils can be purchased/manufactured to these arbitrary dimensions. The foil selection then proceeded to the fourth step, producing the two lists mentioned in Section 5.5.1, as well as their first match, which was a list of $\vartheta = 20$ foils according to the notation in Section 5.5.

The prototype foil selector program also had a more primitive noise threshold than was presented in Section 5.6.4 criterion 5.53 (or proposed in Section 5.7 criterion 5.60). Instead of measuring whether a photopeak was detectable or not, it attempted to determine whether a radionuclide X in the decay chains of the products produced by the root nuclides of all neutron-induced reactions in the foil was detectable. The sum of all its photopeak counts must exceed the user's inputted value of the minimum number of counts to be considered detectable.

6.1 Experiment design

Expressed mathematically,

$$\sum_k^{\text{All root nuclides } k \text{ that contains } X \text{ in its decay chain}} \sum_l^{\text{All gamma-ray lines of } X} \left(\epsilon_{\text{abs}}(E_l) I_{(X, E_l)} \int_b^c \lambda_X \frac{\overline{P_{k \dots X}}(t)}{P_{k,0}} dt \right) N_k \geq \text{User-inputted value} \quad (6.4)$$

$$\sum_k^{\text{All root nuclides } k \text{ that contains } X \text{ in its decay chain}} \sum_l^{\text{All gamma-ray lines of } X} \left(\epsilon_{\text{abs}}(E_l) I_{(X, E_l)} \int_b^c \lambda_X \frac{\overline{P_{k \dots X}}(t)}{P_{k,0}} dt \right) \left(\sum_i^n R_{ki} \phi_{0,i} \right) \geq \text{User-inputted value}$$

$$\sum_k^{\text{All root nuclides } k \text{ that contains } X \text{ in its decay chain}} \sum_l^{\text{All gamma-ray lines of } X} \left(\epsilon_{\text{abs}}(E_l) I_{(X, E_l)} \int_b^c \lambda_X \frac{\overline{P_{k \dots X}}(t)}{P_{k,0}} dt \right) \left(\sum_i^n R_{ki} \phi_{0,i} \right) \geq 500. \quad (6.5)$$

Where $\frac{\overline{P_{k \dots X}}(t)}{P_{k,0}}$ was given by Equation G.4. This value was inputted as 500, i.e. all nuclides with fewer than 500 counts in all its photopeaks expected would be disregarded from the analysis when building matrix $\underline{\underline{\mathbf{G}}}$. This threshold later proved to be too low, including too many radionuclides that were undetectable when performing the experiment analysis, as shown by Section 6.3.

Similarly, due to the incomplete development of the foil selector at the time, the weight vector $\mathbf{w}$ for the **precision** metric was chosen as $\mathbf{1}$ instead of the recommended value in Equation 5.44. For the CCFE-709 group structure where neutron bin sizes are logarithmically weighted (i.e. larger at higher energy), setting $\mathbf{w} = \mathbf{1}$ is identical to applying a logarithmic weight to the neutron spectrum and simultaneously assuming it is piecewise constant in lethargy space, which means it is equivalent to setting $\mathbf{w}$ to the values given by Equation 5.45. This implies that, even if the experiment was carried out perfectly and all gamma-ray spectra peaks identified by the foil selector were actually detected (i.e. even if noise due to Compton scattering was not an issue), then the response matrix would still only favour the precision of the unfolded spectrum at low energy range, going against the premise of focusing on the high energy part of the neutron spectrum as stated at the beginning of this chapter.

Out of the $\vartheta = 20$ foils recommended by the prototype foil selector program, 6 had to be removed due to budget limitations. Therefore only 14 foils were purchased. These 14 foils were combined with four more foils from the existing inventory of foils available to the UKAEA and ChipIr research teams to form a set of 18 foils shown in Table 6.1. Sample 18-CdSnCd was created by two thick cadmium foils sandwiching a piece of tin foil, to reduce the sensitivity of the tin foil to thermal and epithermal neutrons, in an attempt to modify the response matrix. However, due to the difficulty of the analysis process (Section 6.2), it was not possible to extract

6.1 Experiment design

Table 6.1: The dimensions of each foil, compared to the upper limits recommended by criteria 5.50 and 5.29, calculated by simulating irradiation for 1 hour by ϕ_0 , waiting five minutes, and then placing it directly on the casing of the HPGe detector. Some foil samples are reused to improve the precision of the result, hence the repeated elements across some samples.

Sample ID – Element	upper limit for this foil material		Measured value		Note
	mass (g)	thickness (mm)	mass (g)	thickness (mm)	
01-Cu	0.786	0.0807	8.233	1	
02-Cu	0.786	0.0807	0.818	0.1	
03-In	0.00264	0.00185	0.477	0.1	Indium foils could not be purchased in the dimensions close to the upper limits.
04-Sn	2.92	0.901	1.225	0.109	
05-Ge	4.10	0.421	0.277	0.5	
06-Zn	1.10	0.835	0.130	0.025	
07-Pt	1.08	0.0136	0.138	0.01	
08-Hf	1.02	0.00343	0.151	0.02	
10-Yb	1.43	0.0435	0.096	0.025	
11-Mo	3.64	0.0516	0.021	0.003	
12-Cd	2.99	0.000244	0.251	0.05	
13-Er	0.732	0.00480	0.129	0.025	
14-Nd	1.55	0.0182	0.338	0.025	
15-Bi	0.340	15.9	1.400	Not measured	Not selected by the prototype foil selector program/ retrieved from existing inventory of foils.
16-Au	0.129	0.000901	0.240	Not measured	Not selected by the prototype foil selector program/ retrieved from existing inventory of foils.
17-Er	0.732	0.00480	0.129	0.025	Same foil as 13-Er.
18-CdSnCd	Not calculated		Cd: 5.676 and 5.481; Sn: 0.349.	Cd: 1.0 and 1.0; Sn: 0.109.	Made of two cadmium foils sandwiching a tin foil.
19-W	0.846	0.00862	0.374	Not measured	Not selected by the prototype foil selector program/ retrieved from existing inventory of foils.




Figure 6.3: Photo taken inside the irradiation chamber. Neutrons exited perpendicularly to the wall mounted with the neutron detector (black box), and struck the foil held in the polyethylene bag, activating the foil.

its values, and thus it was discarded in the end. Sample 15-Bi, 16-Au, and 19-W were not foils selected by the foil selector, but were used anyway as they were available; this was done because any additional foil used would only strictly increase the **accuracy** of the unfolded result (as explained in Section 5.5.1).

In Table 6.1, the column of masses upper limits are calculated with criterion 5.50 using the efficiency curve at the closest gamma-ray spectroscopy measurement position; if the actual mass of the foil exceeds this limit, then criterion 5.50 can still be met by simply placing the foil further away from the HPGe detector, thus reducing the detection efficiency. The only exception to this is the indium foil: indium foils could not be manufactured to the size desired (as the small size would present difficulty handling), thus the purchased foil far exceeded the upper limit, and was in fact too radioactive to handle for several hours after irradiation (see Table 6.2) even after deliberately reducing the irradiation period a from one hour to half an hour.

The thickness values given in Table 6.1 only give the upper bound if *all* reactions need to be detected without altering the response matrix due to self-shielding; otherwise, offending reactions can simply be ignored in the analysis procedure, as explained in the final paragraph of Section 5.1.5.

For each foil during its irradiation phase, it was placed in bags and positioned in the path of the neutron beam in Figure 6.3. Its position is confirmed by a laser guide system. When the shutter was opened, neutrons exiting the target were collimated by the aperture shown in Figure 6.1b and passed through the neutron detector mounted on the wall in Figure 6.3 to activate the foil samples. The neutron detector records the accumulated counts of detected neutron-induced pulses every three seconds: this record was used to estimate the start and end of irradiation periods, as well as monitor for any anomalous drops in beam current.

Finally, when the planned irradiation period is finished at time $t = a$, the shutter is closed, and the foil is retrieved from the irradiation chamber to be placed on an HPGe gamma-ray detector from time $t = b - c$.

6.2 Data analysis

During the four days of irradiation, the proton beam supplied by the synchrotron experienced occasional drops followed by complete stoppage, which led to lower-than-usual or zero flux being irradiated onto the foil despite the shutter to the irradiation chamber remaining open. As changes in the current to current violated the time-invariant neutron spectrum assumption stated in Section 5.1.3, the data affected by those were repaired by modifying the top-hat function to include these down-time gaps where possible; and discarded where the repair was not possible (i.e. where the beam current was stuck at less-than-full-power mode giving a neutron spectrum that $\neq \phi_{\text{true}}$ and $\neq 0$ for significant periods, i.e. > 10 s).

Instead of following the schedule outlined by the values of (a, b, c) given by Equation 6.1 to 6.3 exactly, in practice, a was sometimes shortened due to this beam availability issue. c was also often lengthened to increase the number of measured counts to reduce the fractional error on $\mathcal{N}_l$ when some extra idling time of the HPGe detector was available due to schedule slippage. To boost the efficiency even further by maximising the number of counts measured by the HPGe, some foils (13-Er and 15-Bi) were also measured more than once, giving 13-Er-later and 15-Bi-later. Sample 09-Bag was a control, where a plastic bag was irradiated to see if it produced a significant amount of background radiation that could have interfered with the analysis (the result of which was negative).

The absolute detection efficiency curve for each foil was determined by modelling the gamma-ray detection setup in MCNP, accurately incorporating the dimensions of both the foil and the HPGe crystal. Direct validation of these simulated efficiency curves was not feasible, as there were no extended sources with known activities that matched the dimensions of the foils available

6.2 Data analysis

Table 6.2: Irradiation and measurement schedule of all foils. The neutron flux “stuttered” as the beam current dropped to zero occasionally; these occurrences are recorded in the last column. They are accounted for by modifying the top-hat function that $P_{k...X}$ is convolved with in Equation 5.21.

sample name	a	b	c	Note
01-Cu	0:28:37	2:24:47	20:09:47	Beam was turned off at 0:18:34 and 0:25:58 for 3 s each.
02-Cu-elevated	1:00:12	1:03:48	1:38:48	Terminated at 1:38:48 to move foil closer to HPGe detector.
02-Cu-contact	1:00:12	1:46:52	1:51:52	Same foil as 02-Cu-elevated; but at 1:38:48 the foil was placed closer to the HPGe detector to increase the counting efficiency, which means a different efficiency curve must be used, thus it is analysed as a separate spectrum.
03-In	0:34:50	8:20:59	23:05:58	
04-Sn	1:00:16	1:04:34	1:59:34	
05-Ge	0:38:35	0:51:25	1:51:25	
06-Zn	0:54:09	0:59:38	2:24:38	Beam was turned off at 0:48:17 for 84 s, and again at 0:53:00 for 48 s.
07-Pt	1:00:06	1:04:41	2:16:41	
08-Hf	1:02:24	1:07:01	2:02:01	
09-Bag	15:29:47	15:40:15	16:55:15	Beam was turned off at 4:06:22 for 17 min and 51 s, at 5:13:56 for 15 s, at 5:20:24 for 12 s.
10-Yb	0:59:57	1:05:42	2:18:42	Beam was turned off at 0:38:47 for 66 s.
11-Mo	1:03:27	1:08:40	2:27:40	
12-Cd	1:00:06	1:03:51	1:40:51	
13-Er	0:24:40	0:34:31	2:14:31	
13-Er-later	0:24:40	23 days, 20:21:04	23 days, 21:09:23	Same foil sample as 13-Er.
14-Nd	1:19:52	1:23:53	2:26:53	Beam was turned off at 1:11:18 for 30 s.
15-Bi	1:00:03	1:03:09	2:03:09	
15-Bi-later	1:00:03	23:09:57	7 days, 15:49:55	Same foil sample as 15-Bi.
16-Au	0:59:57	1:03:40	2:20:40	Beam was turned off at 0:01:00 for 30 s, at 0:33:20 for 21 s.
17-Er	1:10:04	1:17:53	2:57:53	Beam was turned off at 0:05:42 for 7 min and 6 s, at 0:13:21 for 12 s, at 0:13:39 for 45 s, at 0:14:36 for 15 s, at 0:14:54 for 9 s, at 0:15:06 for 75 s, at 0:16:25 for 6 s, at 0:20:16 for 81 s, at 0:38:05 for 75 s, at 0:42:08 for 4 min. It was also briefly irradiated in the same facility for 24 min and 40 s at $t = -24$ days. (See sample 13-Er).
18-CdSnCd	0:43:15	0:47:54	1:27:54	
19-W	0:31:05	0:35:09	17:15:07	It was also briefly irradiated in the same facility at $t = -4$ days.




Figure 6.4: Validation of the MCNP model by comparing its calculated efficiency curve against measured efficiencies at a constant source-detector separation distance.

in the laboratory. However, the validity of the MCNP model was verified by comparing the simulated efficiency curve, which modelled the source as a point-source, to measured efficiencies. The simulated efficiency curve closely matched the efficiency datapoints obtained from measuring point-like reference sources with known activities, as shown in Figure 6.4. This close match confirmed the accuracy of the MCNP model of the HPGe detector, thereby providing confidence in the absolute detection efficiency curves generated for each foil.

6.2.1 Extracting the number of radionuclides decayed using Genie 2000

The gamma-ray detector recorded the incoming gamma-rays in a pseudo-time-resolved manner, by saving and restarting a new gamma-spectrum acquisition at 5-minute intervals. In theory, this would allow us to compare the reaction rates using Equation H.7. However, the > 3000 spectra generated as a result were too numerous to be analysed by hand, and would have required bespoke software to be written to procedurally analyse all of them, adding to the already long list of technical challenges faced in this experiment. Therefore, to reduce the complexity of the data to a manageable level, spectra that are acquired back-to-back are simply added together to reconstruct the gamma-spectrum as if it were one continuous acquisition. This yielded the 22 spectra in Table 6.2.

To extract the number of decays of each species of radionuclide in each foil, the following “Analysis Sequence” was performed in gamma-ray spectroscopy software Genie 2000 [153]:

1. Peak Locate – Library (Gamma M) — the centroid energy and width of peaks were identified at this step, using radionuclide decay information provided in the form of an .NLB file.
2. Peak Area – Library (Gamma M) — A fit was performed on each peak to find the area for each peak (as well as their associated uncertainties) using the result of the previous step.
3. Efficiency Correction – Standard — Corrected for the absolute efficiency of the gamma-ray detector $\epsilon_{\text{abs}}(E_l)$. Using the “Interpolated” option here, the efficiency curve produced by MCNP was supplied.
4. Nuclear Identification – NID plus Interference Correction — The same .NLB files were supplied. To perform cascade correction in this step, Genie was provided with an ISOCS [154] model of the gamma-ray detection set-up.

The .NLB files discussed here were created with a program called the Nuclide Library Editor [155]. For each gamma-ray spectrum, the names of all decaying isotopes and their gamma lines (including gamma-ray energy in keV, intensity in %, and the uncertainties of both) were entered into the Nuclide Library Editor, which then saved the data as .NLB files, which is a type of binary file. Since the Nuclide Library Editor did not have a command-line interface (CLI), the data had to be entered manually through its graphical user interface. Given the large number of gamma-ray lines that needed to be entered for all the foils, manual entry was impractical. To automate this process, AutoHotkey [156] was used to simulate keyboard and mouse inputs. The AutoHotkey script automated the entry of information copied from a list of nuclides and their associated gamma-ray lines, which were expected to be detectable by the foil selector program (i.e., those expected to produce > 500 counts at their photopeaks according to criterion 6.5).

During the analysis process, it was noted that the Analysis Sequence often fails when the .NLB files were too big, i.e. contained too many gamma lines and nuclides, giving enigmatic errors². To overcome this, lines of the least number of expected counts are iteratively removed from the list of nuclides and gamma-ray lines, until the resulting .NLB file no longer causes an error during the Analysis Sequence. Additionally, lines at $E_l = 511 \text{ keV} \pm 1 \text{ keV}$ were also removed as this is known to be a particularly “crowded” line, i.e. many radioisotopes emit a high intensity of γ -rays near this energy, thus knowing the number of counts accumulated in the 511 keV peak across time $t = b - c$ would not be sufficient to uniquely determine the

²The details of these two errors are presented in <https://youtu.be/DEC7ahRSTDY?t=139>, but these details are not presented in this thesis directly as it does not add any material value to the discussion.

number of decays of any nuclides (without employing the time-resolved technique mentioned in Equation H.7, which had been concluded to be infeasible at the beginning paragraph of Section 6.2.1).

In step 4 of the Analysis Sequence, the ISOCS model of the HPGe detector had to be selected from a list of preset dimensions, which did not include an exact match of the dimensions of the HPGe crystal of the detector used. Therefore the 70 mm $\times$ 50.3 mm crystal used in the actual detector was simulated by a 70 mm $\times$ 70 mm crystal in the simulation instead. This slight inaccuracy in dimensions was deemed acceptable as this model was only used for coincidence correction.

The resulting activities of all radioisotopes created in each foil were extracted, forming a vector $\boldsymbol{\eta}$, which became the vector to be unfolded by the unfolding algorithm. Again, due to the discrepancy in the functionality of the prototype foil selector and foil selector framework described by Chapter 5, this new “response vector” $\boldsymbol{\eta}$ is a different response vector than $\mathbf{N}$ (original definition of the response vector) or $\mathbf{N}$ (effective response vector), as η_X represents the number of nuclide X that decayed between time $t = b - c$.

$$\eta_X = \sum_k^m \sum_{\substack{\text{all radionuclides} \\ \text{which are} \\ \text{descendants of} \\ \text{root nuclide } k}}^X \left(\lambda_X \int_b^c \frac{\overline{P_{k \dots X}}(t)}{P_{k,0}} dt \right) N_k. \quad (6.6)$$

The expression $\frac{\overline{P_{k \dots X}}(t)}{P_{k,0}}$ is defined by Equation G.4.

Referring to Equation 3.6, one can express the gross number of radionuclides produced N_k in terms of neutron flux ϕ_i :

$$\eta_X = \sum_i^n \left(\sum_k^m \sum_{\substack{\text{all radionuclides} \\ \text{which are} \\ \text{descendants of} \\ \text{root nuclide } k}}^X \left(\lambda_X \int_b^c \frac{\overline{P_{k \dots X}}(t)}{P_{k,0}} dt \right) R_{ki} \right) \phi_i \quad (6.7)$$

$$= \sum_i^n \rho_{Xi} \phi_i \quad (6.8)$$

The unfolding algorithm in the following section (Section 6.2.2) uses $\underline{\boldsymbol{\rho}}$ in place of the response matrix, $\boldsymbol{\eta}$ in place of the response vector (and its uncertainty $\underline{\text{cov}}(\boldsymbol{\eta})$ in place of the uncertainty on the response vector), to perform the unfolding.

To conclude this subsection, the process of extracting $\boldsymbol{\eta}$ from each gamma-ray spectrum is re-iterated chronologically below:

1. All of the relevant inputs for that foil are inputted into the prototype foil selector program (number of atoms of each species, the nuclear data for their microscopic cross sections and decay, and the schedule, etc.). The prototype foil selector program then creates the list of radionuclides and their
2. MCNP and ISOCS models of the HPGe detector were created to get the efficiency correction and cascade correction for Genie.
3. An .NLB file is created from this list;
4. This .NLB file is inputted into Genie for the appropriate Analysis Sequence. If the Analysis Sequence fails, we fall back to the previous step and remove some lines from the list before re-creating the .NLB file.

This process is time-consuming, requiring several hours per spectrum, the majority of which is spent on steps 3 and 4, which were already optimised by using AutoHotKey script instead of manual entry. This explains why, at the beginning of this subsection, it was stated that performing a time-resolved analysis was infeasible, as it would have required > 6000 hours of labour.

6.2.2 Results

Following the procedure explained in the previous section, the number of decays of each species of radionuclide is obtained. In addition to the usual uncertainty value associated with the number of decays, Genie also outputs a confidence value between 0 to 1 to indicate how confident it is that the isotope has been correctly identified. Radionuclides whose confidence value < 0.95 are discarded, as they may be incorrectly identified radionuclides due to interfering gamma-ray peaks. The number of species of radionuclides detected from each foil is given in Table 6.3. The response vectors $\boldsymbol{\eta}$ for each foil is plotted in Appendix J. The gamma-ray spectrum 02-Cu-contact and 13-Er-later were discarded as they had too few detected counts to be useful, and thus were not worth the time spent performing the Analysis Sequence described in Section 6.2.1. Even though there are 393 foil-and-radionuclides pairings shown in Table 6.3 (and therefore 393 rows in the matrix $\underline{\boldsymbol{\rho}}$ inputted as the response matrix to the unfolding algorithm), the rank of the matrix still only reached 124. This is because many nuclides originated from the same few root radionuclides in each foil, which was the bottleneck of information, as the maximum rank that $\underline{\boldsymbol{\rho}}$ can take is the number of root radionuclides m .

