Sheffield Hallam University

Classifying UK charities’ activities by charitable cause: a new classification system

DAMM, Christopher <http://orcid.org/0000-0002-7355-3496> and KANE, David

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

https://shura.shu.ac.uk/31791/

This document is the Accepted Version [AM]

Citation:

DAMM, Christopher and KANE, David (2022). Classifying UK charities’ activities by charitable cause: a new classification system. Voluntary Sector Review, 1-27. [Article]

Copyright and re-use policy

See http://shura.shu.ac.uk/information.html

Sheffield Hallam University Research Archive
http://shura.shu.ac.uk

Title:

Classifying UK charities' activities by charitable cause: a new classification system.

Abstract:

This article presents a new system for classifying UK charities' activities according to their charitable purposes. It also outlines our attempts to use keyword search rules to apply these classifications to the various UK charity registers. The classification results and code, which are made freely available online, help to address the limitations of existing classification schemes in the UK context. Depending on the scheme, these include a lack of detail and coverage of important sub-sectors, a lack of systematic data collection, and limits on the number of classifications per charity. We discuss the pros and cons of different approaches and show that the keyword searching method provides a sufficiently accurate and transparent approach. We also present some preliminary results on how commonly each 'tag' is matched against UK charities, as well as exploring how the results compare to existing classifications in the register of charities for England and Wales.

Authors:

Christopher Damm (Sheffield Hallam University)

David Kane (independent researcher).

Keywords:

Classification, charities, purposes, activities, UK-CAT

Acknowledgements

This article draws on a collaboration between the Centre for Regional Economic and Social Research, NCVO and David Kane, funded by a grant from the Esmée Fairbairn Foundation. We thank them for their support.

We would like also to acknowledge Oliver Chan for his work as a research assistant on the project and especially his tireless work helping to code the manually classified dataset.

Thanks also to Caitlin Curran for some additional help coding the manual dataset as part of her placement at CRESR, Lisa Hornung and her initial involvement in setting up the project at NCVO, Lorna Dowrick for detailed comments on the UK-CAT, and two anonymous reviewers for their helpful comments on a draft of this paper.

Finally, we would also like to acknowledge the work of Lester Salamon, who passed away in August 2021. Lester was a professor at Johns Hopkins University who first designed the ICNPO classification along with Helmut Anheier.

Introduction

Formal, independent voluntary organisations are united by their non-profit status, and depending on the definition used, their voluntarism and social mission (Salamon and Anheier, 1992b). They can, however, differ wildly in terms of just about everything else (Kendall and Knapp, 1995). Analyses at the level of entire sectors may, therefore, obscure interesting differences, such as in voluntary organisations' cause, or social mission.

There are good reasons to make these distinctions. As Kendall (2003) has argued, voluntary organisations' regulatory and policy environments, which shape many of their experiences, are often defined primarily by their activities. They determine which other organisations voluntary organisations interact with, creating fields in which norms and institutions spread more easily (DiMaggio, 1983; Fligstein and McAdam, 2011). In practical terms, researchers may wish to sample case studies from within a particular activity area, include 'activity-area' as a covariate within statistical modelling, compare one sub-sector with another, or compare for- and non-profit providers within a single activity area. Classifications of different activity areas are also used by government agencies in the composition of national accounts.

A focus on the purposes of charitable activities may also be partly due an interest in the level of economic activity within the sector (Salamon and Anheier, 1992b; Kendall and Knapp, 1995). Although this somewhat instrumentalist emphasis is sometimes criticised as reductive, or disproportionately benefiting elites (Barman, 2013; Nickel and Eikenberry, 2015), the idea is that unless rigorously mapped, the size, scope, and importance of voluntary sector activity is often overlooked (Salamon, 2010). In the UK, some have suggested that a lack of data on these activities may have contributed to the non-profit sector losing out on COVID-19 related government relief (Kenley and Wilding, 2021).

A focus on what voluntary organisations *do* can also point towards the social benefit they provide. Mapping these different activities, and their purposes, can paint a more vivid picture of where voluntary organisations expend their energy. Many attempts to map the sector by charitable cause originate from within voluntary organisations and their own infrastructure bodies (NCVO, 2012; Newbigging et al., 2017), often with the explicit goal of drawing more attention to the importance of a particular sub-sector.

Despite ample interest, in several countries prominent classification schemes have notable limitations. In the US, a small but growing literature addresses concerns with the National Taxonomy of Exempt Entities (NTEE) and its use of Form 990 data (Ma, 2020; Fyall et al., 2020; Lampkin et al., 2001). In European countries such as Austria, attempts have focussed on applying the International Classification of Non-profit and Third Sector Organisations (ICNPTSO) (Litofcenko et al., 2020). We are also aware of attempts to apply bespoke classification systems to Australian voluntary organisations (Our Community, 2020).

In the UK, which provides the main empirical focus of this article, the most comprehensive data sources on voluntary organisations are the regulatory registers of charities in England and Wales, Scotland, and Northern Ireland. As with other countries, however, there are several potential classificatory challenges. The classifications included in the registers, and alternatives such as the ICNPTSO, map poorly onto some of the key areas of interest in a UK context. ICNPTSO categories are not collected directly from charities, while the self-selection of the regulatory categories brings its own challenges.

This article outlines our efforts to address some of these challenges. First, we present a new classification system, the ‘UK Charitable Activity Tags’ (UK-CAT), developed specifically to capture the local context of UK charitable organisations. This scheme has more categories than the existing alternatives and allows multiple classifications per charity, providing a higher level of detail. Second, we also outline our attempts to use automated keyword matching to apply this new system to the various UK charity registers, in contrast to either self-selection by charities or allocation by a researcher or regulator.

The first contribution of this article is, therefore, to the growing, international, methodological literature on classifying voluntary sector organisations, reflecting on the advantages and disadvantages of the new UK-CAT and comparing it to alternatives. In addition, we present preliminary, descriptive findings on the number of charities matching each tag, their combined income, and how these findings compare to existing results from the register of charities in England and Wales. Finally, all the python code has been made available online via Github[1], as well as an implementation of the classification system, applied to the various UK charity registers. This should allow researchers to replicate and expand the UK-CAT, as well as use the categories in their own research.

[1] https://github.com/charity-classification

Literature review

This article focuses specifically on the classification of voluntary organisations by their charitable cause, purpose, or mission. These concepts are largely interchangeable and are usually articulated as a particular activity, beneficiary group, or a targeted problem, which best encapsulates a voluntary organisation's motivation for conducting its activities. In the UK, by far the most data is available on formally registered charities, so for pragmatic reasons this is where we focus here.

In this section, we first present a brief review of the current classification systems most relevant to the UK case. Second, we outline some of the different data sources to which these classifications have been applied, and third, the methods used to apply them. And finally, we address an ongoing debate concerning the ethics of classification.

Existing classification schemes

Several prominent classification systems already exist to classify UK voluntary organisations. As part of their registration process, charities in England and Wales are asked to select from several drop-down lists, identifying what the charity does, who it helps, and how it operates. In practice, there is overlap between these lists. For example, 'disability' appears in more than one drop-down. Some categories are also very broad, including 'general charitable purposes'. More specific categories, however, important in the UK context, are not included, such as food banks, homelessness support, or medical research[2].

The Northern Irish process is similar, with two lists both essentially identifying charitable purposes, and a third identifying beneficiary groups. Again, there is some overlap and repetition. Finally, the Scottish Register of Charities contains a very similar set of three lists, but these are applied post-hoc by officials at the Office of the Scottish Charity Regulator (OSCR).

The National Taxonomy of Exempt Entities (NTEE) is also worth noting. As the most widely used classification system of tax-exempt non-profits' purposes in the US, it has relatively high prominence in the academic literature (Fyall, et al. 2018; Lampkin et al., 2001; Ma, 2020). In 1995, the Internal Revenue Service (IRS), the US's tax collection agency, took responsibility for assigning the classifications to non-profits when they apply for tax-exempt status (Ma, 2020).

In contrast to the nationally specific schemes found in the UK, US and elsewhere, the ICNPTSO is designed to be applied internationally, to enable cross country comparisons (Salamon and Anheier, 1992). In the US or UK, and to our knowledge internationally, ICNPTSO categories are not recorded as part of any formal registration process, which means they are allocated retrospectively. To work across many national contexts, the categories are quite broad, and the system is, again, sometimes a poor fit for some UK charities.