Table 6.3: Number of radionuclides that were detected by Genie using the .NLB files created in Section 6.2.1.

Foil sample	Number of types of radionuclides identified in the sample by Genie
01-Cu	20
02-Cu-elevated	8
03-In	25
04-Sn	40
05-Ge	22
06-Zn	14
07-Pt	22
08-Hf	18
10-Yb	28
11-Mo	22
12-Cd	32
13-Er	22
14-Nd	40
15-Bi	25
15-Bi-later	15
16-Au	18
17-Er	19
19-W	3

Using the unfoldingsuite, the target χ^2 was manually lowered for each algorithm until the program failed to converge due to negative fluxes, which is a sign that the minimum χ^2 is reached. This yielded the unfolded spectra ϕ_{sol} 's in Figure 6.5. The purely iterative algorithms GRAVEL and SAND-II never decreased the χ^2 at all, while the resulting spectrum unfolded by the Pseudo-Inverse algorithm showed negative fluxes and was clearly unphysical, thus adding no value to the discussion. Therefore the ϕ_{sol} of all of these algorithms were not shown. As expected, since the response matrix has more rows than there are columns, the minimum χ^2 achievable is > 0 . These values of minimum $\frac{\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})}{\text{degree of freedom}=393}$ achieved in Figure 6.5 are $\gg 1$. They can be explained by improper measurements of the $\mathbf{N}_{\text{meas}}$ (due to the challenging number-of-decay extraction process shown in Analysis Sequence of Section 6.2.1), and/or incorrectly calculated $\underline{\underline{\rho}}$ (due to the unstable beam, and therefore possibly an unstable ϕ_{true}).

It was obvious that none of the unfolding attempts were successful: IMAXED's and AMAXED's solutions were simply rescaled copies of the *a priori* (save for a single unphysical artefact in AMAXED's solution spectrum around 800 eV); while AMAXED-Regularization did manage to lower the χ^2 even further, its ϕ_{sol} showed clearly unphysical features of oscillations in $\phi(E)$ w.r.t E . These oscillations, which are imprints of resonance peaks in the response matrix, were observed to grow in amplitude when the target χ^2 was lowered.




Figure 6.5: The ChipIr spectrum unfolded using various algorithms of the unfoldingsuite.

In a drastic measure to reduce the complexity of the problem in the hopes of converging to a more realistic neutron spectrum, the entire analysis was repeated with a simplified group structure of only 49 bins, created by

1. merging all bins covering the energy range $< \frac{1}{40}$ eV into a single bin³; then
2. all bins covering the energy range $\geq \frac{1}{40}$ eV were merged in groups of 10 bins.

This yields a 49-group structure that was used to perform the unfolding in Figure 6.6, where the $\underline{\rho}$ has rank = 46, making it a weakly-underdetermined problem. Again, as the target χ^2 was lowered for each unfolding algorithm, some of them reached a lower $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})$ than others; however, none of them appears to be physically realistic, showing oscillations in $\phi(E)$ w.r.t E .

6.3 Review of the experiment: Lessons learned

Let us begin by discussing the successes: The HPGe detector was the correct choice of gamma-ray detectors due to its high resolution, which enabled photopeaks to be differentiated from one another. This capability allowed for the measurement of more root nuclides. Switching to a NaI detector for efficiency gains would not have been advantageous, as the improvement in detection efficiency (that can similarly be achieved on the HPGe by simply using a wider foil and placing it closer to the HPGe crystal) would not have justified the significant loss in resolution, as many

³ $\frac{1}{40}$ eV is approximately the energy of neutrons at room temperature.




Figure 6.6: The ChipIr spectrum unfolded in the same manner as in Figure 6.5 but with a simplified group structure.

peaks would merge and become indistinguishable as a result, some peaks' separations were already near the limits of the HPGe detector's resolution.

In contrast, a few other suboptimal choices were made that could have contributed to the unsatisfactory result in Section 6.2.2, one of which is the choice of nuclear data library. The microscopic-cross section values for incident neutron energies ≥ 30 MeV are stored in a different location (MF=10, MT=5) than the usual (MF=3, MT=*). While reading in the data, the foil selector program simply concatenates these two records together to form a continuous cross section function that spans from 0 MeV to > 30 MeV; and if data is available for the production of a root radionuclide k in MF=3 but not MF=10, then the cross section $\sigma(E)$ is simply set to 0 at $E > 30$ MeV; vice versa, if data is available for the production of k in MF=10 but not MF=3, then the cross section $\sigma(E)$ was simply set to 0 for $E < 30$ MeV. This was where discontinuity in the microscopic cross section was often observed, calling into question the validity of TENDL's data, which would have led to inaccurate values on the response matrix and hence the incorrectly unfolded neutron spectrum. In the absence of any ability to read the uncertainty of the nuclear data (MF=33 files, as explained in Section 5.7), it is difficult to test this hypothesis. Because of the unreliability of nuclear data at this energy range, measuring high-energy neutrons is fundamentally different to measuring fusion neutrons, and its poor performance in this one high-energy neutron spectrum measurement experiment is not indicative of the foil selector program's capability to suggest foils to be used for designing fusion neutron spectrum unfolding experiments.

Another plausible explanation for the poor quality of the unfolded neutron spectrum is the poor choice of thresholds used for selecting foils when deciding whether or not nuclides are detectable. Criterion 6.5 was used, but it is a criterion that does not consider how much background noise is there near E_l , and therefore whether the neutron spectrum is detectable or not. The prototype foil selector program also failed to ignore $E_l = 511$ keV gamma-ray lines when checking against criterion 6.5. A lot of nuclides were falsely assumed to be detectable by the prototype foil selector program as a result, even though the majority of their qualifying 500 counts were hidden in the 511 keV peak where they could not be disentangled from the 511 keV gamma-rays from other nuclides, or behind the background noise formed by Compton continua in low energies. This information could, in theory, be recovered if the spectra were analysed in a time-resolved manner, as the Compton continua caused by high rates at the beginning of the acquisition (close to $t = b$) would not drown out peaks emerging later in the acquisition period (close to $t = c$); and the method presented in Appendix H would have allowed the contributions from different radioisotopes to the 511 keV peak to be separated from one another using the differences in their half-lives.

Finally, the poor quality of the unfolded spectrum could have also been the result of a poor choice of group structure and weight vector $\mathbf{w}$ used when optimising the foil selection. The weight vector $\mathbf{w}$ was incorrectly chosen as explained in Section 6.1, thus the chosen response matrix was optimised to be extra sensitive to deviations of the neutron spectrum at low energy ranges, where the neutron bin widths were small. If instead the correct weight vector was used (Equation 5.44), or if a different (linearly spaced out) group structure was chosen, then the chosen set of foils (and hence its response matrix) should be more sensitive to the higher energy part of the neutron spectrum and would have been less affected by noise in these regions.

Alternatively, if this experiment analysis were to be repeated using the same set of data, we could still improve the quality of the unfolding by repeating the analysis by removing all reactions sensitive to < 10 MeV neutrons from the response matrix, leaving behind only reactions that have negligible cross sections below 10 MeV, and then amending the group structure to exclude bins below 10 MeV. This strategy forces the unfolding algorithm to focus on high-energy bins ($E > 10$ MeV), improving the degree of determination and therefore the quality of the unfolded solution by narrowing the search space. When removing bins below 10 MeV, it is crucial to ensure that the remaining rows of the response matrix are insensitive to neutron fluxes in this range, i.e.

$$\sum_i R_{ki} \phi_i \approx 0 \quad \forall \quad E_i < 10 \text{ MeV} \quad (6.9)$$

Otherwise, the gross production of the k^{th} nuclide N_k would be underestimated due to the unaccounted-for contributions from the excluded energy range.

Finally, if the experiment were to be repeated with the explicit goal of measuring > 10 MeV neutrons, one can modify the response matrix by applying shielding materials in front of the foil: Instead of searching for naturally occurring isotopes whose cross section for radionuclide productions satisfies the condition in Equation 6.9, we can artificially engineer a set-up that ensures this: By placing shielding materials that are transparent to high-energy neutrons but opaque to low-energy neutrons in front of the activation foil, the response matrix can be modified to be sensitive to only high-energy neutrons. However, this type of setup risks further altering the neutron spectrum and violating assumptions from Section 5.1.4, thus must be carried out with great care.

In either case, the foil selector program and the experiment analysis procedure would have to be modified quite thoroughly to achieve such a set-up, which is beyond the scope of this thesis: as neutron spectrum unfolding foil selection is still in its infancy, it is far too ambitious to attempt such modifications to the response matrix within the course of this PhD, as it would have required more work than could be done in the time available.

6.3.1 How the foil selector's organization impacted the experiment analysis

One of the main approaches identified as a potential time-saver was the use of existing software. Specifically, Genie 2000 was used to perform the gamma-ray spectrum analysis to identify the number of radionuclides that decayed during the measurement duration. However, the analysis procedure could have been restructured to reduce the work even further.

For this experiment, the following steps were used:

1. Using Genie's Analysis Sequence to convert the gamma-ray spectrum to the number of radionuclides decayed η ;
2. Calculate $\underline{\underline{\rho}}$ by re-using the fourth step (Section 5.6.4) of the prototype foil selector program.
 - This requires a different input of (a, b, c) and efficiency curve ϵ_{abs} for each foil. Thus this step was repeated 18 times, once for each of the 18 spectra in Table 6.3.
3. η is unfolded with $\underline{\underline{\rho}}$ using an unfolding algorithm to give ϕ_{sol} .

This required intensive use of the prototype foil selector program, which was still under active development during the analysis procedure. It was not designed to produce 18 different response matrices quickly one after the other, thus the haphazardly put-together prototype foil selector program grew to an unmanageable size and became difficult to debug.

Instead, the following scheme would have been preferable:

1. Use some kind of gamma-ray spectrum unfolding software (either Genie or any other simple peak-identifying software) to identify peaks in the gamma-ray spectra, thus extracting $\mathcal{N}$;
2. Create $\underline{\underline{\mathcal{R}}}$ using the FISPACT-II using this procedure: Consider the i^{th} column of the effective response matrix, where $1 \leq i \leq n$, to obtain the values of $\mathcal{R}_{li}$:
 - (a) Begin by creating a ϕ_{test} such that $\phi_{\text{test},j} = \delta_{ji}$, i.e. a spectrum of zero flux in every bin except for the i^{th} bin, which has $1 \text{ cm}^{-2}\text{s}^{-1}$ flux.
 - (b) Use FISPACT-II to simulate the irradiation of each foil by ϕ_{test} .
 - (c) Then, using the **SPECTRUM** keyword, $\frac{1}{\epsilon_{\text{abs}}(E_l)} \frac{d\mathcal{N}_l}{dt}$ for $\mathcal{N} = \underline{\underline{\mathcal{R}}}\phi_{\text{test}}$ is obtained. This can be integrated w.r.t. time to obtain $\frac{1}{\epsilon_{\text{abs}}(E_l)}\mathcal{N}_l$.
 - (d) Multiply this quantity by the efficiency curve $\epsilon_{\text{abs}}(E_l)$ for that foil.
 - (e) This vector is now equivalent to the i^{th} column of the effective response matrix, $(\mathcal{R}^T)_i$.
3. Unfolding $\mathcal{N}$ with $\underline{\underline{\mathcal{R}}}$ using an unfolding algorithm to give ϕ_{sol} .

The main advantage of this approach is that since FISPACT-II is a well-maintained program, it sidesteps the difficulty surrounding nuclear data reading and decay calculations (detailed in Appendix H). It also offers a lot more functionality than was implemented into the prototype foil selector program, thus it would have been possible to determine the source of error when the expected response vector and measured response vector disagree, and perform decay pathway analysis to determine whether the measured response vector extracted was sensible or not.

The latter scheme has the additional advantage that it is easier to capture the covariances between different counts of gamma peaks (i.e. $\underline{\underline{\text{cov}}}(\mathcal{N})$), thus the result's uncertainty $\underline{\underline{\text{cov}}}(\phi_{\text{sol}})$ would be reduced compared to the first method using Genie, where the covariance between the number of decays of different radioisotopes $\text{cov}(\eta_x, \eta_{x'}) = 0 \ \forall \ x \neq x'$ are simply ignored, such that the unfolded solution ϕ_{sol} would likely have had a lower precision as measured by Equation 5.42. It has the serendipitous benefit of reducing the complexity of the foil selector program, as $\mathcal{N}$ will not have to be calculated explicitly, skipping the calculation of (3) in the diagram of Figure 5.4.

6.3.2 Reducing the difficulty of experiment planning

It is no secret that scientific experiments are difficult as experimentalists must be prepared to react to the myriad of ways in which things can go wrong, and foil activation experiments are no exceptions. During the execution of a foil activation experiment, there are two types of variables that may be changed by the experimentalist: the schedules for irradiating and measuring each foil, and the gamma-ray spectroscopy settings. Deviations from the experiment plans (e.g. when the sizes of foils being used do not lie within the limits set out by the foil selector, as was the case in Table 6.1) would lead to degradation in the quality of the results, and the experimentalist may opt to remedy these mistakes by changing the irradiation and measurement schedules and/or gamma-ray spectroscopy set-up; External factors such as availability of the neutron beam, personnel, samples, or equipment would also inevitably affect the schedules. Changes in the schedule lead to changes in the activities of the foil, which in turn requires changes to the gamma-ray spectroscopy settings.

An experimentalist who is proficient in carrying out foil activation experiments should have a good understanding of how they should react to these changes, and this understanding can be enhanced by computer program(s). Section 5.6.5 of the foil selector program attempted to do this by creating simulated gamma-ray spectra that the experimentalist can then use to determine whether the pre-amplifier and amplifier settings and foil placement position can be changed to maximise the number of counts in the photopeaks. This program may be extended to plot the expected changes to the activity with respect to changes in the irradiation schedules, and therefore what changes to the foil-detector-separation distance would be required to keep the gamma-ray spectroscopy set-up optimal. This program should also be able to plot the gains in **accuracy** and **precision** with respect to the increase in the measurement duration of each foil, allowing the experimentalist to weigh their options and react accordingly. If the experimentalist has achieved mastery of this type of experiment, they would be able to proactively anticipate these changes and make the necessary adjustments to hedge against their negative impacts including planning the experiment with the appropriate amount of float to account for schedule slippage. From this, we can conclude that there is room for improvement in the arsenal of tools available for planning and executing activation foil neutron spectrum unfolding experiments: namely, software may be developed to help experimentalists understand the impact of changes to their experiments, and react accordingly.

Alternatively, if an automated sample irradiation and retrieval system is installed, such as the pneumatic sample retrieval system in FRASCATI [58], where reliable irradiations and

6.3 Review of the experiment: Lessons learned

measurements can be consistently administered, then the number of failure modes of the experiment will be drastically reduced. This would negate the need for the aforementioned hypothetical programs to aid the experimentalist in conducting the experiment, as there is no need to deviate from the experiment plan and react to changes in the schedule. This method of carrying out foil activation experiments is likely the most user-friendly one, as the success of the experiment is less tied to the experimentalist's experience in gamma spectrometry, neutron activation and radioactive materials handling; and the analysis procedure would also be easier, as the schedule (values of a , b , and c) can be controlled to be exactly the same as the inputs to the foil selector program in Section 5.6.3 for every foil, thus the effective response matrix outputted by the foil selector program $\underline{\mathcal{R}}$ can be reused directly without modification. However, such an automated system is neither easy nor cheap to install and therefore has the drawback that it is not flexible enough to be deployed everywhere.

Chapter 7

Conclusion

Due to its unique ability to withstand the high temperatures, intense neutron flux, and strong magnetic fields at the first wall of fusion reactors, activation foil neutron spectrum unfolding remains the prime choice of instrument to be deployed for neutron spectrum measurement. This thesis delved into this technique, refining and upgrading it for application in fusion reactors, such that the neutron spectrum information that it extracts may help inform designs of future fusion power plants and improve the operation of fusion reactors (Figure 1.2). This work identified two key areas for improvement in activation foil neutron spectrum unfolding: the mathematical algorithms through which activation data are unfolded into the required spectra, and the experimental design process used to obtain this activation data.

After establishing the language necessary to discuss neutron spectrum unfolding in Chapter 3, a critical review of the existing unfolding algorithm was made in Chapter 4, where unfolding algorithms were classified into one of four categories: machine learning, parametric, regularising, and purely iterative algorithms (Figure 4.1). The difficulty of benchmarking algorithms for unfolding fusion neutron spectra is then explained using Figure 4.3. In the context of neutron spectrum unfolding where prior information is still sparse due to the limited understanding of the possible neutron spectra of fusion reactors, machine learning and parametric unfolding algorithms were not considered good choices of unfolding algorithms; and purely iterative algorithms lack the mathematical rigour needed to handle a complex and highly underdetermined system such as a activation foil neutron spectra unfolding. Therefore, regularising unfolding algorithms were identified as the most suitable category of algorithm to use for fusion neutron spectrum unfolding; within which, cross-entropy was identified as the most suitable regularising function. Several algorithms were rigorously derived based on this regularising function: IMAXED, AMAXED, and AMAXED-Regularization, forming the first major original contribution to knowledge of

this PhD. They had many advantages over the existing algorithms, including quick convergence, guaranteed convergence to the optimal solution, protection against unphysical features, and allowing analytical uncertainty propagation. These algorithms were shown to perform well in synthetic data. AMAXED-Regularization was shown to be most successful in particular, able to unfold the neutron spectrum of a p+Be neutron generator from the measured data using a stilbene scintillator detector. These algorithms, along with MAXED, SAND-II, GRAVEL, and a special class of purely iterative algorithms known as the Pseudo-Inverse unfolding algorithm, were implemented as a Python program known as the *unfoldingsuite* [129]. This *unfoldingsuite* leverages the advancement in computing capabilities in the past several decades, successfully modernising these algorithms in the process.

This investigation into the mathematics also revealed a clearer understanding of the difference between over-, fully- or underdetermined problems, which allowed the resolution of activation foil neutron spectrum unfolding to be quantified definitively: It was found that in a fully-determined set-up, it can achieve a maximum resolution of $\frac{E_{\max}-E_{\min}}{m}$ in the absence of noise, which for fusion neutron spectra, $\approx \frac{17 \text{ MeV}}{m}$, where m = rank of the response matrix, which is $\leq$ number of root radionuclides produced by each foil in the activation foil set. This sheds new insights into a previously under-explored area in existing research, as the difference that degree of determination makes had rarely been openly discussed, and researchers had generally been hesitant to quantify the resolution numerically.

As for the experiment design process, more practical factors had to be considered as we exit the world of pure mathematics and enter the world of experimental physics. In Chapter 5, the experiment design parameters of activation foil neutron spectrum unfolding experiments were identified, along with their constraints. Among all of these design parameters that control the quality of the unfolding experiment result, the parameter of foil selection was isolated as the target to be optimised. A foil selector framework was created to do so, where the quality of the unfolded result is quantified by two parameters: **accuracy** and **precision**. A multi-objective optimization algorithm was designed to choose a foil set that simultaneously optimises for both. This framework and multi-objective optimization algorithm is novel work that had not been seen before and therefore is another original contribution to knowledge from this PhD degree. This framework was implemented as a Python program known as the foil selector program.

The foil selector framework development was a time-consuming Herculean undertaking; as such the foil selector program was still under active development when it was deployed for the planning of the activation foil neutron spectrum unfolding experiment at the ChipIr neutron beam in Chapter 6. A close-to-optimal set of foils was selected using a prototype version of

the foil selector program, and they were activated by the neutron beam. Due to the unstable beam, incorrect sizes of foils being used, and the low confidence in the gamma-ray spectrum analysis result, the experiment analysis was more complicated than it otherwise would have been: The response matrix calculation had to be corrected for the complex irradiation schedule and decays during the measurement, and the response vector consisted of radioisotope decay numbers of low confidence levels due to the unexplained errors. As such, the unfolded spectrum was unsatisfactory and did not yield a meaningful decrease in χ^2 when unfolded.

The investigation into the mathematics of unfolding and the subsequently produced unfolding algorithms forms a self-contained part of this thesis — both empirical and theoretical arguments have been presented in a rigorous and self-consistent manner for the use of AMAXED-Regularization, AMAXED, and IMAXED over existing algorithms, proving their stability and ease of convergence. They have also been released as an open-source program for public review and critique, thus no significant further development is required for these algorithms. Under the assumption that our knowledge of fusion neutron spectra remains limited (and that only foils with reliable cross section values are used for activation experiments), then these regularising algorithms will remain the most suitable unfolding algorithms for fusion neutron spectrum unfolding using activation foils for the foreseeable future. This also highlights the need for strong collaborations between experts in plasma physics and neutronics in fusion: while the neutron spectra of specific configurations of fusion plasmas (e.g. conventional aspect ratio tokamak's plasma) are well understood, the neutron spectra in other configurations of plasma (e.g. large spherical tokamaks) are poorly understood; therefore at the moment, it is not possible to predict the neutron spectra at the first wall of future fusion power plants by extrapolation of existing plasma physics models. However, when a plasma model with more predictive power about the spectrum of plasma neutrons is eventually developed, plasma physicists and neutronics scientists would need to collaborate to develop a model of the neutron spectrum at the first walls of fusion devices, which would enable machine learning and parametric unfolding algorithms to be used in activation foil neutron spectrum unfolding experiments. These algorithms may be able to unfold to accuracies and precisions that surpass the regularising unfolding algorithms presented in this thesis and will be able to perform prior-less unfolding, thus would be another leap forward in the field of fusion neutron spectrum unfolding.