All these systems have their advantages and disadvantages. We argue that there is a clear gap within the classificatory infrastructure, however, for a scheme that provides greater detail for the UK context. This detail should prove valuable to researchers, even at the cost of more categories, less international generalisability, and relying on post-hoc classification.

[2] During the course of this research project, we have been involved in a consultation process run by the Charity Commission for England and Wales, to update and expand the list of categories that they use. In some cases, the new categories may help to address some, though not all, of these gaps.

A second important distinction between different systems is whether the categories are mutually exclusive. Both the NTEE and the ICNPTSO systems have traditionally applied a single classification per organisation, which avoids any double counting of economic activity and makes it more straightforward to incorporate the results into quantitative modelling. However, a single category can fail to capture the multiple or combined purposes of many voluntary organisations (Fyall et al. 2018; Lampkin et al. 2001; Ma, 2020). On the other hand, when organisations self-select many options, and there is no ranking system, it can be difficult to interpret which choices are most meaningful. Whether multiple or single classifications are best, therefore, depends on the research task at hand.

Data sources

The second major limiting factor when applying a classification scheme is the data available. At one extreme, we might only have organisations' names, which can be uninformative (Litofcenko et al., 2020). At the other end of the spectrum, if regulatory staff apply classifications, they have complete application information. The staff from the individual voluntary organisations, of course, have first-hand knowledge to draw on.

Those wishing to apply classifications to UK charities post-hoc, however, generally have access to organisations' charitable 'objects', a legally required paragraph within the charity's governing documentation setting out their purpose and objectives. In the case of England and Wales, and Northern Ireland, we also have access to a written description of their charitable activities. Compared to the formal objects, these may use more modern language, be more up to date, and contain less legalese (Leung, 2020). Both these text fields are comparable to the 'mission statement' in Form 990 data in the US (Fyall et al. 2018).

Whilst very useful for the purposes of classification, these textual data sources can introduce several challenges (Fyall, et al. 2018; Lampkin et al. 2001; Leung, 2020; Ma, 2020). First, charities may omit important aspects of their work, or even fail to complete the relevant section. Second, the quality is uneven. Charities will sometimes write extremely general, uninformative clauses, such as 'general charitable purposes'. Third, these records are rarely updated, meaning that the records can become significantly out of date.

Classification methods

Regardless of the textual data available, there are several options for how to apply classifications (Ma, 2020). First, voluntary organisations, or external individuals, can apply the classifications manually. Alternatively, human coders or a supervised machine learning process can create a programmed set of rules to automatically apply a predefined set of categories. Or finally, an unsupervised machine learning model can derive its own categories, based on recurring patterns in the data.

The first, and most self-explanatory method, is self-selection, as conducted by charities in England, Wales and Northern Ireland. The person completing the relevant form should know the charity well, and the amount of work per charity is relatively small (though not negligible). On the other hand, organisations may select a high number of categories, lack expertise on the classification system, and their reasoning for choosing a category may not always be clear.

A second option is manual classification by a researcher or the regulator (Ma, 2020). This individual may still have access to a reasonable amount of data and be familiar with the classification system. The main limitation is likely to be their time. As the number of classifications needed increases in size, the less feasible human classification becomes.

Litofcenko, Karner and Maier (2020) created a manually classified dataset of 5,000 Austrian non-profits (as well as a separate sample of 1,000 German non-profits), to help assess the results of their automated methods. The 'correct' classification was allocated by consensus. Individual coders achieved between 79 and 87 per cent agreement against this final allocation. The authors highlight the relatively high level of expertise needed, extensive amount of time spent, and the relatively low transparency.

In contrast, there are a range of automated classification options. Ma (2020) distinguishes between 'dictionary methods', and both supervised and unsupervised machine learning. Automated methods can be run consistently and indefinitely, using as much textual data as is available.

Dictionary methods search for keywords that either increase or decrease the probability of a category being relevant. Fyall, Moore and Gugerty (2018) used keywords to determine the probability of a voluntary organisation providing homelessness accommodation. Litofcenko, Karner and Maier (2020) also used keywords as part of multiple, tiered if-then statements, to allocate each voluntary organisation to an ICNPTSO category.

The success of dictionary methods depends partly on the quality of the underlying data, how clearly defined the categories are, and the effectiveness of the chosen keywords. Inevitably, there will be both false positives and false negatives (Fyall et al. 2018). Litofcenko, Karner and Maier (2020), using just organisational names, suggest that their keyword algorithm was correct 85 per cent of the time, similar to the results achieved by an individual human coder. Fyall, Moore and Gugerty (2018) did not have a manually classified dataset to compare against but did achieve significantly more matches amongst their sample of Washington State non-profits than relying on NTEE categories or regulatory listings alone.

In an unsupervised machine learning model, an algorithm uncovers patterns in the text without using a prior set of classifications (Ma, 2020). Leung (2020) used an unsupervised method of natural language processing and clustering to identify recurring terms within the activities and objects of 'arts, culture, heritage or science' organisations. The result was an automatically generated, hierarchical taxonomy of keywords. Although Leung's categories are generally meaningful and very detailed, there is inevitably a greater risk with unsupervised models that the categories may not be as theoretically useful as those developed manually (Ma, 2020).

Finally, as with the dictionary method, a supervised machine learning model uses an existing list of categories. The rules and keywords used, however, are derived inductively by a machine learning algorithm using a training dataset, containing 'correct' classifications. This is then tested, before being used on new cases. Classification based on text is a relatively common machine learning task (Ma, 2020; Lantz, 2015).

In addition to their keyword-based rules, Litofcenko, Karner and Maier (2020) experimented with a decision tree machine learning algorithm. They found the results unsatisfactory and suggested this may be because they were forced to rely on names and web scraped data. None of their models classified more than 50 per cent of the test sample correctly.

Ma (2020), using relatively advanced machine learning methods, experimented with several different models and parameters to apply NTEE categories using Form 990 textual data. The most successful is called a BERT classifier (Bidirectional Encoder Representations from Transformers) and achieved 90 per cent overall accuracy using the nine broad NTEE categories and 88 per cent for the 25 major groups. This suggests that a high level of success may be possible, given enough cases, high-quality textual data, and relevant expertise.

Ethics of classification

Classification is intrinsic to human understanding, used to divide an endless mass of phenomena into workable, comprehensible definitions (Bowker and Star, 1999; Barman, 2013). Nevertheless, it can be a controversial process. Classifications are socially constructed, and will, therefore, inevitably embody a particular set of ethical and political values (Bowker and Star, 1999). Some entities are highlighted. Others are made less visible, forced under ill-suited headings, or submerged within 'other' categories.

Barman (2013) argues that classification systems are always the product of struggles between groups for various forms of capital and lead to inequality and domination. For example, Barman suggests that the NTEE was formed as part of a larger struggle by powerful elites and scholars to protect the tax-exempt status of foundations. This argument is similar to those who critique the 'mapping' of non-profits, especially as a prelude to state interference (Nickel and Eikenberry, 2016).

We agree that classifications are socially constructed, inevitably subject to bias, and often reflect wider power relations. Appe (2012) is correct to point out that who measures, and why they measure, matters. LePere-Schloop et al. (2021) add that the 'how' is also important, particularly as some methods are more comprehensible to outside observers than others. We do not necessarily believe, however, that classification must be to the detriment of voluntary organisations, or to the primary benefit of elite groups. Arguably, some critiques underplay more benign motivations, such as an intrinsic interest in the sector, or in the case of voluntary sector infrastructure bodies, desiring improvements for frontline voluntary organisations.

We are aware, however, that our own, largely self-appointed role as classifiers, involves choices with ethical implications. Bowker and Star (1999) offer helpful advice. First, a healthy degree of self-reflexivity is important, constantly asking which groups are being made more or less visible and reflecting on our own positionality (LePere-Schloop et al. 2021).

Second, transparency ensures that systems can be critiqued and therefore improved. As such, we seek to present an extensive account of our methods, especially our keyword choices, as well as making all the python code and classification results from this project open source[3]. We have also tried to include as much relevant information as possible on the project website[4].