In contrast, much room for improvement in the foil selector framework is identified in Section 5.7. Apart from the changes to the user interface to improve the user experience as well as the ease of maintenance, two changes are proposed: One is that the noise threshold shall be refined such that the remaining gamma-ray lines unrejected by the noise threshold more

closely match the detectable list of gamma-ray photopeaks in reality. The other change is to the dependency of the foil selector program on OpenMC: another, smaller, open-source package (endf-parserpy) may be used, which would expand the range of readable files, and potentially allow the foil selector program to also take into consideration the uncertainty associated with the nuclear data of the foil chosen when performing the foil selection.

Section 5.7 also pointed to two open questions in Appendix H.1 and I.1, where in the former, a question about the optimal integration limits t_k ($1 \leq k < m$) for the γ -ray count rate measurements was posed; while in the latter, a question about the effect of simplifying a constant used in the non-self-shielding assumption (Equation 5.29) is posed. Both are excellent starting points for smaller, self-contained research projects.

Section 6.3 concluded the lessons learned from performing the experiment at the ChipIr neutron beam. Many isotopes (and therefore many photopeaks) are expected to be detected on the same spectrum in activation foil neutron spectrum unfolding experiments; hence HPGe was the correct choice of gamma-ray detector due to its high resolution that allows peaks to be differentiated from one another.

The conclusion was also drawn that there was a subtle difference in the challenge of foil selection for activation foil experiments that characterise high-energy neutron sources compared to activation foil experiments that characterise fusion neutron sources. The nuclear data applicable in each case is different since the former's nuclear data has larger uncertainties. Selecting foils for high-energy neutron sources may also require a different group structure, weighting ($\mathbf{w}$), or even different techniques of modifying the response matrix, such as by adding shielding materials in front of foils. This means that the poor performance in Section 6.2.2 is not representative of the performance of the foil selector program in fusion applications.

The experiment analysis procedure also could have been restructured: FISPACT-II could have been used to generate the effective response matrix $\underline{\mathbf{R}}$ to calculate the number of counts in each photopeak per unit flux in each bin of the neutron spectrum. This sidesteps the need to perform an Analysis Sequence in Genie to calculate the number of decays of each isotope. This approach would have the added benefit of being able to compare the effective response matrix $\underline{\mathbf{R}}$ obtained by the foil selector program against the effective response matrix $\underline{\mathbf{R}}$ obtained by FISPACT-II to validate the former's correctness.

Finally, additional modules to the foil selector program are suggested in Section 6.3.2 to aid the activation foil neutron spectrum unfolding experiment planning further, by providing tools that may help the experimentalist to make decisions about changes to the preparation of foils, irradiation and measurement schedule, and gamma-ray spectroscopy set-up.

This concludes a hopeful future for the field of activation foil fusion neutron spectrum unfolding: this thesis has laid down the foundational framework to understand the mathematics of unfolding and foil activation experiment design, on which further innovations can grow. This thesis developed a set of new unfolding algorithms and presented the arguments for why these algorithms are the best. Further testing and application of these algorithms on fusion neutron spectra data would help prove their value empirically. As for the foil activation experiment design, the foil selector framework created was the first of its kind. It tackled many difficult challenges, such as quantifying “**accuracy**” and “**precision**” of the unfolded result, as well as formalising the definition of the effective response matrix $\underline{\mathcal{R}}$. On the flip side, precisely because of its pioneering nature, this development was time-consuming and fraught with false starts and dead ends (most of which were ultimately not presented in the thesis). Therefore, regrettably, there was not an opportunity to test the foil selector program on a fusion neutron system. Nonetheless, this sets the foil selector framework up excellently for further research: as it is now a fully developed framework, implemented as a program, further research can be built on these to select foils, allowing it to be applied to more experiments, and demonstrate its value in fusion neutron spectra measurement experiments.

Appendix A

The minimum value of χ^2

This chapter concerns itself with the epistemology of science (in particular, how χ^2 of the solution should be interpreted). Therefore it is beyond the scope of the main text, but still deserves to be addressed, thus it is placed in the appendix.

A.1 Clearing up misconceptions around the minimum χ^2 in an overdetermined v.s. underdetermined problem.

This section started out as a spin-off document addressing the common misconception and mistakes when choosing the χ^2 , where users mistakenly assume that χ^2 should equal to certain values. The reason for such mistakes is because physicists are often trained to work in overdetermined systems, where χ^2 of the solution reflects meaningful information. This section will show that when the system becomes underdetermined, χ^2 has a completely different behaviour, and therefore no longer conveys any useful information about the system.

Consider an experiment where m datapoints were obtained, and the experimentalist attempts to use these data to calibrate the coefficients of n degrees of freedom. When $n \geq m$, then this set of data and model together forms a over-/fully-determined system. (In practice, it is often trivial to increase m by taking more datapoints or reducing the complexity of the model to lower n to achieve $n \geq m$.) Assuming that there are no covariance between different datapoints, χ^2 can be calculated by Equation 3.17, repeated below for ease of reference:

$$\chi^2 = \sum_i^m \frac{(x_{\text{meas},i} - x_{\text{model},i})^2}{\text{var}(x_{\text{meas},i})}. \quad (\text{A.1})$$

A.1 Clearing up misconceptions around the minimum χ^2 in an overdetermined v.s. underdetermined problem.




Figure A.1: After performing an experiment where $\nu = 5$, a **minimum** value of χ^2 is obtained by fitting the data to the model. By repeating this experiment 1000 times, these values of minimum χ^2 's is shown to follow the $\chi^2_{\nu=5}$ distribution. The same is repeated for $\nu = 3$, yielding the same result, i.e. simulation agrees with theory.

Any hypothesis has an associated value of χ^2 when compared against the dataset, computed by Equation A.1. However, only the **best-fit** hypothesis would give the **minimum** χ^2 . A competent experimentalist would find this best-fit hypothesis by tuning the model parameters (which should change the values of $x_{\text{model},i} \forall 1 \leq i \leq m$, and that in turn changes the value of χ^2) until the χ^2 cannot be lowered any further.

Assuming the correct model is chosen, then the **minimum** χ^2 can be described by the χ^2 distribution, shown in Figure 3.4. This is exemplified further by Figure A.1. Experimental physicists are often taught that if their experiment has $\frac{\chi^2}{m-n} = 1$ (where $\frac{\chi^2}{m-n}$ is often termed the “reduced χ^2 ”), then it is a good sign that the experiment was successful with correctly measured data and uncertainties, and the model was correctly fitted. These figures reveals the reason: the mean value for the ν^{th} χ^2 distribution (where $\nu = m - n$) is ν when $\nu > 0$. Therefore the average $\frac{\chi^2}{m-n}$ achieved by an experiment after fitting is around 1.

As mentioned in point 3 of Section 3.3: when $\nu > 0$, the mean value of **minimum** achievable $\chi^2 > 0$, and the probability of having $\chi^2 = 0$ is vanishingly small; but when $\nu = 0$, there will be one hypothesis whose $\chi^2 = 0$; and when $\nu < 0$, there will be multiple (infinitely many) points where $\chi^2 = 0$. Therefore this gives us a way to meaningfully continue the definition of the χ^2_{ν} distributions into nonpositive integers: if we define the χ^2_{ν} distribution as the probability distribution of the minimum values of χ^2 attainable when infinitely many experiments are repeatedly performed on a system with degree of determination $= \nu$, then $\chi^2_{\nu} =$ the Dirac-delta function when $\nu \leq 0$, where $\chi^2 = 0$ with 100% probability.

In other words, the minimum χ^2 will *always* equal 0 when the system is fully- or underdetermined (assuming the correct model is chosen). This clears up the misconception that some people hold that “If χ^2 is too small then the experiment has failed/hypothesis is wrong”, as this rule does not apply in an underdetermined experiment. But it also means that the χ^2 of the fitted model no longer carries any useful indication of whether there are any anomalies in the data and their uncertainties.

A.2 Fixing the solution ϕ_{sol} to a user-specified value of χ^2

Some algorithms (MAXED, IMAXED, AMAXED, as well as some implementations of purely iterative algorithms [73]) requires the user to input a value of χ^2 (as defined by Equation 3.20) that the converged solution ϕ_{sol} is expected to meet. However, as we shall see in this appendix, deciding this value of χ^2 for our solution raises even more questions. These algorithms are effectively asking the user “what value of χ^2 would you like the solution to have?”, but this is just as nonsensical as asking a holiday-goer “how much would you like the hotel for your destination to cost?” — it subverts causality.

In an experiment, we are always looking for the lowest χ^2 in a given parameter space, and we would pick the solution that corresponds to that minimum χ^2 , in the same way that when searching for a hotel, one would be searching for the cheapest hotel given a certain set of criteria (destination city, star ratings, etc.), and they would book the hotel offering that lowest price. When the user is asked by the algorithm to input the required χ^2 , a rational user should input “0, if you can please; otherwise as low as possible” — the same answer that a rational consumer should input if asked the question “how much would you like the hotel for this destination to cost?”.

The issue is, of course, unfolding programs generally do not have “as low as possible” provided as one of the input χ^2 options. Instead, experimentalists would have to resort to using trial and error to find the minimum χ^2 that successfully converge without crashing the program, as that indicates a solution that matches the required χ^2 exists and can be found. This is a problematic approach for many reasons, one of which is that it is impractical at finding the true minimum χ^2 .

A.3 A better approach: Tikhonov Regularization

Now, some experimentalists might object to the notion that we want the solution ϕ_{sol} to have a $\chi^2 =$ to the minimum possible χ^2 implied in the previous sections, as this means we are prioritising the reduction of χ^2 over the alignment of ϕ_{sol} to any *a priori* information (such as by reducing its regularising value) (Section 3.1.3). And this is a very valid argument — this forces all *a priori* information provided to take a back-seat when it comes to deciding the final solution ϕ_{sol} . But this is a natural consequence of how MAXED, IMAXED and AMAXED etc. algorithms has formulated their premises: they intend to minimise the regularising function while fixing the χ^2 , completely decoupling the effect of lowering χ^2 from the dependency on the *a priori* information in its reasoning, hence yielding the perverse directive to “lower χ^2 as much as possible” when setting the desired χ^2 .

I argue that to resolve this dilemma, instead of trying to figure out what is the correct value of χ^2 to fix the optimization at, we should use a loss function that minimises both χ^2 and the regularising function simultaneously, without forcing either one into the back-seat. Without resorting to more complicated multi-objective optimization methods, the simplest way to do so is via Tikhonov regularization algorithms such AMAXED-Regularization (Section 4.3.3). Going back to the hotel analogy, using such an algorithm would be the same as the holiday-goer requesting that the program to “find a hotel that is both cheap and good quality”.

Appendix B

IMAXED

B.1 Algorithm derivation

Instead of optimising the scalar-valued function obtained by “integrating” the vector-valued function (to obtain 4.22); we take one step back, and search for the root(s) of the vector-valued function (4.20). Also, we extended the definition of Equation 4.22 a little such that it can account for covariances when calculating χ^2 . This modifies Equation 4.22 and 4.20 into

$$Z = -\phi'_0 \cdot \exp(-\lambda \cdot \underline{\underline{\mathbf{R}}}) - \sqrt{\Omega(\lambda \cdot \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \lambda)} - \lambda \cdot \mathbf{N}, \quad (\text{B.1})$$

$$\nabla(Z) = \mathbf{0} = \underline{\underline{\mathbf{R}}}(\phi'_0 \circ \exp(-\lambda \cdot \underline{\underline{\mathbf{R}}})) - \mathbf{N} - \left(\sqrt{\frac{\Omega}{\lambda \cdot \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \lambda}} \right) \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \lambda, \quad (\text{B.2})$$

Where $\underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} = (\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}}))^{-1}$ is the inverse of the covariance matrix of the measurements $\mathbf{N}_{\text{meas}}$, and is a symmetric matrix. In case of no off-diagonal covariance values, $\underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}$ is simply $\underline{\underline{\text{Diag}}}(\sigma_N \circ \sigma_N)$.

$\nabla(Z)$ ($= \frac{\partial Z}{\partial \lambda}$) is then a vector-valued function, with vector input λ .

Using Newton’s method, we can update our current vector $\lambda \leftarrow \lambda - (\underline{\underline{\nabla^2}}(Z))^{-1}(\nabla(Z))$ at each iteration to approach the final solution.

B.1.1 Newton’s method

The second-order derivative (i.e. Hessian matrix of Z) is given as

$$\begin{aligned} \underline{\underline{\nabla^2}}(Z) = & -\underline{\underline{\mathbf{R}}} \cdot (\phi'_0 \circ \exp(-\underline{\underline{\boldsymbol{\lambda}}} \cdot \underline{\underline{\mathbf{R}}}) \cdot \underline{\underline{\mathbf{R}}}^T) \\ & + \sqrt{\frac{\Omega}{(\underline{\underline{\boldsymbol{\lambda}}} \cdot \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}})^3}} \frac{(\underline{\underline{\boldsymbol{\lambda}}} \cdot \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}) \otimes (\underline{\underline{\mathbf{S}}}_{N_{\text{meas}}} \underline{\underline{\boldsymbol{\lambda}}})}{\underline{\underline{\boldsymbol{\lambda}}} \cdot \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}} + \sqrt{\frac{\Omega}{\underline{\underline{\boldsymbol{\lambda}}} \cdot \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}}} \underline{\underline{\mathbf{S}}}_{N_{\text{meas}}}, \end{aligned} \quad (\text{B.3})$$

where $\otimes$ is the outerproduct, i.e. if $\underline{\underline{\mathbf{c}}} = \underline{\underline{\mathbf{a}}}^T \otimes \underline{\underline{\mathbf{b}}}$ then $c_{ij} = a_i b_j$. This matrix is negative definite matrix and hence invertible, but it is difficult to show its definite-ness by examination of the equation alone. For more details see Section B.2.

Since the gradient (Equation B.2) and its derivative (Equation B.3) are already calculated, When using the Newton's method, the difference in $\underline{\underline{\boldsymbol{\lambda}}}$ consecutive steps can be expressed as $\underline{\underline{\Delta\boldsymbol{\lambda}}}$:

$$\underline{\underline{\Delta\boldsymbol{\lambda}}} = -(\underline{\underline{\nabla^2}}(Z))^{-1}(\underline{\underline{\nabla}}(Z)), \quad (\text{B.4})$$

So that the next iteration's value of $\underline{\underline{\boldsymbol{\lambda}}} =$

$$\underline{\underline{\boldsymbol{\lambda}}} \leftarrow \underline{\underline{\boldsymbol{\lambda}}} + \beta \underline{\underline{\Delta\boldsymbol{\lambda}}}. \quad (\text{B.5})$$

where $\beta = 1$ for the unmodified Newton's method.

However, since the Newton's method only takes into account 2^{nd} order derivative information about the function Z at the current $\underline{\underline{\boldsymbol{\lambda}}}$, it will only update $\underline{\underline{\boldsymbol{\lambda}}}$ to the stationary point identified by the local second-order approximation of the function Z . Error in approximation of Z due to higher order dependence on $Z(\underline{\underline{\boldsymbol{\lambda}}})$ is not accounted for, so the updated $\underline{\underline{\boldsymbol{\lambda}}}$ may not lie exactly at the true global maximum. We correct for this by adjusting the step size: by finding the value of β that gives the maximum $Z(\underline{\underline{\boldsymbol{\lambda}}} + \beta \underline{\underline{\Delta\boldsymbol{\lambda}}})$.

B.1.2 Line search (Wolfe condition)

To search for the correct value of step size β which will give the maximum Z along the direction $\underline{\underline{\Delta\boldsymbol{\lambda}}}$, we must find the value of β where $\frac{\partial Z(\underline{\underline{\boldsymbol{\lambda}}} + \beta \underline{\underline{\Delta\boldsymbol{\lambda}}})}{\partial \beta} = 0$.

$$\begin{aligned}
\frac{\partial Z(\boldsymbol{\lambda} + \beta \Delta \boldsymbol{\lambda})}{\partial \beta} = 0 = & + \phi'_0 \circ \exp(-\boldsymbol{\lambda} \cdot \underline{\mathbf{R}}) \circ (\Delta \boldsymbol{\lambda} \underline{\mathbf{R}}) \cdot \exp(-\beta \Delta \boldsymbol{\lambda} \underline{\mathbf{R}}) \\
& - \sqrt{\Omega} \frac{(\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda}) + \beta (\Delta \boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda})}{\sqrt{(\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \boldsymbol{\lambda}) + 2\beta (\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda}) + \beta^2 (\Delta \boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda})}} \quad (\text{B.6}) \\
& - \mathbf{N} \cdot \Delta \boldsymbol{\lambda}.
\end{aligned}$$

Using the standard `root_scalar` algorithm from `scipy.optimize`, the correct value of β can be found. We can confirm that the new β does indeed result in a maximum value of Z by checking its second derivative, which should be negative:

$$\begin{aligned}
& \frac{\partial^2 Z(\boldsymbol{\lambda} + \beta \Delta \boldsymbol{\lambda})}{\partial \beta^2} \\
= & - \phi'_0 \circ \exp(-\boldsymbol{\lambda} \cdot \underline{\mathbf{R}}) \circ (\Delta \boldsymbol{\lambda} \underline{\mathbf{R}}) \circ (\Delta \boldsymbol{\lambda} \underline{\mathbf{R}}) \cdot \exp(-\beta \Delta \boldsymbol{\lambda} \underline{\mathbf{R}}) \quad (\text{B.7}) \\
& - \sqrt{\Omega} \frac{(\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \boldsymbol{\lambda})(\Delta \boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda}) - (\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda})^2}{\sqrt{(\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \boldsymbol{\lambda}) + 2\beta (\boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda}) + \beta^2 (\Delta \boldsymbol{\lambda} \cdot \underline{\mathbf{S}}_{N_{\text{meas}}} \Delta \boldsymbol{\lambda})}^3}.
\end{aligned}$$

B.2 Guaranteed convergence to a maximum

The technique of Lagrangian multipliers is used when creating the IMAXED (and MAXED) unfolding algorithm. The Lagrangian function embodies constraints and an objective function, such that all stationary points of the Lagrangian function are local maximum, saddle point, or minimum in the objective function obeying the constraints laid out. However, we are interested in only the minimiser of cross-entropy, not local maxima or saddle points. This minimiser of cross-entropy would also be the maximiser of Z .

To ensure that the solution $\boldsymbol{\lambda}$ is a maximiser of Z , it only requires us to check that the second derivative of Z with respect to $\boldsymbol{\lambda}$ is a negative definite matrix, i.e. it must have negative curvature in every direction. This is done by taking the singular decomposition of the Hessian matrix in Equation B.7, and confirming that all of its singular values are negative.