[3] https://github.com/charity-classification
[4] https://charityclassification.org.uk/

Finally, “the only good classification is a living classification” (Bowker and Star, 1999, p.326). We hope, based on feedback, to revise and improve the classification system presented here, and its automated application, in the years ahead. We would warn again any temptation to see the results as definitive, rather than just one constructed viewpoint, albeit one which we hope will be of use more widely.

Methodology

In response to the issues raised with the existing classification schemes, this project aimed to develop a new classification system tailored to the UK charitable context. We have called this schema the UK Charity Activity Tags (UK-CAT). It has considerably more categories than either the classifications from the UK registers of charities or the ICNPTSO, providing greater depth, at the expense of parsimony.

Secondly, we aimed to design a means of automatically applying the UK-CAT to all charities registered in the UK, given the impracticality of applying them manually. We also allowed multiple tags to be applied per charity, similar to the UK regulatory classifications, but in contrast to most implementations of the ICNPTSO.

To summarise, Table 1 summarises how the UK-CAT fits alongside the range of existing schemes currently used in the UK.

Table 1: summary of the main classification systems used in the UK

Classification system	Application method	Single or multiple choice	Structure
Register of Charities - England and Wales	Self-selected by charities	Multiple	Three 'types' of category: 'what does your charity do?' (17 options), 'Who does your charity help?' (seven options) and 'how does your charity operate' (ten options).
Register of Charities - Northern Ireland	Self-selected by charities	Multiple	Three 'types' of category: 'what the charity does' (12 options), 'Who does your charity help?' (32 options) and 'how does your charity operate' (33 options).
Register of Charities - Scotland	Applied by the regulator	Multiple	Three 'types' of category: 'charitable purposes' (16 options), 'beneficiaries' (seven options) and 'type of activity' (four options).
ICNPTSO	Keyword matching	Generally single	12 'sections', 50 'groups', and 65 'sub-groups'.
UK-CAT	Keyword matching	Multiple	24 'categories', 17 optional 'sub-categories' and 211 lower level 'tags'. Covers a variety of charitable 'causes'.

The UK-CAT consists of a list of 'charity activity tags', so called because the categories are applied to the activity descriptions of UK charities. The classifications themselves are perhaps better described more broadly as 'causes', or descriptions of a charity's 'mission'. In other words, the subject matter that motivates them to operate. As such, categories include a variety of different *types* of things, reflecting the variety with which charities articulate their missions. Sometimes a charitable cause is articulated as a problem to tackle, such as 'poverty'. Sometimes it is something generally positive, such as a sport. At others times it is a particular group of beneficiaries. More rarely, we have allowed a type of activity, such as grant making or acting as an umbrella body, or a facility such as playing fields.

Combining all these different types of tag has helped to eliminate some of the duplication and overlap found within other classification systems. For analytical purposes, however, it may sometimes make sense to use only parts of the UK-CAT, homing in on different types of beneficiary group, for example.

Training dataset and the UK-CAT

The first step was to create a combined population dataset of all active UK charities (N = 201,963), using the various national registers of charities. Keyword searches were used to remove a small number of very easily identifiable groups, such as scout associations, based on their name only (totalling 25,852 charities), to ensure more variety in the remaining charities.

A sample of 4,200 charities was then selected, including 1,325 charities with an income of £100,000 or over and 2,875 with a lower income. This over-sampling of larger charities was primarily because some charities are disproportionately found amongst this size range (for example medical research). In contrast, many smaller charities fall into a relatively small number of categories, such as places of worship, small grant makers, and community associations. Without oversampling larger charities, some important categories might be missed entirely.

We derived the UK-CAT iteratively from the sample, manually classifying each charity and updating a central list of classifications with newly derived tags. The three main coders met regularly to refine and clarify the categories. Because the UK-CAT was derived and assigned in parallel, this process inevitably involved some backwards revisions. Fortunately, most changes could be applied retrospectively relatively easily, such as changes to tag names, merging or removing tags.