B.3 Uncertainty propagation

The uncertainty of the solution obtained by IMAXED (as well as MAXED) can be analytically derived. Consider the curvature of Z in $\boldsymbol{\lambda}$ space, given by Equation B.3:

$$\frac{\partial^2 Z}{\partial \boldsymbol{\lambda}^2} = \underline{\underline{\nabla^2}}(Z). \quad (\text{B.8})$$

B.3 Uncertainty propagation

We can perform a change of basis to get the curvature in ϕ space:

$$\frac{\partial^2 Z}{\partial \phi_{\text{sol},i} \partial \phi_{\text{sol},j}} = \sum_{\kappa}^m \frac{\partial \lambda_{\kappa}}{\partial \phi_{\text{sol},i}} \sum_k^m \frac{\partial^2 Z}{\partial \lambda_{\kappa} \partial \lambda_k} \frac{\partial \lambda_k}{\partial \phi_{\text{sol},j}}. \quad (\text{B.9})$$

To make this change of basis easier, we define a matrix $\underline{\underline{\mathbf{J}}}$, such that

$$J_{ki} = \frac{\sum_j^n \phi_{0,j}}{R_{ki} \phi_{\text{sol},i}}, \quad (\text{B.10})$$

or, in a more concise form,

$$\underline{\underline{\mathbf{J}}} = \frac{\sum_j^n \phi_{0,j}}{\underline{\underline{\mathbf{R}}} \circ \phi_{\text{sol}}}, \quad (\text{B.11})$$

thereby turning Equation B.9 into a more cleaner, easier to read form:

$$\frac{\partial^2 Z}{\partial \phi_{\text{sol}}^2} = \left(\frac{\partial \boldsymbol{\lambda}}{\partial \phi_{\text{sol}}} \right)^T \left(\frac{\partial^2 Z}{\partial \boldsymbol{\lambda}^2} \right) \left(\frac{\partial \boldsymbol{\lambda}}{\partial \phi_{\text{sol}}} \right) \quad (\text{B.12})$$

$$= \underline{\underline{\mathbf{J}}}^T \left(\frac{\partial^2 Z}{\partial \boldsymbol{\lambda}^2} \right) \underline{\underline{\mathbf{J}}}. \quad (\text{B.13})$$

Finally, we can get the inverse of $\underline{\underline{\text{cov}}}(\phi_{\text{sol}})$, denoted with $\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}$ in terms of $\underline{\underline{\mathbf{J}}}$ and the curvature of Z :

$$\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}} = - \frac{\partial^2 Z}{\partial \phi_{\text{sol}}^2} \quad (\text{B.14})$$

$$= - \underline{\underline{\mathbf{J}}}^T \underline{\underline{\nabla^2}}(Z) \underline{\underline{\mathbf{J}}}. \quad (\text{B.15})$$

For the sake of brevity the Hessian $\underline{\underline{\nabla^2}}(Z)$ in Equation B.15 is not explicitly written, but can be expanded from its definition in Equation B.3 if necessary.

However, there is a problem: while $\frac{\partial^2 Z}{\partial \boldsymbol{\lambda}^2}$ is guaranteed to be a full-rank matrix (i.e. has rank = m , meaning the function Z has negative curvature in every direction in $\boldsymbol{\lambda}$ space) due to its negative definite nature proved in Section B.3, $\frac{\partial^2 Z}{\partial \phi_{\text{sol}}^2}$ is not. If the system is underdetermined, i.e. $m < n$, then $\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}$ will still be a singular matrix as $\underline{\underline{\mathbf{J}}}$ has fewer than full rank (rank = m as opposed to rank = n).

Despite this, in practice, experimental physicists are often interested in questions such as “what is the variance/standard deviation on bin i ”, rather than “what is the full covariance matrix”. Fortunately, we can still obtain the variance on each bin of the unfolded neutron

B.3 Uncertainty propagation

spectrum by taking the main diagonal of the pseudo-inverse $\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}$:

$$\text{var}(\phi_{\text{sol},i}) = \mathbf{Diag} \left((\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}})^\dagger \right)_i, \quad (\text{B.16})$$

achieving uncertainty propagation. In fact, this method can be extended to obtain the entire covariance matrix $\underline{\underline{\mathbf{cov}}}(\phi_{\text{sol}})$ as shown in Equation 4.48.

Appendix C

AMAXED

C.1 Algorithm derivation

We begin by attempting to minimise the scalar value D_{KL} as defined by Equation 4.34, while adhering to the condition that $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})$ as defined by Equation 3.19. Using μ as the Lagrangian multiplier, we obtain the Lagrangian expression:

$$\mathcal{L} = \mu \left(\left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\mathbf{S}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) - \Omega \right) - \phi'_0 \cdot \log\left(\frac{\phi'_{\text{sol}}}{\phi'_0}\right), \quad (\text{C.1})$$

where $\phi'_{\text{sol}} = \frac{\phi_{\text{sol}}}{\phi_{\text{sol}} \cdot \mathbf{1}^n}$ is the normalized distribution of neutrons in the solution spectrum.

$$\begin{aligned} \mathcal{L} = & \mu \left(\left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\mathbf{S}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) - \Omega \right) \\ & - \phi'_0 \cdot \log(\phi'_{\text{sol}}) + \phi'_0 \cdot \log(\phi'_0). \end{aligned} \quad (\text{C.2})$$

Removing the constant term at the end as it will not be affecting our optimization, and expanding the penultimate term,

$$\begin{aligned} \mathcal{L} = & \mu \left(\left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\mathbf{S}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) - \Omega \right) \\ & - \phi'_0 \cdot \log(\phi_{\text{sol}}) + \left(\log(\phi_{\text{sol}} \cdot \mathbf{1}^n) \right) \phi'_0 \cdot \mathbf{1}^n. \end{aligned} \quad (\text{C.3})$$

Taking advantage of the fact that ϕ_0 is also normalized, i.e. $\phi_0 \cdot \mathbf{1}^n = 1$,

$$\begin{aligned} \mathcal{L}(\phi_{\text{sol}}, \mu) = & \mu \left(\left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\mathbf{S}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\mathbf{R}}\phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) - \Omega \right) \\ & - \phi'_0 \cdot \log(\phi_{\text{sol}}) + \log(\phi_{\text{sol}} \cdot \mathbf{1}^n). \end{aligned} \quad (\text{C.4})$$

C.1 Algorithm derivation

We then take the first and second derivative of $\mathcal{L}$ against the vector of (ϕ_{sol}, μ) , which is an $(n+1)$ -dimensional variable with the first n dimensions = $\phi_{\text{sol},i}$ and the final $n+1^{\text{th}}$ dimension representing μ .

To make the following derivations shorter, we use the following shorthand:

$$\mathbf{a} = \mathbf{N}_{\text{meas}} \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \quad (\text{C.5})$$

$$\mathbf{b} = \phi_{\text{sol}} \cdot \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}}. \quad (\text{C.6})$$

The set of equations to compute the first derivative is given by:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \phi_{\text{sol}}} &= \frac{\mathbf{1}^n}{\phi_{\text{sol}} \cdot \mathbf{1}^n} - \frac{\phi'_0}{\phi_{\text{sol}}} + 2\mu(\mathbf{b} - \mathbf{a}) \\ \frac{\partial \mathcal{L}}{\partial \mu} &= (\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} (\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}}) - \Omega. \end{aligned} \quad (\text{C.7})$$

The set of equations to compute the second derivatives is given by:

$$\begin{aligned} \frac{\partial^2 \mathcal{L}}{\partial \phi_{\text{sol}}^2} &= -\left(\frac{\mathbf{1}^n \otimes \mathbf{1}^n}{\phi_{\text{sol}} \cdot \phi_{\text{sol}}} \right) + \underline{\underline{\text{Diag}}} \left(\frac{\phi'_0}{\phi_{\text{sol}} \circ \phi_{\text{sol}}} \right) + 2\mu \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \\ \frac{\partial^2 \mathcal{L}}{\partial \phi_{\text{sol}} \partial \mu} &= \left(\frac{\partial^2 \mathcal{L}}{\partial \mu \partial \phi_{\text{sol}}} \right)^T = 2(\mathbf{b} - \mathbf{a}) \\ \frac{\partial^2 \mathcal{L}}{\partial \mu^2} &= 0. \end{aligned} \quad (\text{C.8})$$

Starting with an initial vector $(\phi_{\text{sol}}, \mu) = (\phi_0, \mu_0)$, where

$$\mu_0 = \frac{(\mathbf{b} - 2\mathbf{a}) \cdot \phi'_0 + \frac{1}{\phi_{\text{sol}} \cdot \mathbf{1}} (\mathbf{N}_{\text{meas}} \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \mathbf{N}_{\text{meas}} - \Omega)}{2[(\mathbf{b} \circ \mathbf{b}) \cdot \phi_{\text{sol}} - 3(\mathbf{b} \circ \mathbf{a}) \cdot \phi_{\text{sol}} + 2(\mathbf{a} \circ \mathbf{a}) \cdot \phi_{\text{sol}}]}, \quad (\text{C.9})$$

same as Equation B.4 and B.5, we can begin iteratively updating the solution spectrum by

$$(\Delta \phi, \Delta \mu) = -\left(\frac{\partial^2 \mathcal{L}(\phi_{\text{sol}}, \mu)}{\partial (\phi_{\text{sol}}, \mu)^2} \right)^{-1} \frac{\partial \mathcal{L}(\phi_{\text{sol}}, \mu)}{\partial (\phi_{\text{sol}}, \mu)} \quad (\text{C.10})$$

$$(\phi_{\text{sol}}, \mu) \leftarrow (\phi_{\text{sol}} + \beta \Delta \phi, \mu + \beta \Delta \mu). \quad (\text{C.11})$$

C.1.1 Line search (Wolfe condition)

To search for the correct value of step size β which will give the minimum $\mathcal{L}$ along the direction $(\Delta\phi, \Delta\mu)$, it requires $\frac{\partial \mathcal{L}}{\partial \beta} = 0$:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \beta} = & \frac{\Delta\phi \cdot \mathbf{1}^n}{(\phi + \beta\Delta\phi) \cdot \mathbf{1}^n} - \frac{\phi'_0}{\phi + \beta\Delta\phi} \cdot \Delta\phi \\ & + \Delta\mu \left[\left(\underline{\underline{\mathbf{R}}}(\phi + \beta\Delta\phi) - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\underline{\mathbf{R}}}(\phi + \beta\Delta\phi) - \mathbf{N}_{\text{meas}} \right) - \Omega \right] \\ & + 2(\mu + \beta\Delta\mu) \left((\mathbf{b} - \mathbf{a}) \cdot \Delta\phi + \beta\Delta\phi \cdot \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \Delta\phi \right) \end{aligned} \quad (\text{C.12})$$

Same as in Section B.1.2, we can confirm this point is a minimum by checking that $\frac{\partial^2 \mathcal{L}}{\partial \beta^2} > 0$:

$$\frac{\partial^2 \mathcal{L}}{\partial \beta^2} = - \left(\frac{\Delta\phi \cdot \mathbf{1}^n}{(\phi + \beta\Delta\phi) \cdot \mathbf{1}^n} \right)^2 + \frac{\phi'_0 \circ \Delta\phi}{(\phi + \beta\Delta\phi) \circ (\phi + \beta\Delta\phi)} \cdot \Delta\phi \quad (\text{C.13})$$

$$\begin{aligned} & + 2\Delta\mu (\underline{\underline{\mathbf{R}}}(\phi + \beta\Delta\phi) - \mathbf{N}_{\text{meas}}) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \Delta\phi \\ & + 2\Delta\mu \left((\mathbf{b} - \mathbf{a}) \cdot \Delta\phi + \beta\Delta\phi \cdot \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \Delta\phi \right) \\ & + 2(\mu + \beta\Delta\mu) \left(\Delta\phi \cdot \underline{\underline{\mathbf{R}}}^T \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} \Delta\phi \right) \end{aligned} \quad (\text{C.14})$$

C.2 Guaranteed convergence to a minimum

As explained in Section B.2, the Lagrangian function alone does not guarantee that the converged solution is a global minimum. However, we have set up the algorithm in a way that ensures it converges on the stationary point nearest to ϕ_0 , which must be the global minimum. This is because we start the optimization at ϕ_0 , which is a point of lowest D_{KL} ; and we take steps that increase D_{KL} to approach the constraint $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}}) = \Omega$, using the Wolfe condition (Section C.1.1) to avoid overshooting/missing the desired minimum.

AMAXED also checks that the solution obtained is indeed a minimum by ensuring the Hessian matrix in Equation C.8 is positive definite, i.e. all its singular values are positive.

C.3 Uncertainty Propagation

In this algorithm, the uncertainty on the solution spectrum can be obtained directly from Equation C.8:

$$\left(\underline{\underline{\text{cov}}}(\phi_{\text{sol}}) \right)^{-1} = \underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}} = \frac{\partial^2 \mathcal{L}}{\partial \phi_{\text{sol}}^2}, \quad (\text{C.15})$$

C.3 Uncertainty Propagation

which we have ensured to be nonsingular (and in fact, is positive definite) from Section C.2. Therefore we can obtain $\underline{\underline{\mathbf{cov}}}(\phi_{\text{sol}})$ by directly inverting the matrix:

$$\underline{\underline{\mathbf{cov}}}(\phi_{\text{sol}}) = \left(\frac{\partial^2 \mathcal{L}}{\partial \phi_{\text{sol}}^2} \right)^{-1}. \quad (\text{C.16})$$

Appendix D

AMAXED-Regularization

D.1 Algorithm derivation

The compact form of the loss function chosen is shown in Equation 4.36. Expanding it, we have:

$$\text{loss} = \left(\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) + \tau \phi'_0 \cdot \log\left(\frac{\phi'_0}{\phi'_{\text{sol}}}\right), \quad (\text{D.1})$$

which can expanded as

$$\text{loss} = \left(\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \left(\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) - \tau \phi'_0 \cdot \log(\phi_{\text{sol}}) + \tau \log(\phi_{\text{sol}} \cdot \mathbf{1}^n) + \tau \phi'_0 \cdot \log(\phi'_0). \quad (\text{D.2})$$

The last term is a constant multiplied by τ .

Keeping τ constant, we can find the gradient and curvature of the loss function as:

$$\underline{\underline{\nabla}}(\text{loss}) = 2 \left(\underline{\underline{\mathbf{R}}} \phi_{\text{sol}} - \mathbf{N}_{\text{meas}} \right) \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} - \tau \frac{\phi'_0}{\phi_{\text{sol}}} + \tau \frac{\mathbf{1}^n}{\phi_{\text{sol}} \cdot \mathbf{1}^n} \quad (\text{D.3})$$

$$\underline{\underline{\nabla}}^2(\text{loss}) = 2 \underline{\underline{\mathbf{R}}}^T \cdot \underline{\underline{\mathbf{S}}}_{\mathbf{N}_{\text{meas}}} \underline{\underline{\mathbf{R}}} + \tau \underline{\underline{\text{Diag}}} \left(\frac{\phi'_0}{\phi_{\text{sol}} \circ \phi_{\text{sol}}} \right) - \tau \left(\frac{1}{\phi_{\text{sol}} \cdot \mathbf{1}^n} \right) \underline{\underline{\mathbf{1}^n \otimes \mathbf{1}^n}}. \quad (\text{D.4})$$

Beginning the first iteration at $\phi_{\text{sol}} = \phi_0$, we can apply Newton's method iteratively to approach a stationary point:

$$\Delta \phi_{\text{sol}} = -\underline{\underline{\nabla}}^2(\text{loss})^{-1} \underline{\underline{\nabla}}(\text{loss}) \quad (\text{D.5})$$

$$\phi_{\text{sol}} \leftarrow \phi_{\text{sol}} + \beta \Delta \phi_{\text{sol}}. \quad (\text{D.6})$$

D.1.1 Line search (Wolfe condition)

Similar to previous section, $\frac{\partial \text{loss}}{\partial \beta}$ and $\frac{\partial^2 \text{loss}}{\partial \beta^2}$ were be derived, but to include them in this section would add length to this thesis without without adding any insight. Therefore it was excluded, but the keen reader can still find them at [129].

D.2 Uniqueness of minimum

Before we begin our prove, it is useful to bear in mind that ϕ and τ can only take positive values, i.e.

$$\phi \in (\mathbb{R}^+)^n, \quad \tau \in \mathbb{R}^+. \quad (\text{D.7})$$

Consider the two components of the loss function, $\chi^2(\mathbf{N}_{\text{sol}}, \mathbf{N}_{\text{meas}})$ and $D_{KL}(\phi'_0, \phi'_{\text{sol}})$. Their contribution to the loss function, gradient of the loss function, and curvature of the loss function are shown in Equation D.2, D.3, and D.4 respectively. The 1st terms in these equations are the contributions from χ^2 , while the 2nd and 3rd terms in these equations are the contributions from D_{KL} . To find the singular and nonsingular directions of the matrices in Equation D.4, we can interrogate the matrix by computing the scalar $(\mathbf{v} \cdot \underline{\underline{\mathbf{M}}}\mathbf{v})$: if $\mathbf{v}$ is a singular direction of matrix $\underline{\underline{\mathbf{M}}}$, then $\mathbf{v} \cdot \underline{\underline{\mathbf{M}}}\mathbf{v}$ should equal to 0. Otherwise, $\mathbf{v}$ is a nonsingular direction.

1. In any given nonsingular directions of the χ^2 component, the loss-function contribution is quadratic in nature, i.e. the gradient in the loss function contributed by the χ^2 component linearly increases proportional to the deviation from the point of minimum χ^2 on that line. In any singular direction of χ^2 , i.e. a direction $\mathbf{n}$ where the dot products of $\mathbf{n}$ with each of the nonsingular directions are all equal to 0, the χ^2 contribution to the loss function gradient remains unchanged. For this term, its singular and nonsingular directions are independent of where it is in the domain of ϕ .
2. Meanwhile, the loss function gradient contribution by D_{KL} has only one singular direction. This direction, expressed as a normalized vector is $\frac{\phi}{\|\phi\|_2}$, which is evidently dependent on ϕ . In all other directions, the D_{KL} contribution to the loss function gradient (and the derivative of its gradient) increases strictly monotonically.

Then we can begin our proof. Assume that we have already found the point of lowest loss value, $\phi = \phi_{\text{sol}}$. By definition, $\nabla(\text{loss})|_{\phi_{\text{sol}}} = \mathbf{0}$. Taking a step $\Delta\phi$ in any direction except $\pm \frac{\phi}{\|\phi\|_2}$ increases the loss value strictly monotonically, due to the contribution of D_{KL} as explained

D.3 Uncertainty Propagation

in point 2, compounded by the linear increase in gradient contributed by χ^2 as explained in point 1. Therefore, there can be no other stationary points on any other lines passing through ϕ_{sol} , pointed any other directions than $\pm \frac{\phi}{\|\phi\|_2}$.

However, one (or rather, infinitely many) candidate(s) remains that may interfere with the uniqueness of our global minimum: If the loss function has a constant value along the line $\phi = \alpha \phi_{\text{sol}}$ (the singular direction of D_{KL}), then we may have potentially infinitely many points that can be called the “global minima”. But we can prove that this will never happen in practice: the only way for this to happen is if $\frac{\phi}{\|\phi\|_2}$ is one of the singular directions of the χ^2 contribution to the loss function gradient, as explained in point 1. However, this will never happen in practice as the response matrix is nonnegative, i.e. $(\mathbf{R}_k \cdot \phi_{\text{sol}}) > 0 \forall 1 \leq k \leq m$, where $\mathbf{R}_k$ is the k^{th} row of the response matrix, and ϕ_{sol} must be a positive vector by definition.

To conclude, the properties of D_{KL} ensures the Hessian matrix is nonsingular in all directions except in the direction of ϕ_{sol} ; and we have also proven that in practice, this direction is also nonsingular due to the nonnegative nature of $\underline{\mathbf{R}}$. Therefore there should only be one stationary point in the loss function of over the domain of $\phi \in (\mathbb{R}^+)^n$.

D.3 Uncertainty Propagation

As before, the uncertainty about the solution ϕ_{sol} is given by $\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}}$:

$$\underline{\underline{\mathbf{S}}}_{\phi_{\text{sol}}} = \frac{\partial^2(\text{loss})}{\partial \phi_{\text{sol}}^2} \quad (\text{D.8})$$

$$\underline{\underline{\text{cov}}}(\phi_{\text{sol}}) = \left(\underline{\underline{\nabla^2}}(\text{loss}) \right)^{-1} \quad (\text{D.9})$$

And we have proven that the Hessian $\underline{\underline{\nabla^2}}(\text{loss})$ is a nonsingular matrix when ϕ_{sol} is a positive vector in Section D.2, therefore it is invertible, so there should be no issue calculating its inverse in Equation D.9.

Appendix E

SAND-II weight matrix

The following chapter records the steps in the SAND-II algorithm used to calculate the weight matrix $\underline{\underline{\mathbf{W}}}$. Due to its rather lengthy equations that detracts from the main point of the text in the thesis, it was decided that it should be placed in the appendix.

While early version of SAND-II ([54]) calculates a weight matrix simply using

$$\underline{\underline{\mathbf{W}}} = \frac{1}{N_{\text{meas}}} \circ \underline{\underline{\mathbf{R}}} \circ \phi_{\text{sol}}, \quad (\text{E.1})$$

the same equation used to calculate the weight matrix of GRAVEL (Equation 4.28), it was deemed unsatisfactory as this method has the tendency of amplifying “spurious structures” in the neutron spectrum, meaning that the solution neutron spectra inherit the same peaks as the resonance peaks found in the cross sections of the activation foils used to measure them. Therefore the latest versions of SAND-II, created in 1968, requires user to also put in a smoothing degree, an integer N_s , to force the weight matrix to apply extra constraint to stabilise the resulting neutron spectrum.