Because of the iterative way that the UK-CAT was developed and assigned, as well as the fact that multiple tags could be applied, a stringent test of inter-coder reliability using the UK-CAT was not possible. We did, however, conduct a small, alternative test, with all three coders assigning a single ICNPTSO category to 100 charities. After retrospectively determining the 'correct' category by consensus, the individual coder agreement rates were 84, 82, and 78 per cent, similar to Litofcenko, Karner and Maier (2020), though using more textual data.

Once completed, the UK-CAT had 252 tags. The lower-level tags sit within a hierarchy of 24 top-level 'groups', such as health, education, and social welfare. In some cases, there are also mid-level sub-groups to help provide further structure[5].

Keyword based classifications

Manually applying classifications becomes less feasible as the number of charities increases. To automate the process, we had the option of either a 'dictionary method', using human-designed keyword matching, or machine learning methods, which generate these rules based on the statistical

[5] For details of the UK-CAT full classification scheme and the number of matches for every category, please see Appendix A and https://github.com/charity-classification/ukcat/blob/main/data/ukcat.csv .

properties of the data (Ma, 2020). In the first instance, we opted for the dictionary method. One of the main advantages is transparency, as it is both simple to understand and anyone can look up the keywords that led to a particular assignment.

The 4,200 manually classified entries provided a pool of baseline data from which to start developing the keywords, using code written in the Python programming language[6]. We initially used the training dataset to provide suggestions for relevant keywords, using frequencies on common words and pairs of words (bigrams), though many of the search terms were also derived from our own knowledge or online research.

To ensure an effective list of keywords, it is useful to use a technique to match plurals, different tenses, and variations in spelling. One option is 'word stemming', for example using the term 'famil' to match both 'family' and 'families'. This process can cause unintended consequences, however, if terms such as 'training' were shortened to 'train', which of course has an alternative meaning. Instead, we relied on regular expressions (regex), a way of specifying search terms flexibly using special characters. For example the regular expression "wom[ae]n'?s?" would match women, woman, woman's, or even a misspelling such as 'womans'.

After completing an initial set of search terms, we examined the results against the training dataset, paying close attention to any 'false negatives'. These were charities which we had matched manually to a tag, but which were not yet being matched by our search terms. This usually revealed necessary modifications. At the same time, we kept a close eye on those charities that were being included, particularly those not matched by the human coders. In many cases, these were reasonable, if slightly less central to the charity's main mission. In others, modification was again necessary.

Matching performance

Applying the search rules to the population dataset of 201,963 active charities resulted in 807,782 keyword matches across all 252 UK-CAT tags. Figure 1 shows the distribution of direct UK-CAT matches per charity. The mean and median number of matches was four. There is a modest amount of skew caused by some charities with many matches, but most have a relatively small number. It is difficult to avoid a minority of cases being 'over-tagged', as some community organisations provide long descriptions outlining many activities.

Four per cent of charities (7,419) have no matches at all. In some cases, this is due to the name or activities being written in Welsh, though in future we hope to explore the option of using an automatic translation service. Some charities may also have the potential for a positive match, with further improvement to the regular expressions. Finally, it may also be possible to reduce the number of zero matches by supplementing the results with other data.

An initial 'eyeball' examination suggested that the results were providing a reasonable summary of each charity's activities (see Table 2 for three examples).