To remind readers about the weight matrix, it is a matrix with m rows $\times n$ columns. The first column and last column of the weight matrix $\underline{\underline{\mathbf{W}}}$ is obtained by an operation akin to convolution with a kernel of the form $[-1, 2, 5, 2, -1]$:

$$W_{i1} = \frac{5R_{i,1}\phi_{\text{sol},1} + 2R_{i,2}\phi_{\text{sol},2} - R_{i,3}\phi_{\text{sol},3}}{6N_{\text{sol},i}} \quad (\text{E.2})$$

$$W_{im} = \frac{5R_{i,m}\phi_{\text{sol},m} + 2R_{i,(m-1)}\phi_{\text{sol},m-1} - R_{i,(m-2)}\phi_{\text{sol},m-2}}{6N_{\text{sol},i}}. \quad (\text{E.3})$$

Then, using the user-defined smoothing factor N_s , the remaining leftmost $\frac{N_s-1}{2}$ columns are given as

$$W_{ij} = \frac{N_{\text{sol},i}}{(2j-1)R_{i,j}\phi_{\text{sol},j}} \quad \text{for } 1 < j \leq \frac{N_s-1}{2}, \quad (\text{E.4})$$

And the remaining rightmost $\frac{N_s-1}{2}$ columns are given as

$$W_{ij} = \frac{N_{\text{sol},i}}{(2m-2j+1)R_{i,j}\phi_{\text{sol},j}} \quad \text{for } (m - \frac{N_s-1}{2}) < j < m. \quad (\text{E.5})$$

And then the remaining columns at the center, not covered by the edge cases already mentioned, are given as:

$$W_{ij} = \frac{N_{\text{sol},i}}{(N_s)R_{i,j}\phi_{\text{sol},j}} \quad \text{for } \frac{N_s-1}{2} < j \leq (m - \frac{N_s-1}{2}). \quad (\text{E.6})$$

Please note that the series of steps in Equation E.2 to E.6 above are simply re-arrangement of equations in Appendix E of reference [157] and Appendix C of [158], and do not constitute original work carried out in this PhD investigation. They are simply placed in this appendix to allow readers to more easily reference and compare different unfolding algorithms, without having to consult a separate document.

The `unfolding` suite implementation of SAND-II allows for two run-mode, `None` or `smoothing`. The default run-mode `None` follows the earlier version of SAND-II whose weight matrix is shown in Equation E.1, while the `smoothing` mode's weight matrix is shown in Equation E.2 to E.6. In the main text of this thesis, when SAND-II is used, it runs in the `None` mode unless otherwise specified, in which case the degree of smoothing N_s would be provided as well. Note that Equation E.2 and E.3 requires the neutron spectrum to have at least 3 bins.

Appendix F

Loss function of GRAVEL

The original document of GRAVEL ([83]) shows that the author Matzke set out with the goal of minimising

$$\text{loss} = \left(\log(\mathbf{N}_{\text{sol}}) - \log(\mathbf{N}) \right) \cdot \underline{\underline{\mathbf{S}}}_{\log(\mathbf{N}_{\text{meas}})} \left(\log(\mathbf{N}_{\text{sol}}) - \log(\mathbf{N}) \right), \quad (\text{F.1})$$

where $\underline{\underline{\mathbf{S}}}_{\log(\mathbf{N}_{\text{meas}})}$ is the inverse of the covariance matrix associated with the errors of $\mathbf{N}_{\text{meas}}$. The k^{th} main diagonal elements of this covariance matrix is given by $\frac{\text{variance on } (N_{\text{meas},k})}{(N_{\text{meas},k})^2}$.

However he attempted to invert a singular (i.e. non-invertible) matrix $\underline{\underline{\mathbf{B}}}$ between Equation (34) on page 11 and Equation (37) on page 12, by incorrectly asserting $(\mathbf{B}^{-1})_{ij} = \frac{1}{B_{ij}}$, such that it has deviated from its initial goal of minimising the expression in Equation F.1, so it is now unclear what form of loss function it ends up minimising. Regardless, the code still function sufficiently well as an unfolding program to be used in many studies (including fusion neutronics) and generate citations, hence it is noteworthy enough to be included in this thesis.

Appendix G

Decomposing branching decay chains

A decay chain of a root nuclide A is used to refer to the graph of all descendants that are generated when A decays, and the daughter nuclides of A decays, and their daughter nuclides decays, etc. It is perhaps more apt to call such a graph a decay tree due to the possibility of branches (one nuclide producing more than 1 daughter) occurring in the chain, thus in general, a decay “chain” is not linear.

A decay chain is always acyclic as nuclides always decay from nuclides of higher potential energy to nuclides of lower potential energy. This property means that there is a finite number of ways to get the root node (the root nuclide) to any other nodes (descendant nuclides).

We can then linearise this directed acyclic graph (DAG) by performing depth-first search on it to discover the various components of the decay chain. For the example in Figure G.1, this is broken down in to the set of constituent chains in (G.1):




Figure G.1: An example decay chain, where A is the root nuclide.

$$\begin{aligned}
 & (f_{AB}f_{BC}f_{CE})\text{times of } \textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{C} \xrightarrow{\lambda_C} \textcircled{E} \xrightarrow{\lambda_E} \\
 & (f_{AB}f_{BC}f_{CF}f_{FE})\text{times of } \textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{C} \xrightarrow{\lambda_C} \textcircled{F} \xrightarrow{\lambda_F} \textcircled{E} \xrightarrow{\lambda_E} \\
 & (f_{AB}f_{BD}f_{DC}f_{CE})\text{times of } \textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{D} \xrightarrow{\lambda_D} \textcircled{C} \xrightarrow{\lambda_C} \textcircled{E} \xrightarrow{\lambda_E} \\
 & (f_{AB}f_{BD}f_{DC}f_{CF}f_{FE})\text{times of } \textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{D} \xrightarrow{\lambda_D} \textcircled{C} \xrightarrow{\lambda_C} \textcircled{F} \xrightarrow{\lambda_F} \textcircled{E} \xrightarrow{\lambda_E} \\
 & (f_{AB}f_{BD}f_{DF}f_{FE})\text{times of } \textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{D} \xrightarrow{\lambda_D} \textcircled{F} \xrightarrow{\lambda_F} \textcircled{E} \xrightarrow{\lambda_E}
 \end{aligned} \tag{G.1}$$

where f_{YZ} is the branching fraction of nuclide $\textcircled{Y}$ decaying to nuclide $\textcircled{Z}$. The branching fraction of all decays from a nuclide must add up to 1.0, hence $f_{AB} = 1.0 = f_{FE}$ as they are the only decay pathways that $\textcircled{A}$ and $\textcircled{E}$ can take.

By breaking the decay chain down to its constituent parts as shown in (G.1), we can identify all decay pathways that leads to a specific nuclide. For example, all pathways that leads to is listed below:

1. 4th generation in the chain $(f_{AB}f_{BC}f_{CF})$ times of $\textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{C} \xrightarrow{\lambda_C} \textcircled{F} \xrightarrow{\lambda_F}$
2. 5th generation in the chain $(f_{AB}f_{BD}f_{DC}f_{CF})$ times of $\textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{D} \xrightarrow{\lambda_D} \textcircled{C} \xrightarrow{\lambda_C} \textcircled{F} \xrightarrow{\lambda_F}$
3. 4th generation in the chain $(f_{AB}f_{BD}f_{DF})$ times of $\textcircled{A} \xrightarrow{\lambda_A} \textcircled{B} \xrightarrow{\lambda_B} \textcircled{D} \xrightarrow{\lambda_D} \textcircled{F} \xrightarrow{\lambda_F}$

The population of F at time $t > 0$ can then be calculated as the sum of several Bateman equations $\mathcal{B}$:

$$\begin{aligned}
 P_F(t) &= P_{A,0} && f_{AB}f_{BC}f_{CF}\mathcal{B}_{ABCF}(t) \\
 &+ P_{A,0} && f_{AB}f_{BD}f_{DC}f_{CF}\mathcal{B}_{ABDCF}(t) \\
 &+ P_{A,0} && f_{AB}f_{BD}f_{DF}\mathcal{B}_{ABDF}(t)
 \end{aligned} \tag{G.2}$$

Equation G.2 can be generalised to any nuclide X , given as Equation G.3:

$$P_{A \dots X}(t) = P_{A,0} \sum_{\substack{\text{all linearised} \\ \text{decay chains}}} \left(\prod_{\substack{\text{linearised} \\ \text{decay chain } j}} \left(\begin{array}{c} \text{All steps in } j, \\ \text{ending with} \\ Z=X \\ \prod_{\substack{(Y,Z), \\ \text{starting with} \\ Y=A}} (f_{YZ}) \end{array} \right) \mathcal{B}_j(t) \right). \tag{G.3}$$

Appendix G. Decomposing branching decay chains

The area under the curve $P_{A\dots X} = \frac{P_{A,0}}{\lambda_X}$. To normalize this such that the area under the curve is independent of the initial population of A ,

$$\frac{P_{A\dots X}(t)}{P_{A,0}} = \sum_{\substack{\text{all linearised} \\ \text{decay chain } j}} \left(\prod_{\substack{\text{All steps in } j, \\ \text{ending with} \\ Z=X \\ (Y,Z), \\ \text{starting with} \\ Y=A}} (f_{YZ}) \mathcal{B}_j(t) \right). \quad (\text{G.4})$$

The expression in the bracket is only dependent on all the decay constants in the linearised decay chain $(\lambda_A, \lambda_B, \dots, \lambda_X)$.

Appendix H

Overlapping decay chains

Let there be two different neutron-induced reactions in the same foil, reaction y and reaction z , producing root nuclides A^* and A . Let's say reaction y produces the excited state nuclide A^* , which then decays as follows:

$$y: [\text{reactant}] \xrightarrow{n} A^* \xrightarrow{\lambda_{A^*}} A \xrightarrow{\lambda_A} B, \quad (\text{H.1})$$

while reaction z produces the ground state nuclide A , so it has a shorter decay chain:

$$z: [\text{reactant}] \xrightarrow{n} A \xrightarrow{\lambda_A} B. \quad (\text{H.2})$$

Let's say that the internal conversion from $A^* \xrightarrow{IT} A$ only emits gamma-rays and X-rays of < 20 keV, outside of the detector's calibrated range of neutron energies; the only detectable gamma-ray line for both reactions is the positron emission: $A \xrightarrow{\beta^+} B$ which emits two 511 keV positron annihilation gamma-rays upon decaying.

The number of (time-integrated) counts measured at the 511 peak in the gamma-ray spectrum for this foil, labelled the l^{th} peak below, is given as a scalar value:

$$\mathcal{N}_l = \epsilon_{\text{abs}}(511 \text{ keV}) I_{(A, 511 \text{ keV})} \left(N_y \int_b^c \lambda_A (\Pi_a dt * \mathcal{B}_{A^*A}) (t) dt + N_z \int_b^c \lambda_A (\Pi_a * \mathcal{B}_A) (t) dt \right) \quad (\text{H.3})$$

If the experiment proceeds by acquiring a time-unresolved gamma-ray spectrum, then the contribution of y and z are added together, such that N_y cannot be separated from N_z , removing one row from the response matrix $\underline{\mathbf{R}}$ and reducing its rank by one.




Figure H.1: Using two hypothetical reaction rates as the decay constants λ_{A^*} and λ_A , an example of a decay rate curve that can be separated into two components is shown. The decay rates are only plotted at time between $b < t \leq c$.

However, the count **rate** tells a richer story, allowing us to determine whether each incoming 511 keV gamma-ray originated from reaction y or reaction z :

$$\begin{aligned} \left. \frac{d\mathcal{N}_l}{dt} \right|_{b \leq t \leq c} &= \epsilon_{\text{abs}}(511 \text{ keV}) I_{(A, 511 \text{ keV})} \left(\lambda_A (\Pi_a * \mathcal{B}_{A^*A}) (t) N_y \right) \\ &+ \epsilon_{\text{abs}}(511 \text{ keV}) I_{(A, 511 \text{ keV})} \left(\lambda_A (\Pi_a * \mathcal{B}_A) (t) N_z \right) \end{aligned} \quad (\text{H.4})$$

ϵ_{abs} , I_l , and the decay constants λ are all known constants extracted from nuclear data libraries. The convolved Bateman equations (constructed from exponentiations of various decay constants λ convolved with the top hat function given by Equation 5.1.1) can be evaluated at any time between $t = b$ and $t = c$.

Figure H.1 gives a graphical intuition: the total decay rate (**black** line) can be decomposed into the decay rates of A atoms indirectly produced by reaction y (**blue** line) and the decay rates of A atoms directly produced by reaction z (**orange** line). Note that when $\frac{\lambda_{A^*}}{\lambda_A} < \varepsilon$ where ε is a small positive number, the difference between these two curves becomes negligible and therefore the decay rate cannot be separated into its constituent components.

If the data was acquired in a time-resolved manner, then it can be split into multiple chunks. Continuing with the above example, let's say the data were acquired with enough time-resolution

to allow it to be split into the following three time periods:

$$\mathcal{N}_{511 \text{ keV, meas}} = \begin{pmatrix} \text{counts at the 511 keV between } t = b - t_1 \\ \text{counts at the 511 keV between } t = t_1 - t_2 \\ \text{counts at the 511 keV between } t = t_2 - c \end{pmatrix} \quad (\text{H.5})$$

where $b < t_1 < t_2 < c$.

Similarly, the expected number of counts in these periods expressed with gross number of reaction of y and z is:

$$\mathcal{N}_{511 \text{ keV, theo}} = \begin{pmatrix} \int_b^{t_1} \left(\frac{d\mathcal{N}_l}{dt} \right) dt \\ \int_{t_1}^{t_2} \left(\frac{d\mathcal{N}_l}{dt} \right) dt \\ \int_{t_2}^c \left(\frac{d\mathcal{N}_l}{dt} \right) dt \end{pmatrix} \quad (\text{H.6})$$

where “theo” stands for theoretical.

The two unknowns N_y and N_z in Equation H.4 can be calculated by performing a least-squares fit of $\mathcal{N}_{511 \text{ keV}}$ onto $\mathcal{N}_{511 \text{ keV, meas}}$. The equation used to perform this least-squares fit is exactly the same as Equation 4.46, only difference being that the variables are change to $\underline{\mathbf{S}}_{\mathcal{N}_{\text{meas}}} \rightarrow \underline{\mathbf{S}}_{\mathcal{N}_{511 \text{ keV, meas}}}$, $\mathbf{N}_{\text{meas}} \rightarrow \mathcal{N}_{511 \text{ keV, meas}}$, and $\underline{\mathbf{R}} \rightarrow$ the appropriate matrix by integrating Equation H.4 to the limits stated in Equation H.6. It no coincidence that the unfolding equation of in the Chapter 4 concerning neutron spectrum unfolding reappears this context of gamma-ray spectrum unfolding, as both are unfolding problems which are solved by the algorithms proposed in Chapter 4 equally well.

Equation H.4 can be generalised to represent the count-rate at the l^{th} line found in the foil’s gamma-ray spectrum spectrum:

$$\left. \frac{d\mathcal{N}_l}{dt} \right|_t = \sum_{\substack{\text{All isotopes } X \\ \text{that emits } \gamma\text{-rays} \\ \text{at energy } E_l}} \epsilon_{\text{abs}}(E_l) I_{(X, E_l)} \sum_{\substack{\text{All root nuclides} \\ \text{whose decay chain} \\ \text{includes } X}} (\lambda_X \overline{P_{k \dots X}}(t)) \quad \{\text{where } b \leq t \leq c\}. \quad (\text{H.7})$$

The function $P_{k \dots X}(t)$ is defined in Equation G.3, where $PA, 0 \equiv N_k$.

Then, for any arbitrarily complex decay chain of isotopes, with the appropriate choice of cutoff times of the measurement t_k where $1 \leq k \leq m$, the $\mathcal{N}_{l, \text{theo}}$ can be constructed for each gamma-ray line l , and by fitting to the corresponding $\underline{\mathbf{R}}_{l, \text{meas}}$, the gamma-ray spectrum $\mathcal{N}$ can be unfolded to the response vector (gross number of radionuclides produced), achieving the inverse of Equation 5.20.

H.1 Future direction of investigation

Further research would be required to find the best values of the cutoff times of the measurement t_k ($1 \leq k < m$), such that the matrix built by integrating Equation H.7 to these cutoff times can be used to optimally unfold the gamma-ray spectrum $\mathcal{N}$ to achieve a response vector $\mathbf{N}_{\text{meas}}$ with the smallest error $\underline{\underline{\text{cov}}}(\mathbf{N}_{\text{meas}})$.

Appendix I

Calculating $\Sigma_{k,\max}$

In order to calculate the maximum macroscopic cross section $\Sigma_{k,\max}$, one can apply the $\max_{E_n}$ operator onto the definition of macroscopic cross section (Equation 5.3):

$$\Sigma_{k,\max} = \max_{E_n}(\Sigma_k(E_n)) = \max_{E_n} \left(\sum_u (N_{d,u} \times \sigma_{k,u}(E_n)) \right). \quad (\text{I.1})$$

which sums over all of the relevant reactants u .

Self-shielding issues are typically caused by resonance peaks that are 2 magnitudes or more above than the average microscopic-cross section values.

Consider a material consisting of two hypothetical reactant isotopes ψ and v , both of them produce root nuclide k . Assume none of the resonance peak of $\sigma_{k,\psi}(E_n)$ overlap with the resonance peaks of $\sigma_{k,v}(E_n)$, which is a valid assumptions as resonance peaks tend to be quite narrow. Then at the tip of every resonance peak of ψ , $\Sigma_k(E_n) = N_{d,\psi}\sigma_{k,\psi}(E_n) + N_{d,\psi}\sigma_{k,v}(E_n) \approx N_{d,\psi}\sigma_{k,\psi}(E_n)$, and vice versa. Therefore, we can simplify Equation I.1 as:

$$\Sigma_{k,\max} = \max_u \left(N_{d,u} \times \max_{E_n} (\sigma_{k,u}(E_n)) \right). \quad (\text{I.2})$$

This means only a handful of scalar values $\max_{E_n} (\sigma_k(E_n))$ needs to be stored after reading the nuclear data library at Section 5.6.2, instead of storing the full array of cross section data (before rebinning) for each reactant u , reducing memory usage.

This definition of $\Sigma_{k,\max}$ can then be substituted into criterion 5.29 in Section 5.6.3 to calculate the maximum allowable thickness.

I.1 Future direction of investigation

The current method of calculating the allowable thickness may be too stringent as many of the resonance peaks have a sharp tip whose height hasn't been precisely measured, and due to its narrowness its height doesn't actually meaningfully affect the overall probability of reaction of neutrons in that group. A better way to calculate the neutron spectrum in the future is to simply find the maximum macroscopic cross section after it has been rebinned into the new group structure,

$$\Sigma_{k,\max} = \max(\Sigma_{ki}) \forall 1 \leq i \leq n, \quad (\text{I.3})$$

so that

$$T \leq \frac{-\log(1 - P_{\text{threshold}})}{\Sigma_{k,\max}} \quad (\text{I.4})$$

$$\leq \frac{-\log(1 - P_{\text{threshold}})}{\Sigma_{ki}} \forall 1 \leq i \leq n. \quad (\text{I.5})$$

This method has the added benefit that no extra variable $\max_{E_n} \sigma_k(E_n)$ needs to be extracted and stored between steps, as the Σ_{ki} is already part of the data transferred from Section 5.6.2 to Section 5.6.3.

However, further investigation is required to examine the impact of this simplification: it remains to be seen whether the underestimation of self-shielding by Equation I.3 would be significant enough to impact the final gross production of nuclides or not.

Appendix J

Supplementary plots

Due to the numerous data points, the plots of the response vector $\boldsymbol{\eta}$ for the experiment are shown here in the appendix rather than the main text. Section 6.2 includes more details about each foil, including why foil 09 and 18 are excluded from consideration. The expected number of decays (blue dots ●) is obtained by first folding the *a priori* spectrum ϕ_{sol} with the response matrix $\underline{\mathbf{R}}$ to obtain $N_0 = \underline{\mathbf{R}}\phi_0$, then converted to $\boldsymbol{\eta}$ according to Equation 6.6. The unfolded $\boldsymbol{\eta}$ (orange dots ●) are obtained in a similar fashion, by converting $\mathbf{N}_{\text{sol}}$ to $\boldsymbol{\eta}$ according to Equation 6.6. In general, the unfolding algorithm is able to adjust the neutron spectrum such the unfolded $\boldsymbol{\eta}$ moves from the *a priori-generated* $\boldsymbol{\eta}$ towards the measured $\boldsymbol{\eta}$.

Despite the apparent skewness of the distribution of deviations of measured $\boldsymbol{\eta}$ from *a priori-generated* $\boldsymbol{\eta}$, the *a priori* has indeed been correctly normalized, as verified graphically and numerically.




Figure J.1: Number of decays of radioisotopes in foil 01-Cu during the measurement period.




Figure J.2: Number of decays of radioisotopes in foil 02-Cu-contact during the measurement period.




Figure J.3: Number of decays of radioisotopes in foil 03-In during the measurement period.