Figure 1: Number of charities with different numbers of tag matches

[6] https://github.com/charity-classification

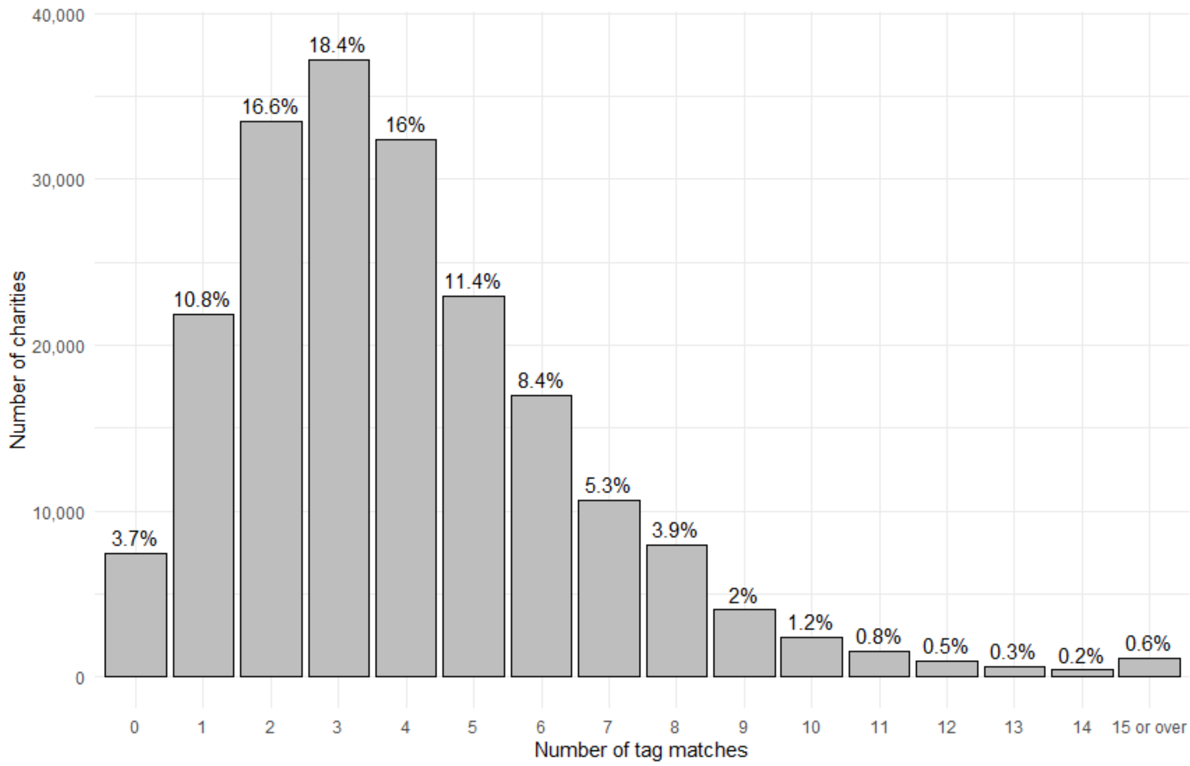

Table 2: Three example charities and matched UK-CAT classifications

Charity name	Activities	UK-CAT tags
Corporation of The High School of Dundee	"The advancement of education."	Education, Schools, Secondary education
Northampton Scottish Association Fund	"Providing charitable donation to local charities on an annual basis"	Associations, Charity and VCS support
Craven police charity fund	"Supporting local causes in raising money"	Fundraising, Emergency services

To conduct a more quantifiable examination, we created a new sample of 100 randomly selected charities and classified them manually using the UK-CAT. The research team allocated 135 tags in total, applying no more than two tags per charity (though no upper limit was imposed). The keyword rules matched 347 tags directly, about 2.5 times more than the research team. Perhaps unsurprisingly, the keyword matching process is therefore less parsimonious.

In the final results, additional group and sub-group tags are applied automatically, based on matches against lower-level tags. So, if 'Museum' is matched, the group in sits within, 'Heritage', is also matched automatically. For this test, however, only the tags applied directly using the keyword searching were considered, to avoid inflating agreement levels. This does, however, make the comparison quite strict.

Of the 135 tags applied by the research team, 105 were also matched using the keyword searching (78 per cent). Of the other 30, in three cases, the activities listed a high number of activities, but the human coders summarised the charity as a community association. In another two cases, a spelling or formatting error prevented the match. In 14 cases, the keyword did not pick up on a key phrase but could potentially do so with further modification of the keyword search rules.

Finally, eleven non-matches were more difficult to resolve, due to words having multiple meanings or human coders being able to infer additional context. For example, one charity listed only 'General Charitable Purposes' as its activities. Another refers to a 'people's park', which is a green space. The keyword 'park', however, cannot be included in the relevant regular expression, as it occurs too often in the names of schools and other organisations (as well as the occasional reference to a car park).