Figure J.4: Number of decays of radioisotopes in foil 04-Sn during the measurement period.




Figure J.5: Number of decays of radioisotopes in foil 05-Ge during the measurement period.




Figure J.6: Number of decays of radioisotopes in foil 06-Zn during the measurement period.




Figure J.7: Number of decays of radioisotopes in foil 07-Pt during the measurement period.




Figure J.8: Number of decays of radioisotopes in foil 08-Hf during the measurement period.




Figure J.9: Number of decays of radioisotopes in foil 10-Yb during the measurement period.




Figure J.10: Number of decays of radioisotopes in foil 11-Mo during the measurement period.




Figure J.11: Number of decays of radioisotopes in foil 12-Cd during the measurement period.




Figure J.12: Number of decays of radioisotopes in foil 13-Er during the measurement period.




Figure J.13: Number of decays of radioisotopes in foil 14-Nd during the measurement period.




Figure J.14: Number of decays of radioisotopes in foil 15-Bi during the measurement period.




Figure J.15: Number of decays of radioisotopes in foil 15-Bi-later during the measurement period.




Figure J.16: Number of decays of radioisotopes in foil 16-Au during the measurement period.




Figure J.17: Number of decays of radioisotopes in foil 17-Er during the measurement period.




Figure J.18: Number of decays of radioisotopes in foil 19-W-decayed during the measurement period.

Appendix K

Notations

N_k = The gross number of root radionuclides produced by the k^{th} reaction, where $1 \leq k \leq m$.

ϕ_i = The (discretized) neutron flux ($\text{cm}^{-2}\text{s}^{-1}$) in the i^{th} bin, where $1 \leq i \leq n$. Obtained by integrating the un-discretized scalar flux over the energy group, see Equation 3.1.

R_{ki} = The gross number of root radionuclides k produced per unit flux of neutron in the i^{th} bin.

n = number of energy groups, i.e. length of the group structure. See Equation 3.1 for more details.

m = number of types of reactions that produces measurable radionuclides and are used for neutron spectrum unfolding.

p = dimensionality of a manifold P , expected to be a number smaller than m .

P = a set of all of physically realistic fusion neutron spectra. This forms a manifold of lower dimensionality than m . See Appendix L.4 for “**manifold**” (mathematics).

$\text{cov}(x, y)$ = covariance between two scalar random variables. If $x = y$, then $\text{cov}(x, y)$ = variance of x .

$\underline{\underline{\text{cov}}}(\mathbf{x})$ = covariance matrix associated with $\mathbf{x}$. See Equation 3.12 for more details.

$\underline{\underline{\mathbf{S}}}_v$ = Inverse of the covariance matrix of vector variable $\mathbf{v}$. See Equation 3.15.

$\phi_{\text{true}}, \phi_0, \phi_{\text{sol}}$ = the true neutron spectrum, the *a priori*, and the solution spectrum. The differences between them are detailed in Table 3.1 and Section 4.2.1.

$\mathbf{N}_{\text{meas}}, \mathbf{N}_{\text{sol}}$ = the measured response vector, and the would-have-been response vector if ϕ_{sol} was indeed the neutron spectrum with no noise. See the difference between them at Table 3.1.

ε = the noise (accumulated during the measurement of $\mathbf{N}_{\text{meas}}$) that causes the measured response vector $\mathbf{N}_{\text{meas}}$ to deviate from $\underline{\mathbf{R}}\phi_{\text{true}}$ (Equation 3.7).

E = the energy of a neutron.

E_l = the energy of a γ -ray.

K.1 Number sets

$\mathbb{N}$ = The set of all natural numbers, i.e. all positive integers (which excludes zero).

$\mathbb{R}$ = The set of all real numbers.

$\mathbb{R}^+$ = The set of all positive real numbers (which excludes zero).

$\mathbb{R}^n$ = The set of all vectors made of n real numbers. This forms the n -dimensional Cartesian coordinates, where any point in this space represents a vector of length $= n$.

$v \in \mathbf{S}$ means that variable v is an element in the set $\mathbf{S}$.

$\forall$ = for all.

K.2 Generic information about the activation foil

σ_{ki} = Microscopic cross section for reaction k due to an incident neutron in energy bin i . See Equation 3.4. The unit usually used is barn = 10^{-24}cm^2 , and reflects the reaction rate upon being bombarded by a unit flux of neutrons. Note the notation of microscopic cross section σ clashes with the notation for standard deviation σ . Therefore in this thesis, unless otherwise specified, σ refers to microscopic cross section by default.

A = Area of the foil.

T = Thickness of the foil.

N_d = Number density of a reactant isotope.

K.4 Information Theory

Σ_{ki} = Macroscopic cross section. In this thesis it is differentiated from the summation sign $\sum_{i=1}^n$ by the position of the subscript, where macroscopic cross section has the indices on the bottom right instead of a dummy variable directly below it.

$\mathcal{N}$ = the effective response vector, where $\mathcal{N}_l$ is the number of counts at the l^{th} gamma-ray photopeak.

$\underline{\underline{\mathcal{R}}}$ = the effective response matrix, where $\underline{\underline{\mathcal{R}}}_{li}$ = the number of counts from the l^{th} gamma-ray photopeak for every $\text{cm}^{-2}\text{s}^{-1}$ of neutron flux in the i^{th} energy bin incident on the foil. See Equation 5.23.

η = the number of decays of all radionuclides between time $t = b - c$. This was fed to the neutron spectrum unfolding algorithm in Section 6.2.2, in lieu of $\mathbf{N}_{\text{meas}}$.

$\underline{\underline{\rho}}$ = the response matrix that converts the neutron spectrum into the number of decays of all radionuclides between time $t = b - c$. This was fed to the neutron spectrum unfolding algorithm in Section 6.2.2, in lieu of $\underline{\underline{\mathbf{R}}}$.

K.3 un-discretized fluxes

$\Phi(\mathbf{r}, E, \Omega; t)$ = Angular flux. It represents the number of neutrons passing through a unit area at the specified position ($\mathbf{r}$), energy (E) and angle (Ω), and is a time (t) dependent quantity.

$\phi(\mathbf{r}, E; t) = \iint_{4\pi} \Phi(\mathbf{r}, E, \Omega; t) d\Omega$ = scalar flux. It represents the number of neutrons of a certain energy (E) passing through a unit area at the specified position ($\mathbf{r}$). Same as above, it is a time-dependent (t) quantity.

$\phi(\mathbf{r}; t) = \int_0^\infty \phi(\mathbf{r}, E; t) dE$ = standard flux. It represents the total number of neutrons passing through a unit area at the specified position ($\mathbf{r}$). Same as above, it is a time-dependent (t) quantity.

K.4 Information Theory

$H(\mathbf{p})$ = (self-)entropy of one normalized distribution. See Equation 4.13.

$H(\mathbf{p}, \mathbf{q})$ = cross-entropy of two normalized distributions. See Equation 4.14.

$D_{KL}(\mathbf{p}, \mathbf{q})$ = Kullback–Leibler Divergence of two normalized distributions. See Equation 4.15.

ϕ' = a normalized distribution of ϕ , such that $\phi'_i = \frac{\phi_i}{\sum_j \phi_j}$.

K.5 Linear Algebra operators

$\underline{\underline{\mathbf{I}}}$ = the identity matrix.

$\mathbf{x} \cdot \mathbf{y}$ = dot product of vector $\mathbf{x}$ with vector $\mathbf{y}$.

$\mathbf{x} \cdot \underline{\underline{\mathbf{M}}}$ = matrix $\underline{\underline{\mathbf{M}}}$ left-multiplied by vector $\mathbf{x}$. This is the same operation as $\mathbf{x}^T \underline{\underline{\mathbf{M}}}$. If matrix is of the shape $m \times n$, then vector $\mathbf{x}$ is expected to be an m -dimensional vector, and the output of this operation is expected to be an n -dimensional vector.

δ_{ij} = Kronecker delta, so that δ_{ij} only = 1 if $i = j$, otherwise $\delta_{ij} = 0$.

$\underline{\underline{\text{Diag}}}(\mathbf{x})$ = Diagonal matrix built from vector $\mathbf{x}$, where the resulting matrix $\underline{\underline{\mathbf{M}}}$ is given by $M_{ij} = \delta_{ij}x_i$.

$\underline{\underline{\text{Diag}}}(\underline{\underline{\mathbf{M}}})$ = vector of built from the main diagonal of the square matrix $\underline{\underline{\mathbf{M}}}$, where the resulting vector $\mathbf{x}$ is given by $x_i = M_{ii}$.

$\mathbf{x} \circ \mathbf{y}$ = Hadamard product of vector $\mathbf{x}$ with vector $\mathbf{y} = \underline{\underline{\text{Diag}}}(\mathbf{x})\mathbf{y}$.

$\mathbf{x} \circ \underline{\underline{\mathbf{Y}}}$ = Hadamard product of vector $\mathbf{x}$ with matrix $\underline{\underline{\mathbf{Y}}}$, so that the resulting matrix $\underline{\underline{\mathbf{M}}}$ is given by $M_{ij} = x_i Y_{ij}$.

$\underline{\underline{\mathbf{X}}} \circ \mathbf{y}$ = Hadamard product of matrix $\underline{\underline{\mathbf{X}}}$ with vector $\mathbf{y}$, so that the resulting matrix $\underline{\underline{\mathbf{M}}}$ is given by $M_{ij} = X_{ij} y_j$.

$\mathbf{x} \otimes \mathbf{y}$ = outerproduct of vector $\mathbf{x}$ with vector $\mathbf{y}$, so that the resulting matrix $\underline{\underline{\mathbf{M}}}$ is given by $M_{ij} = x_i y_j$.

$\frac{1}{\mathbf{x}}$ = element-wise inverse of $\mathbf{x}$.

$\exp(\mathbf{x})$ = element-wise exponentiation of vector $\mathbf{x}$.

$\log(\mathbf{x})$ = element-wise logarithm of vector $\mathbf{x}$.

$\sqrt{\mathbf{x}}$ = element-wise square root of vector $\mathbf{x}$.

$\mathbf{1}^m$ = A vector of m -dimensions, made of m 1's.

$\|\mathbf{x}\|_2$ = the L_2 norm of the vector $\mathbf{x} = \sqrt{\sum_i (x_i^2)}$.

K.6 Transmutation equations

f_{AB} = Branching ratio of radionuclide A to nuclide B , i.e. what fraction of radionuclide A decays to nuclide B .

$\mathcal{B}_{XYZ}(t)$ = Bateman decay equation explaining the number of atoms Z , where Z is situated at the end of the decay chain $X \rightarrow Y \rightarrow Z$.

$P_X(t)$ = population of nuclide X , which is time-dependent.

$\overline{P_X}(t) = (\Pi_a * P_X)(t)$.

$\Pi_a(t)$ = the top-hat function that is equal to $\frac{1}{a}$ between $t = 0$ and $t = a$, and 0 everywhere else.

K.7 Gamma-ray spectroscopy

$\mathcal{R}$ = the effective response matrix with L rows and n columns, converting the neutron spectrum into gamma-ray peaks.

L = number of usable distinct gamma-ray peaks on all of the gamma-ray spectra of a set of activation foils.

Θ = the total number of candidate foils supplied to the foil selector framework.

ϑ = the number of foils that the user input requires the foil selector framework to output.

$$1 \leq \vartheta < \Theta$$

Appendix L

Glossary

“**neutron spectrum unfolding with activation foil/activation foil (neutron spectrum) unfolding**”: The very technique that is discussed in this paper. While the former (“neutron spectrum unfolding with activation foil”) describes the technique fully and unambiguously, it is very verbose. Therefore it is sometimes abbreviated to “activation foil neutron spectrum unfolding”, or even more briefly, “activation foil unfolding”.

L.1 Nuclear physics

“**nucleus**”: The core of a single atom, containing ≥ 1 protons and ≥ 0 neutrons.

“**nuclide**”: The nucleus of a species of atom.

“**radionuclide**”: Nuclide that is radioactive/ The nucleus of a species of atoms with unstable nuclei, whose decay typically emit detectable radiation.

“**isotope**”: An isotope of element A refers to nuclei of A that contains a specific number of neutrons. This term is sometimes used to emphasize the difference between different species of atoms that make up the same chemical element.

“**radioisotope**”: Isotope that is radioactive.

nuclear “**isomer**”: An isomer of isotope ${}^x A$ refers to nuclei of ${}^x A$ occupying a specific excited state.

“**neutron spectrum**”: The number of neutrons in each energy bin that passes through a unit area per unit time, typically $\text{cm}^{-2}\text{s}^{-1}$.

“**neutron fluence**”: the number of neutrons that passes through a unit area, accumulated over a period of time, typically cm^{-2} . Integrating the neutron spectrum over a range of time yields the neutron fluence.

“**angular flux**”: The number of neutrons of a particular energy passing through a unit area per unit time at a specified direction (or solid angle, hence the word “angular”) and energy. See the entry for Φ in Section K.3. Not used in this thesis, only included here for completeness.

“**scalar flux**”: The number of neutrons of a certain energy passing through a unit area (regardless of direction).

“**standard flux**”: The number of neutron passing through a unit area per unit time (regardless of direction or energy).

“group structure”

“Thermal neutrons”:

“slow neutrons”: 1-300 keV neutrons

“fast neutrons”: 0.3 to 20 MeV

“high-energy neutrons”: 20 MeV+

L.1.1 Nuclear data

“**nuclear data library**”: A file or a collection of files containing data about nuclides. Only two types of data concerns the scope of this thesis: microscopic cross section of reaction with incident neutron, and decay data. These data can be found downloaded from [144].

“**ENDF**”: Evaluated Nuclear Data File, a standardized format of storing nuclear data as a plain ascii text file. [142]

“**MT**” number: An identifier number denoting the type of nuclear reaction that may occur in a material recorded in an ENDF file. The full list of MT numbers and what reaction each represents can be found in the cited pdf [145].

L.2 Fusion-specific terminology

“**tokamak**”: the toroidally-shaped vacuum chamber used to contain the plasma used in magnetic confinement fusion.

“**first wall**”: the interior walls of the tokamak that faces the plasma. This is the first component that all particles escaping the plasma hits.

“**pulse height spectrum**”: a histogram of pulse heights recorded by the detector acquisition software. When a pulse is measured by the detection system, it is sorted into one of the bins according to its height, so that each bin in the histogram stores the number of pulses that fall into its range of pulse heights.

“**deuterium**”: atom of the 2nd heavy ${}^2_1\text{H}$.

“**deuteron**”: nucleus of deuterium.

“**manifold**” (fluid mechanics): A section of piping.

“**tritium**”: atom of ${}^3_1\text{H}$.

“**triton**”: nucleus of tritium.

“**TBR**”: Tritium breeding ratio, $\text{TBR} = \frac{\text{tritium produced by fusion reactor}}{\text{tritium consumed by fusion reactor}}$.

L.3 Experimental physics terminology

“**deconvolution**”: Same definition as unfolding

The difference between the usage of the two is mainly semantic in nature: deconvolution seem to be semantically closer to the concept of “sharpening a blurred image”, and its generalization to other number of dimensions. In 1-D its response matrix would resemble a Toeplitz matrix, each row being generated from a Gaussian distribution with a shifted centroid position compared to the previous row.

However, if the response matrix is no longer Toeplitz-like (i.e. each row is not related to the previous), or when Radon transform is involved, then the literature inclines to name it “unfolding” instead.

There is no literature that rigorously made a distinction between the two, and some unfolding programs are equivocally called “deconvolution” programs (e.g. MAXED [56]).

“**Toeplitz matrix**”: A square matrix where each row is a rotated copy of the previous, i.e. its 2nd row = 1st row shifted to the right by 1 cell, 3rd row = 2nd row shifted to the right by 1 cell, etc.

“**FWHM**”: Full-width of a peak at its half-maximum, where the peak (on a histogram) is assumed to be formed by the detection of a collection of monoenergetic particles.

L.4 Mathematical terminology

“**ill-posed**”: A problem that doesn’t have a unique solution, and/or is ill-conditioned. See Section 3.1.3 or reference [81] for more details.

“**ill-conditioned**”: A problem where errors in the input space are easily magnified and projected onto the output space. All ill-conditioned problems are ill-posed. See Section 3.1.3 or reference [81] for more details.

“**uncertainty**”: A numerical quantity, used to quantify the difficulty of asserting how close a true solution is the to the unfolded solution. To capture the lowest-order uncertainty of a vector, a covariance matrix is typically used.

“**uncertainty propagation**”: the process of inferring the uncertainty associated with the output (covariance matrix associated with ϕ_{sol}) given the uncertainty associated the inputs (covariance matrix associated with $\mathbf{N}_{\text{meas}}$ and covariance tensor associated with $\underline{\mathbf{R}}$) to an unfolding algorithm.

“**degree of underdetermination**”: $n - m$. See Appendix K for definition of m and n .

“**degree of determination**”: $m - n$. See Appendix K for definition of m and n .

“**condition number**”: The condition number of a matrix $\underline{\mathbf{R}}$ is given by $\frac{\text{largest singular value of } \underline{\mathbf{R}}}{\text{smallest singular value of } \underline{\mathbf{R}}}$.

“**curse of dimensionality**”: The parameter space’s hypervolume increases exponentially w.r.t. number of parameters p , such that the number of solutions that have to be checked becomes very large at when the p is high. This term thus refers to the intractability issue (i.e. the parameter space that requires checking becomes unfeasibility large) that arises as p increases.

“**manifold**” (mathematics): an object with $< n$ dimensions, while embedded in n dimensions. See Section 3.3.1.

“**hypervolume**”: the generalised of concept of “volume” applied to higher dimensions. The hypervolume spanned by a 1-D object is known as the length; the hypervolume spanned by a 2-D object is known as a line, the hypervolume spanned by a 3-D object is known as the volume, etc.

L.4.1 Mathematical programming terminology

“**loss function**”/“cost function”: a function with a vector-valued domain and scalar-valued codomain. In the context of unfolding, the vector domain is ϕ_{sol} and the scalar codomain is a loss value to be minimised by the unfolding algorithm.

“**loss (value)**”/“cost”: value outputted by the loss function when given a certain point in the solution space (ϕ_{sol}) as the input.

“**MCMC**”: Markov Chain Monte Carlo, a type of algorithm used to sample distributions.

“**Tikhonov/Tykhonov regularization**”: A technique of adding extra conditions into an optimization to ensure that its resulting solution is physically sensible, and to prevent overfitting. See Equation 4.5.

“**Hessian matrix**”: A square matrix of 2nd order derivatives of a scalar-valued function with respect to its vector-valued domain. For example, if $\exists f(\mathbf{x}) \in \mathbb{R} \forall \{\mathbf{x} : \mathbf{x} \in \mathbb{R}^2\}$, then
$$\text{Hess}_{\mathbf{x}}(f(\mathbf{x})) = \begin{pmatrix} \frac{\partial^2 f}{\partial x_1 \partial x_1} & \frac{\partial^2 f}{\partial x_1 \partial x_2} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2 \partial x_2} \end{pmatrix}.$$

“**line search**”: Finding the point of lowest loss value along some straight line that cuts through the loss function landscape.

L.5 Unfolding algorithms

“**MAXED**”: Maximum Entropy Deconvolution, an unfolding algorithm.

“**SAND-II**”: Spectrum Analysis by Neutron Detectors [119].

“**GRAVEL**”: a purely iterative unfolding code that is a successor to SAND-II

“**IMAXED**”: Improved MAXED, where a more efficient optimum-searching algorithm is implemented.

“**AMAXED**”: Alternative definition MAXED, where the definition of cross entropy is reversed compared to MAXED.

“**AMAXED-Regularization**”: Same as AMAXED, but the solution is obtained by adjusting the regularisation parameter τ rather than imposing the user-supplied desired- χ^2 as the constraint.

“**pseudo-inverse**”:

“**direct-inversion**” [87]

L.6 Metrics

“**Fisher information**”: A quantity that measures the average 2nd derivative of a continuous distribution. See Equation 4.10.

“**Kullback–Leibler Divergence**”: Difference between two normalized distributions, defined in Equation 4.15, it is always nonnegative. Note that this value is not symmetric, i.e. in general $D_{KL}(\mathbf{p}, \mathbf{q}) \neq D_{KL}(\mathbf{q}, \mathbf{p})$. This quantity is also sometimes known as “relative entropy” in the literature.

“**Cross Entropy**”: Sum of Kullback–Leibler Divergence and self-entropy, defined in Equation 4.14, it is always nonnegative. Note the asymmetry as well, i.e. in general $H(\mathbf{p}, \mathbf{q}) \neq H(\mathbf{q}, \mathbf{p})$.

“**(Self-)entropy**” (information entropy): A quantity that measures the information content of a normalized distribution.

- It is always non-negative, and is defined in Equation 4.13.
- This is used *purely* in the context of information theory, and has no direct links to thermodynamics applications. The entropy of a distribution is often quoted as the “average surprisal” [118] when a sample is drawn from a distribution.
- To allow ourselves to translate the mathematical jargon into physicists’ language at the cost of some loss of generality, the word “sample” can be replaced with “neutron” and “distribution” can be replaced with “neutron field”, i.e.
 - “On average, how surprised should one be when they see a neutron’s energy reported by a detector immersed in a (known) neutron field.”
 - If one compares this definition to Equation 4.13, it is obvious that “surprisal” is measured by log of probability. Therefore the average surprisal is simply the weighted mean of the log-of-probability.

L.7 Computational terminology

“**GUI**” (Graphical User Interface) program, where the user may use their cursor to interact with the program. This may be more intuitive for less experienced users who have only used Windows/ macOS/ Android/ iOS.

“**CLI**” (Command Line Interface) program, where the user can only communicate with the program by typing into a text-only terminal interface. This usually requires some understanding of the terminal interface typically used on Unix systems.

L.8 Terminology typically used in unfolding experiments

“**ADC**”: Analogue-digital converter

“**Genie 2000**”: A software written by Mirion Technologies to perform gamma-ray spectroscopy and analyse its results. It has a GUI and no CLI, and is only compatible with Windows.

“**FISPACT-II**”: A multi-physics nuclear inventory code that simulates the population of radionuclides during transmutation (i.e. irradiation *and* decay) using an ODE-solver.

“**rabbit/NAS**”: a pneumatic sample transfer system that are used by laboratories worldwide to perform human-contactless sample transport, especially useful for radioactive samples. A capsule containing the sample fits loosely inside a long tube that spans the origin and destination, and is propelled towards the destination end by low pressure jet of air. Some facilities refer to their system as rabbit [139, 140] while others refers to them as Neutron Activation System (NAS) [46, 64, 67, 77]

“**BSS**”: Bonner sphere system, a neutron spectrometry technique.

“**MPR**”: Magnetic proton recoil, a neutron spectrometry technique.

“**ToF**”: Time-of-Flight, a neutron spectrometry technique.

“**PMT**”: Photomultiplier tube.

L.9 Terminology specific to this thesis

“**schedule**”: This refers specifically to 3 positive real numbers: irradiation duration, transport duration, and acquisition duration. See Figure 5.1 for the definition of these three durations.

“**saturation threshold**” count rate of gamma-ray detector: the count rate above which the detector fails to accurately count the number of incident gamma-ray, or log their energies accurately on the gamma-ray spectrum (due to baseline-shift of the detector signal).

“**foil selector framework**”: A collection of abstract ideas of how foil selection should be done, in theory.

“**foil selector program**”: The Python program implementing the ideas in Section 5.6.

“**candidate foils**”: A list of possible foils that the foil selector program can choose from.

“**root (radio)nuclide**”: The radioactive nuclide produced by neutron-induced reaction in the foil. See (2) of Figure 5.4 “Root” is used in this term to reflect that it is the root of the decay chain, when drawn as a tree diagram.

References

- [1] UNFCCC. *The Paris Agreement*. United Nations Framework Convention on Climate Change (UNFCCC), 2017.
- [2] Jeffrey P Freidberg. *Plasma physics and fusion energy*. Cambridge university press, 2008.
- [3] ITER R.A. Question of the week: Will fusion run out of fuel? 2018. online article:<https://www.iter.org/newsline/-/3071>.
- [4] Carley Willis and Joanne Liou. Safety in fusion. *IAEA Bulletin*, 62(2), 5 2021.
- [5] Mark R Gilbert, T Eade, T Rey, R Vale, C Bachmann, U Fischer, and NP Taylor. Waste implications from minor impurities in European DEMO materials. *Nuclear Fusion*, 59(7):076015, 2019.
- [6] Thomas J Dolan, Ralph W Moir, Wallace Manheimer, Lee C Cadwallader, and Martin J Neumann. *Safety and Environment*, volume 19 of *Lecture Notes in Energy*, page 632–633. Springer London, 2013.
- [7] ITER website: Frequently Asked Questions. https://www.iter.org/faq#When_will_ITER_be_operational. Accessed: Sep-2020.
- [8] M. Siccino, W. Biel, M. Cavedon, E. Fable, G. Federici, F. Janky, H. Lux, F. Maviglia, J. Morris, F. Palermo, O. Sauter, F. Subba, and H. Zohm. DEMO physics challenges beyond ITER. *Fusion Engineering and Design*, 156:111603, 2020.
- [9] Yican Wu. *Fusion neutronics*. Springer Singapore, 2017.
- [10] H-S Bosch and GM Hale. Improved formulas for fusion cross-sections and thermal reactivities. *Nuclear fusion*, 32(4):611, 1992.
- [11] Richard J. Pearson, Armando B. Antoniazzi, and William J. Nuttall. Tritium supply and use: a key issue for the development of nuclear fusion energy. *Fusion Engineering and Design*, 136:1140–1148, 2018. Special Issue: Proceedings of the 13th International Symposium on Fusion Nuclear Technology (ISFNT-13).
- [12] UKAEA. EUROfusion Media Library, 2017. Inside the Joint European Torus (JET) at Culham, UK, with a superimposed image of the hot plasma; credit: UKAEA, courtesy of EUROfusion, image shared under the CC BY 4.0. Original url: <https://euro-fusion.org/services/media-library/>.
- [13] H. Brady, J. E. Cook, J. Maddock, M. Murphy, O. Funk, and C. Mould. fusrr: A Rendering Pipeline for Fusion Reactor Designs. GitHub repository. <https://github.com/Fusion-Power-Plant-Framework/fusrr>.
- [14] ITER Physics Expert Group on Diagnostics and ITER Physics Basis Editors. Chapter 7: Measurement of plasma parameters. *Nuclear Fusion*, 39(12):2541–2575, dec 1999.
- [15] Žiga Štancar, Marina Gorelenkova, Sean Conroy, Patrick Sauvan, James Buchanan, Henri Weisen, and Luka Snoj. Multiphysics approach to plasma neutron source modelling at the JET tokamak. *Nuclear Fusion*, 59(9):096020, jul 2019.

- [16] J.-Ch. Sublet, J.W. Eastwood, J.G. Morgan, M.R. Gilbert, M. Fleming, and W. Arter. FISPACT-II: An Advanced Simulation System for Activation, Transmutation and Material Modelling. *Nuclear Data Sheets*, 139:77–137, 2017. Special Issue on Nuclear Reaction Data.
- [17] C. Bachmann, M. Siccino, M. Albino, A. Chiappa, G. Falcitelli, G. Federici, L. Giannini, and C. Luongo. Influence of a high magnetic field to the design of EU DEMO. *Fusion Engineering and Design*, 197:114050, 2023.
- [18] Harald W. Weber. *Irradiation Damage in Superconductors*, pages 853–864. Springer US, Boston, MA, 1986.
- [19] Arunodaya Bhattacharya, Steven J Zinkle, Jean Henry, Samara M Levine, Philip D Edmondson, Mark R Gilbert, Hiroyasu Tanigawa, and Charles E Kessel. Irradiation damage concurrent challenges with RAFM and ODS steels for fusion reactor first-wall/blanket: a review. *Journal of Physics: Energy*, 4(3):034003, jul 2022.
- [20] Bethany Colling, Tim Eade, Malcolm J Joyce, Raul Pampin, Fabien Seyvet, Andrew Turner, and Victor Udintsev. Neutronics analysis for integration of ITER diagnostics port EP10. *Fusion Engineering and Design*, 109:1109–1113, 2016.
- [21] Shaw Martin, Bradley Campbell, James Ballantine, and Kevin Roberts. Development of a High Temperature Neutron Flux Detector. In *EPJ Web of Conferences*, volume 225, page 03003. EDP Sciences, 2020.
- [22] R. Mitteau, R. Eaton, A. Gervash, V. Kuznetsov, V. Davydov, and R. Rulev. Allowable heat load on the edge of the ITER first wall panel beryllium flat tiles. *Nuclear Materials and Energy*, 12:1067 – 1070, 2017. Proceedings of the 22nd International Conference on Plasma Surface Interactions 2016, 22nd PSI.
- [23] P.-H. Rebut. ITER: the first experimental fusion reactor. *Fusion Engineering and Design*, 30(1):85 – 118, 1995.
- [24] Aljaž Čufar, Paola Batistoni, Sean Conroy, Zamir Ghani, Igor Lengar, Sergey Popovichev, Brian Syme, Žiga Štancar, Luka Snoj, and JET Contributors. Calculations to Support In Situ Neutron Yield Calibrations at the Joint European Torus. *Fusion Science and Technology*, 74(4):370–386, 2018.
- [25] R Prokopowicz, M Scholz, A Szydłowski, S Popovichev, and JET EFDA contributors. Measurements of Neutron Yield from Deuterium Plasmas at JET by Activation Techniques. In *AIP Conference Proceedings*, volume 993, pages 251–254. AIP, 2008.
- [26] Brian Syme. 3rd EIROForum School on Instrumentation. https://indico.cern.ch/event/233159/contributions/492771/attachments/388590/540450/EiroForum_2013_NaG_Diagnostics-V1_.pdf, 2013.
- [27] Sean William Conroy. *Diagnosis of fusion products for reactor relevant plasmas*. PhD thesis, Imperial College London, 1990.
- [28] Iwona Klimek. *Modelling and Measurements of MAST Neutron Emission*. PhD thesis, Uppsala University, Department of Physics and Astronomy, 2016.
- [29] M. Cecconello, W. Boeglin, D. Keeling, S. Conroy, I. Klimek, and R.V. Perez and. Discrepancy between estimated and measured fusion product rates on MAST using TRANSP/NUBEAM. *Nuclear Fusion*, 59(1):016006, Nov 2018.
- [30] EFDA-JET. JET Inside Vessel, 11 2007. Image shared under the CC BY 4.0.

- [31] L. W. Packer, P. Batistoni, C. R. Nobs, S. Bradnam, B. Colling, K. Drozdowicz, S. Jednorog, M. R. Gilbert, E. Laszynska, D. Leichtle, J.W.Mietelski, M. Pillon, I.E. Stamatelatos, T. Vasilopoulou, A. Wójcik-Gargula, and JET Contributors. ACT - activation of real ITER materials Report on Deliverables F21-29. internal report, 6 2019.
- [32] V Vylet. Response matrix of an extended Bonner sphere system. *Nuclear Inst. and Methods in Physics Research, A*, 476(1-2):26–30, 2002.
- [33] D.J Thomas and A.V Alevra. Bonner sphere spectrometers—a critical review. *Nuclear Inst. and Methods in Physics Research, A*, 476(1-2):12–20, 2002.
- [34] A. Zimbal, M. Reginatto, H. Schuhmacher, L. Bertalot, B. Esposito, F. Poli, J. M. Adams, S. Popovichev, V. Kiptily, A. Murari, and Contributors to the EFDA-JET work program. Compact NE213 neutron spectrometer with high energy resolution for fusion applications. *Review of Scientific Instruments*, 75(10):3553–3555, 2004.
- [35] Lin-Yue Liu, Xiao Ouyang, Jin-Lu Ruan, Song Bai, and Xiao-Ping Ouyang. Performance Comparison Between SiC and Si Neutron Detectors in Deuterium–Tritium Fusion Neutron Irradiation. *IEEE Transactions on Nuclear Science*, 66(4):737–741, 2019.
- [36] T. Angelescu, A.E. Cheremukhin, V.M. Ghete, N. Ghiordănescu, I.A. Golutvin, S. Lazanu, I. Lazanu, A. Mihul, A. Radu, N.Yu Susova, A. Vasilescu, and N.I. Zamyatin. Radiation hardness studies on silicon detectors in fast neutron fields. *Nuclear instruments & methods in physics research. Section A, Accelerators, spectrometers, detectors and associated equipment*, 357(1):55–63, 1995.
- [37] M. Rebai, D. Rigamonti, S. Cancelli, G. Croci, G. Gorini, E. Perelli Cippo, O. Putignano, M. Tardocchi, C. Altana, M. Angelone, G. Borghi, M. Boscardin, C. Ciampi, G.A.P. Cirrone, A. Fazzi, D. Giove, L. Labate, G. Lanzalone, F. La Via, S. Loreti, A. Muoio, P. Ottanelli, G. Pasquali, M. Pillon, S.M.R. Puglia, A. Santangelo, A. Trifiro, and S. Tudisco. New thick silicon carbide detectors: Response to 14 MeV neutrons and comparison with single-crystal diamonds. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 946:162637, 2019.
- [38] Jan Mlynar, John M. Adams, Luciano Bertalot, and Sean Conroy. First results of Minimum Fisher Regularisation as unfolding method for JET NE213 liquid scintillator neutron spectrometry. *Fusion Engineering and Design*, 74(1):781 – 786, 2005. Proceedings of the 23rd Symposium of Fusion Technology.
- [39] E Andersson Sundén, Henrik Sjöstrand, Sean Conroy, Göran Ericsson, M Gatu Johnson, L Giacomelli, Carl Hellesen, Anders Hjalmarsson, E Ronchi, Matthias Weiszflog, et al. The thin-foil magnetic proton recoil neutron spectrometer MPRu at JET. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 610(3):682–699, 2009.
- [40] Göran Ericsson, L Ballabio, S Conroy, J Frenje, H Henriksson, A Hjalmarsson, J Källne, and M Tardocchi. Neutron emission spectroscopy at JET—results from the magnetic proton recoil spectrometer. *Review of Scientific Instruments*, 72(1):759–766, 2001.
- [41] M Gatu Johnson, L Giacomelli, Anders Hjalmarsson, Jan Källne, Matthias Weiszflog, E Andersson Sundén, Sean Conroy, Göran Ericsson, Carl Hellesen, Emanuele Ronchi, et al. The 2.5-MeV neutron time-of-flight spectrometer TOFOR for experiments at JET. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 591(2):417–430, 2008.
- [42] AV Alevra and DJ Thomas. Neutron spectrometry in mixed fields: multisphere spectrometers. *Radiation Protection Dosimetry*, 107(1-3):33–68, 2003.

- [43] E. Klimenkov, Yu Kashchuk, A. Krasilnikov, AA Oleinikov, V. Sevast'yanov, LA Trykov, and N. Semenova. Radiation Hardness of a Fast Neutron Scintillation Spectrometer with a Stilbene Crystal. *Instruments and Experimental Techniques*, 47:166–167, 03 2004.
- [44] Michal Košťál, Jaroslav Šoltés, Ladislav Viererbl, Zdeněk Matěj, František Cvachovec, Vojtěch Rypar, and Evžen Losa. Measurement of neutron spectra in a silicon filtered neutron beam using stilbene detectors at the LVR-15 research reactor. *Applied Radiation and Isotopes*, 128:41–48, 2017.
- [45] C Cazzaniga, E Andersson Sundén, Federico Binda, G Croci, Göran Ericsson, L Giacomelli, G Gorini, E Griesmayer, G Grosso, G Kaveney, et al. Single crystal diamond detector measurements of deuterium-deuterium and deuterium-tritium neutrons in Joint European Torus fusion plasmas. *Review of Scientific Instruments*, 85(4):043506, 2014.
- [46] Kuo Tian, Patrick Calderoni, Bradut-Eugen Ghidersa, and Axel Klix. Feasibility study of a neutron activation system for EU test blanket systems. *Fusion Engineering and Design*, 109-111:1517–1521, 2016. Proceedings of the 12th International Symposium on Fusion Nuclear Technology-12 (ISFNT-12).
- [47] Bethany Colling, P. Batistoni, S.C. Bradnam, Z. Ghani, M.R. Gilbert, C.R. Nobs, L.W. Packer, M. Pillon, and S. Popovichev. Testing of tritium breeder blanket activation foil spectrometer during JET operations. *Fusion Engineering and Design*, 136:258–264, 2018.
- [48] T. Kin, Y. Sanzen, M. Kamida, K. Aoki, N. Araki, and Y. Watanabe. Artificial Neural Network for Unfolding Accelerator-based Neutron Spectrum by Means of Multiple-foil Activation Method. In *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference, NSS/MIC 2017 - Conference Proceedings*. Institute of Electrical and Electronics Engineers Inc., 2018.
- [49] Davide Chiesa, Massimiliano Nastasi, Carlo Cazzaniga, Marica Rebai, Laura Arcidiacono, Ezio Previtali, Giuseppe Gorini, and Christopher D. Frost. Measurement of the neutron flux at spallation sources using multi-foil activation. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 902:14 – 24, 2018.
- [50] Davide Chiesa, Mario Carta, Valentina Fabrizio, Luca Falconi, Angelo Grossi, Massimiliano Nastasi, Mario Palomba, Stefano Pozzi, Ezio Previtali, Pier Giorgio Rancoita, Barbara Ranghetti, and Mauro Tacconi. Characterization of TRIGA RC-1 neutron irradiation facilities for radiation damage testing. *The European Physical Journal Plus*, 135(4):349, Apr 2020.
- [51] S.P. Tripathy, C. Sunil, M. Nandy, P.K. Sarkar, D.N. Sharma, and B. Mukherjee. Activation foils unfolding for neutron spectrometry: Comparison of different deconvolution methods. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 583(2):421 – 425, 2007.
- [52] Z.A. Kulage, C.H. Castano, S. Usman, and G. Mueller. Characterization of the neutron flux energy spectrum at the Missouri University of Science and Technology Research Reactor (MSTR). *Nuclear Engineering and Design*, 261:174–180, 2013.
- [53] Darrell Francis Newman. Determination of the absolute neutron flux spectrum in the PCTR fast neutron cavity from multiple foil activation measurements. Technical report, Battelle-Northwest, 1970.
- [54] S Berg. Modification of SAND II. Technical report, TRW Systems Group, Redondo Beach Calif., 1968.

- [55] K. C. Humpherys. A REVIEW OF SOME PASSIVE DOSIMETRY METHODS USED IN RADIATION EFFECTS STUDIES WITH FAST BURST REACTORS. Technical Paper Number S-50-TP. In *Proceedings of the National Topical Meeting on Fast Burst Reactors*, The University of New Mexico, Albuquerque, 1 1969. American Nuclear Society.
- [56] M. Reginatto and P. Goldhagen. MAXED, a computer code for maximum entropy deconvolution of multisphere neutron spectrometer data. *Health Phys.*, 77(5):579–83, 1999.
- [57] M. Matzke. Unfolding procedures. *Radiation Protection Dosimetry*, 107(1-3):155–174, 2003.
- [58] MI Savva, T Vasilopoulou, CR Nobs, P Batistoni, B Colling, Z Ghani, MR Gilbert, S Loreti, K Mergia, S Messoloras, et al. VERDI detector benchmark experiment at the ENEA 14 MeV Frascati Neutron Generator. *Fusion Engineering and Design*, 146:1877–1881, 2019.
- [59] O Ficker, J Mlynar, G Bonheure, A Murari, S Popovichev, et al. Unfolding of energies of fusion products measured by the activation probe at JET. In *Proceedings of 18th Conference of Czech and Slovak Physicists*, page 29, 2015.
- [60] L W Packer, P Batistoni, S C Bradnam, S Conroy, Z Ghani, M R Gilbert, E Laszyńska, I Lengar, M Nobs, CRand Pillon, S Popovichev, P Raj, IE Stamatelatos, T Vasilopoulou, A Wójcik-Gargula, R Worrall, and JET-Contributors. Neutron spectrum determination at the ITER material irradiation stations at JET. *UKAEA scientific publications*, 79, 9 2018.
- [61] Igor Lengar, Aljaž ČUfar, Vladimir Radulović, Paola Batistoni, Sergey Popovichev, Lee Packer, Zamir Ghani, Ivan A. Kodeli, Sean Conroy, and Luka Snoj. Activation material selection for multiple foil activation detectors in JET TT campaign. *Fusion Engineering and Design*, 136:988–992, 2018.
- [62] Sabina Griffith. Measuring ITER’s fusion power. *ITER Newslines*, Jun 2011.
- [63] R. Prokopowicz, B. Bienkowska, K. Drozdowicz, S. Jednorog, E. Kowalska-Strzeciwiłk, A. Murari, S. Popovichev, K. Pytel, M. Scholz, A. Szydlowski, B. Syme, and G. Tracz. Measurements of neutrons at JET by means of the activation methods. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 637(1):119–127, 2011.
- [64] Vitaly Krasilnikov, MunSeong Cheon, and Luciano Bertalot. Neutron Activation System for ITER Tokamak. In Lylia Alghem Hamidatou, editor, *Advanced Technologies and Applications of Neutron Activation Analysis*, chapter 5. IntechOpen, Rijeka, 2019.
- [65] L Bertalot, V Krasilnikov, L Core, A Saxena, N Yukhnov, R Barnsley, and M Walsh. Present status of ITER neutron diagnostics development. *Journal of Fusion Energy*, 38(3):283–290, 2019.
- [66] A Klix, A Domula, U Fischer, D Gehre, P Pereslavitsev, and I Rovni. Test facility for a neutron flux spectrometer system based on the foil activation technique for neutronics experiments with the ITER TBM. *Fusion Engineering and Design*, 86(9-11):2322–2325, 2011.
- [67] M S Cheon, Y S Lee, A C England, H S Kim, S Pak, C R Seon, and H G Lee. Diagnostic neutron activation system for KSTAR. *Journal of Instrumentation*, 7(05):C05009, may 2012.
- [68] Marcel Reginatto, Burkhard Wiegel, Andreas Zimbal, and Frank Langner (UMGPlot). NEA-1665 UMG 3.3. <https://www.oecd-neo.org/tools/abstract/detail/nea-1665>, 09 2004.

- [69] Manfred Matzke. HEPROW, unfolding of pulse height spectra using bayes theorem and maximum entropy method. *OECD Nuclear Energy Agency*, Oct 2004. <https://www.oecd-neo.org/tools/abstract/detail/nea-1666/>.
- [70] Roberto Bedogni, Carles Domingo, Adolfo Esposito, and Francisco Fernández. FRUIT: An operational tool for multisphere neutron spectrometry in workplaces. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 580(3):1301 – 1309, 2007.
- [71] David Thomas. Personal communications, Aug 2020.
- [72] FG Perey. Least-squares dosimetry unfolding: The program STAY'SL. Technical report, Oak Ridge National Lab., 1977.
- [73] R. Worrall, B. Colling, M.R. Gilbert, E. Litherland-Smith, C.R. Nobs, L.W. Packer, C. Wilson, and A. Zohar. The development, testing and comparison of unfolding methods in SPECTRA-UF for neutron spectrometry. *Fusion Engineering and Design*, 161:112038, 2020.
- [74] O.N Jarvis. Neutron spectrometry at JET (1983–1999). *Nuclear Inst. and Methods in Physics Research, A*, 476(1-2):474–484, 2002.
- [75] T.D. Beynon. Neutron flux discontinuities due to inelastic scattering. *Journal of Nuclear Energy*, 26(4):177–188, 1972.
- [76] A Turner, A Burns, B Colling, and J Leppanen. Applications of Serpent 2 Monte Carlo Code to ITER Neutronics Analysis. *Fusion Science and Technology*, 74(4):315–320, 2018.
- [77] Prasoon Raj, Steven C Bradnam, Bethany Colling, Axel Klix, Mitja Majerle, Chantal R Nobs, Lee W Packer, Mario Pillon, Milan Štefánik, and JET Contributors. Evaluation of the spectrum unfolding methodology for neutron activation system of fusion devices. *Fusion Engineering and Design*, 146:1272–1275, 2019.
- [78] MJ Fleming, LWG Morgan, and E Shwageraus. Optimization algorithms for multigroup energy structures. *Nuclear Science and Engineering*, 183(2):173–184, 2016.
- [79] Manfred Matzke. Propagation of uncertainties in unfolding procedures. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 476(1):230 – 241, 2002.
- [80] Wolfgang von der Linden. Maximum-entropy data analysis. *Applied Physics A*, 60(2):155–165, 1995.
- [81] Mikael Kuusela. Statistical Issues in Unfolding Methods for High Energy Physics. Master's thesis, Aalto University, July 2012.
- [82] M Matzke. Unfolding procedures. *Radiation protection dosimetry*, 107:155–74, 02 2003.
- [83] Manfred Matzke. Unfolding of pulse height spectra: the HEPRO program system. Technical report, SCAN-9501291, 1994.
- [84] P. Batistoni, U. Fischer, K. Ochiai, L. Petrizzi, K. Seidel, and M. Youssef. Neutronics and nuclear data issues in ITER and their validation. *Fusion Engineering and Design*, 83(7):834–841, 2008. Proceedings of the Eight International Symposium of Fusion Nuclear Technology.
- [85] FISPACT-II wiki Refernce spectra. https://fispact.ukaea.uk/wiki/Reference_input_spectra#ITER-DT. Accessed: Sep-2020.

- [86] Steven J. Zinkle and Anton Möslang. Evaluation of irradiation facility options for fusion materials research and development. *Fusion Engineering and Design*, 88(6):472–482, 2013. Proceedings of the 27th Symposium On Fusion Technology (SOFT-27); Liège, Belgium, September 24-28, 2012.
- [87] Jiří CVACHOVEC. *Algorithms for processing experimental data from radiation measurements*. PhD thesis, Masarykova univerzita, Fakulta informatiky, 2011.
- [88] J.M. Ortiz-Rodríguez, A. Reyes Alfaro, A. Reyes Haro, J.M. Cervantes Viramontes, and H.R. Vega-Carrillo. A neutron spectrum unfolding computer code based on artificial neural networks. *Radiation Physics and Chemistry*, 95:428–431, 2014.
- [89] Amin Asgharzadeh Alvar, Mohammad Reza Deevband, and Meghdad Ashtiyani. Neutron spectrum unfolding using radial basis function neural networks. *Applied Radiation and Isotopes*, 129:35–41, 2017.
- [90] Ma. Del Rosario Martinez-Blanco, Gerardo Ornelas-Vargas, Celina Lizeth Castañeda-Miranda, Luis Octavio Solís-Sánchez, Rodrigo Castañeda-Miranada, Héctor René Vega-Carrillo, Jose M Celaya-Padilla, Idalia Garza-Veloz, Margarita Martínez-Fierro, and José Manuel Ortiz-Rodríguez. A neutron spectrum unfolding code based on generalized regression artificial neural networks. *Applied Radiation and Isotopes*, 117:8–14, 2016.
- [91] Seyed Abolfazl Hosseini. Neutron spectrum unfolding using artificial neural network and modified least square method. *Radiation Physics and Chemistry*, 126:75–84, 2016.
- [92] Emanuele Ronchi. *Neural Networks Applications and Electronics Development for Nuclear Fusion Neutron Diagnostics*. PhD thesis, Uppsala University, Department of Physics and Astronomy, 2009.
- [93] Arun K. Understanding curse of dimensionality. <https://www.mygreatlearning.com/blog/understanding-curse-of-dimensionality/>, 2022.
- [94] O. Wong. Fusion Neutron Activation Spectra Unfolding by Neural Networks(FACTIUNN). Master’s thesis, University of Birmingham, 2019.
- [95] Jie Wang, Yulin Zhou, Zhirong Guo, and Haifeng Liu. Neutron spectrum unfolding using three artificial intelligence optimization methods. *Applied Radiation and Isotopes*, 147:136–143, 2019.
- [96] Vitisha Suman and P.K. Sarkar. Neutron spectrum unfolding using genetic algorithm in a Monte Carlo simulation. *Nuclear Inst. and Methods in Physics Research, A*, 737:76–86, 2014.
- [97] Bhaskar Mukherjee. A high-resolution neutron spectra unfolding method using the Genetic Algorithm technique. *Nuclear Inst. and Methods in Physics Research, A*, 476(1-2):247–251, 2002.
- [98] Ross Worrall. Developing a genetic algorithm for the unfolding of neutron energy spectra. Master’s thesis, University of Birmingham, Culham Science Centre, Abingdon OX14 3EB, 9 2014.
- [99] H. Mazrou and F. Bezoubiri. Evaluation of a neutron spectrum from Bonner spheres measurements using a Bayesian parameter estimation combined with the traditional unfolding methods. *Radiation Physics and Chemistry*, 148:33–42, 2018.
- [100] Jacob Eriksson, Sean Conroy, E Andersson Sundén, and Carl Hellesen. Calculating fusion neutron energy spectra from arbitrary reactant distributions. *Computer Physics Communications*, 199:40–46, 2016.

- [101] D. Moseev, M. Salewski, M. Garcia-Muñoz, B. Geiger, and M. Nocente. Recent progress in fast-ion diagnostics for magnetically confined plasmas. *Reviews of Modern Plasma Physics*, 2(1):7, Sep 2018.
- [102] Jacob Eriksson. *Neutron emission spectrometry for fusion reactor diagnosis: Method development and data analysis*. PhD thesis, Uppsala University, Department of Physics and Astronomy, 2015.
- [103] Žiga Štancar. Personal communications, Jan 2022.
- [104] L. Ballabio, G. Gorini, and J. Kallne. Energy spectrum of thermonuclear neutrons. *Review of Scientific Instruments*, 68(1):585–588, 1997.
- [105] Carlo Cazzaniga. *Fast neutron measurements for fusion and spallation sources*. PhD thesis, University of Milano Bicocca, 2014.
- [106] M. Gatu Johnson, L. Giacomelli, A. Hjalmarsson, M. Weiszflog, E. Andersson Sundén, S. Conroy, G. Ericsson, C. Hellesen, J. Källne, E. Ronchi, H. Sjöstrand, G. Gorini, M. Tardocchi, A. Murari, S. Popovichev, J. Sousa, R. C. Pereira, A. Combo, and N. Cruz. The TOFOR neutron spectrometer and its first use at JET. *Review of Scientific Instruments*, 77(10):10E702, 2006.
- [107] P C Hansen. Numerical tools for analysis and solution of Fredholm integral equations of the first kind. *Inverse Problems*, 8(6):849, dec 1992.
- [108] Per Christian Hansen. The L-curve and its use in the numerical treatment of inverse problems. In *Computational Inverse Problems in Electrocardiology*. WIT Press, 2000. Invite Computational Inverse Problems in Electrocardiology ; Conference date: 01-01-2000.
- [109] Erik Andersson Sundén, S Conroy, G Ericsson, M Gatu Johnson, L Giacomelli, C Hellesen, A Hjalmarsson, J Källne, E Ronchi, H Sjöstrand, et al. Evaluation of spectral unfolding techniques for neutron spectroscopy. In *AIP Conference Proceedings*, volume 988, pages 315–318. American Institute of Physics, 2008.
- [110] Tim Adye. Unfolding algorithms and tests using RooUnfold. *arXiv preprint arXiv:1105.1160*, 2011.
- [111] Olivier Besida. An hybrid Tykhonov method for neutron spectrum unfolding. *arXiv preprint nucl-ex/0506030*, 2005.
- [112] Shigetaka Maeda and Tetsuo Iguchi. A new unfolding code combining maximum entropy and maximum likelihood for neutron spectrum measurement. *Journal of Nuclear Science and Technology*, 50(4):381–386, 2013.
- [113] J Cvachovec and F Cvachovec. Maximum likelihood estimation of a neutron spectrum and associated uncertainties. *Advances in Military Technology*, 1(2):67–79, 2008.
- [114] Fatma Zohra Dehimi, Abdeslam Seghour, and Saddik El Hak Abaidia. Unfolding of neutron energy spectra with fisher regularisation. *IEEE Transactions on Nuclear Science*, 57(2):768–774, 2010.
- [115] Nicholas J. Quartemont and James E. Bevins. Regularized heuristic method for activation monitor neutron spectrum unfolding. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1014:165722, 2021.
- [116] S. F. Gull and J. Skilling. Maximum entropy method in image processing. *IEE Proceedings F - Communications, Radar and Signal Processing*, 131(6):646–659, 1984.

- [117] Arthur Hobson. A new theorem of information theory. *Journal of Statistical Physics*, 1(3):383–391, 1969.
- [118] Kenneth P. Burnham and David Raymond Anderson. *Model selection and multimodel inference : a practical information-theoretic approach*. Springer, New York, London, 2nd ed edition, 2002.
- [119] WN McElroy, S Berg, T Crockett, and RG Hawkins. A computer-automated iterative method for neutron flux spectra determination by foil activation. Volume 1. A study of the iterative method. Technical report, ATOMICS INTERNATIONAL CANOGA PARK CA, 1967.
- [120] G. L. Guthrie and R. L. Simons. Analysis and extension of the SAND-II code in damage function unfolding applications. [BWR; PWR]. In *Conference: 9. symposium on effects of radiation in structural materials*. Westinghouse Hanford Company, jul 1978. HEDL-SA-1475; CONF-780722-5; TRN: 79-000100.
- [121] Gene H. Golub and Charles F. Van Loan. *Matrix Computations (3rd Ed.)*. Johns Hopkins University Press, USA, 1996.
- [122] DW Wootan and F Schmittroth. Comparison of SAND-II and FERRET. Technical report, Hanford Engineering Development Lab., Richland, WA (USA), 1981.
- [123] Riley Murray, James Demmel, Michael W. Mahoney, N. Benjamin Erichson, Maksim Melnichenko, Osman Asif Malik, Laura Grigori, Piotr Luszczek, Michal Derezinski, Miles E. Lopes, Tianyu Liang, Hengrui Luo, and Jack Dongarra. Randomized numerical linear algebra: A perspective on the field with an eye to software. Technical Report UCB/EECS-2023-19, Feb 2023.
- [124] CA Oster. Review of unfolding methods used in the US and their standardization for dosimetry. Technical report, Battelle Pacific Northwest Labs., Richland, WA (USA), 1977.
- [125] W Meyer, W H Miller, and C B Wallace. Standardization of neutron spectrum unfolding codes: the benchmarking problem. In *American Nuclear Society annual meeting*, volume 34. American Nuclear Society, Jun 1980.
- [126] A Klix, A Domula, U Fischer, D Gehre, P Pereslavytsev, and I Rovni. Neutronics diagnostics for European ITER TBMs: Activation foil spectrometer for short measurement cycles. *Fusion Engineering and Design*, 87(7-8):1301–1306, 2012.
- [127] Héctor René Vega Carrillo and M.Pilar Iñiguez de La Torre. Catalogue to select the initial guess spectrum during unfolding. *Nuclear Inst. and Methods in Physics Research, A*, 476(1-2):270–272, 2002.
- [128] Manfred Matzke. Propagation of uncertainties in unfolding procedures. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 476(1-2):230–241, 2002.
- [129] Ocean Wong. unfoldingsuite. <https://doi.org/10.5281/zenodo.10982488>, 04 2024.
- [130] O. Wong, R. Smith, C. R. Nobs, and A. M. Bruce. Optimising Foil Selection for Neutron Activation Systems. *Journal of Fusion Energy*, 41(1):12, Jun 2022.
- [131] Mitja Majerle, Martin Ansorge, Pavel Bém, Daniil Koliadko, Jaromír Mrázek, Jan Novák, Eva Šimečková, Milan Štefánik, Michal Košťál, Zdeněk Matěj, Haridev Chohan, and Ocean Wong. Measurements of the neutron spectra from the p+Be neutron generator of the NPI CAS. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1053:168314, 2023.
- [132] Ocean Wong. foilselector. <https://doi.org/10.5281/zenodo.11194361>, 05 2024.

- [133] LW Packer, P Batistoni, SC Bradnam, B Colling, Sean Conroy, Z Ghani, MR Gilbert, S Jednorog, E Łaszyńska, D Leichtle, et al. Activation of ITER materials in JET: nuclear characterisation experiments for the long-term irradiation station. *Nuclear Fusion*, 58(9):096013, 2018.
- [134] Roberto Capote, Konstantin I. Zolotarev, Vladimir G. Pronyaev, and Andrej Trkov. Updating and Extending the IRDF-2002 Dosimetry Library. In *Reactor Dosimetry: 14th International Symposium*, pages 197–209, West Conshohocken, PA, Aug 2012. ASTM International.
- [135] P. Batistoni, S. Conroy, S. Lilley, J. Naish, B. Obryk, S. Popovichev, I. Stamatelatos, B. Syme, T. Vasilopoulou, and JET contributors. Benchmark experiments on neutron streaming through JET Torus Hall penetrations. *Nuclear Fusion*, 55(5):053028, apr 2015.
- [136] James F. Ziegler, M.D. Ziegler, and J.P. Biersack. SRIM – The stopping and range of ions in matter (2010). *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 268(11):1818–1823, 2010. 19th International Conference on Ion Beam Analysis.
- [137] AB Pashchenko, H Wienke, and S Ganesan. FENDL: International reference nuclear data library for fusion applications. *Journal of nuclear materials*, 233:1601–1606, 1996.
- [138] J-Ch Sublet, Arjan Koning, and Dimitri Rochman. TENDL-2017: The making of multifaceted technological nuclear data library. *Ratio*, 3:103, 2018.
- [139] Davide Chiesa, Ezio Previtali, and Monica Sisti. Bayesian statistics applied to neutron activation data for reactor flux spectrum analysis. *Annals of Nuclear Energy*, 70:157–168, 2014.
- [140] M. Angelone, P. Batistoni, M. Laubenstein, L. Petrizzi, and M. Pillon. Neutronics experiment for the validation of activation properties of DEMO materials using real DT neutron spectrum at JET. *Fusion Engineering and Design*, 81(8):1485–1490, 2006.
- [141] A.S. Jain and S. Meeran. Deterministic job-shop scheduling: Past, present and future. *European Journal of Operational Research*, 113(2):390–434, 1999.
- [142] Cross Section Evaluation Working Group (CSEWG). ENDF-102 data formats and procedures for the evaluated nuclear data file ENDF-6. Technical report, Cross Sections Evaluation Working Group, National Nuclear Data Center, Brookhaven National Laboratory, Upton, NY 11973-5000, 07 2010. Document ENDF-102, Report BNL-90365-2009.
- [143] Paul K. Romano, Nicholas E. Horelik, Bryan R. Herman, Adam G. Nelson, Benoit Forget, and Kord Smith. OpenMC: A state-of-the-art Monte Carlo code for research and development. *Annals of Nuclear Energy*, 82:90–97, 2015. Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms.
- [144] Nuclear data library downloads, 01 2023. <https://fispact.ukaea.uk/nuclear-data/downloads/>.
- [145] Cross Section Evaluation Working Group (CSEWG). ENDF Appendix B: Definition of Reaction Types. https://www.oecd-neo.org/dbdata/data/manual-endf/endf102_MT.pdf. Accessed: 5-3-2022.
- [146] Tomasz Piotrowski, María José Martínez-Echevarría Romero, Piotr Prochoń, Mónica López Alonso, Rafał Michalczyk, Armando Arvizu-Montes, Łukasz Ciupiński, Santiago Becerril Jarque, Kazimierz Józeffiak, Yuefeng Qiu, Martin Ansorge, Hari Chohan, and Magdalena Wojtkowska. Optimization and evaluation of structural and shielding concrete for IFMIF-DONES. *Nuclear Materials and Energy*, 38:101597, 2024.

- [147] G. Schnabel, D. L. Aldama, and R. Capote. How to explain ENDF-6 to computers: A formal ENDF format description language. 2023. arXiv:2312.08249, DOI:10.48550/arXiv.2312.08249.
- [148] Christopher D Frost, Stuart Ansell, and Giuseppe Gorini. A new dedicated neutron facility for accelerated SEE testing at the ISIS facility. In *2009 IEEE International Reliability Physics Symposium*, pages 952–955, 2009.
- [149] S.A. Wender, S. Balestrini, A. Brown, R.C. Haight, C.M. Laymon, T.M. Lee, P.W. Lisowski, W. McCorkle, R.O. Nelson, W. Parker, and N.W. Hill. A fission ionization detector for neutron flux measurements at a spallation source. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 336(1):226–231, 1993.
- [150] Roberto Bedogni, Adolfo Esposito, Carla Andreani, Roberto Senesi, Maria Pia De Pascale, Piergiorgio Picozza, Antonino Pietropaolo, Giuseppe Gorini, Christopher D. Frost, and Stewart Ansell. Characterization of the neutron field at the ISIS-VESUVIO facility by means of a Bonner sphere spectrometer. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 612(1):143–148, 2009.
- [151] J.W.G. Thomason. The ISIS Spallation Neutron and Muon Source—The first thirty-three years. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 917:61 – 67, 2019.
- [152] CCFE-709 group structure - FISPACT-II Wiki. https://fispact.ukaea.uk/wiki/CCFE-709_group_structure. Accessed: Sep-2020.
- [153] Canberra Industries Inc. Genie™ 2000 spectroscopy software, 2004.
- [154] Mirion Technologies. ISOCS. <https://www.mirion.com/isocs>.
- [155] Leen Verheyen and Michel Bruggeman. Genie2k library creator. *SCK-CEN*, 2017. https://www.nks.org/download/verheyen_genie2k_library_creator.pdf.
- [156] Chris Mallett, Steve Gray, A. Ahmed, and Charlie Simmons. AutoHotKey. <https://github.com/AutoHotkey/AutoHotkey>, 2013.
- [157] CA Oster, WN McElroy, and JM Marr. Monte Carlo program for SAND-II error analysis. Technical report, Hanford Engineering Development Lab., Richland, WA (United States), 1973.
- [158] CA Oster, WN McElroy, RL Simons, EP Lippincott, and GR Odette. Modified Monte Carlo program for SAND-II with solution weighting and error analysis. Technical report, Hanford Engineering Development Lab., Richland, WA (United States), 1976.