Of the 209 tags matched by the keyword searching, but not by the research team, these were not necessarily 'incorrect'. Some were near misses, hitting the group or subgroup, but not the exact same tag. In addition, some tags may be superfluous, but not incorrect. For example, we tagged one charity using 'playground', but in contrast to the keyword search, did not also include 'young children'. Manually reviewing all 347 keyword search matches, we found only fourteen (4.0 per cent) that we felt could be classified as 'false positives'. For example, a reference to a 'War Memorial Hospital' matching against 'Monuments; statues and memorials'. Overall, these results appear encouraging, even if they suggest that the keyword search terms still have room for further improvement.

Findings

A central contribution of this article has been to introduce a new classificatory system for UK charities, the UK-CAT, and apply it to charity registration data. To fully assess how successful this process has been, this section outlines how charities are distributed between the different categories, both by number and their overall income. We also briefly compare these results with the classifications from the register of charities in England and Wales.

Number of charities by tag group

All 252 UK-CAT classifications are listed in Appendix A. This table also includes the number of charities matched against each classification, using a combined September 2021 download of the English and Welsh, Scottish and Northern Irish charity registers. To recap, the hierarchy has three levels: 24 'groups, 17 'sub-groups' and 211 'lower-level tags'. If a charity matches the regular expression at a lower level (or sub-group), any higher levels are also automatically matched. But groups and sub-groups can also have their own independent matches.

It is not feasible to explore the results for all 252 tags within this article. Instead, we briefly examine some of the most popular 'groups' of tags identified. The number of charities matched against each group is shown in Figure 2. The groups 'beneficiary group', 'facilities' and 'charitable activities' are not included, as they are not as substantively interesting when aggregated to the group level.

Some caution is needed, however, as some groups contain more lower-level tags and broader keywords than others. Lots of charities will be matched to 'Education', for example, because they mention educating the general public, rather than because they are directly involved with schools (lower-level tags are needed to isolate these cases). The results, therefore, partly reflect definitional and methodological choices. Nevertheless, they help to paint a useful picture of UK charitable activity and causes. They also help to indicate the type of language commonly used within charities' names and activities.

'Education' is the most matched group (**42 per cent** of charities). The common use of educational terms and language reflects the long history of the UK voluntary sector's involvement in this field (Harris, 2010), which continues today through PTAs, private (fee-charging) schools, and adult training. The individual tag 'Schools' is the most common lower-level tag, from any group, matched by 16 per cent of charities.

Similarly, there is a longstanding association between religion and many charities in the UK (McCabe et al. 2016). **22 per cent** of charities match at least one tag within the 'Religion' group. 'Christianity' is the most common specific religion tagged (16 per cent of charities), with Islam second (1.2 per cent).

The keywords for some of the 'Associations' group (**24 per cent** of charities) are very broad, including 'association' and 'club'. This does, however, highlight the associational language used by many charities, perhaps reflecting the membership-based roots of many voluntary organisations (Billis, 2010), and the fact that most charities are relatively small, local, and run by volunteers.

Tags from the 'Social welfare' group were matched by **18 per cent** of charities, with 11 per cent matching the tag for 'Individual poverty', another cornerstone of charitable activity in the UK and perhaps the one most associated with the idea of 'charity'.

Tags in the arts groups are also matched relatively frequently (**15 per cent** of charities), as is the 'Leisure' group (**14 per cent** of charities). This is worth highlighting, as arguably the connection between charity and leisure is sometimes neglected in public discourse, with a greater focus on charities involved in welfare service-based activities (Rochester, 2013).

The language used in charities' names and activities, therefore, clearly points towards education, association, religion, poverty, and leisure. Language associated with economic development, housing, homelessness, unemployment, or other forms of specific welfare services for marginalised groups, is less common. As context, it is worth noting that as of 2010, just over half of all third sector organisations had no staff (IPSOS, 2010), and 44 per cent of general charities currently have an income of less than £10,000 (NCVO, 2021).

Despite their prominence in current policy debates, it is also notable that only two per cent of charities refer explicitly to the 'Environment', only 0.1 per cent to the lower-level tag 'Climate Emergency', 0.1 per cent to 'Racial justice', and 0.8 per cent to various forms of 'Abuse'. One reason may be that because charities' activities and objects are rarely updated, the picture painted here is likely to reflect the historical development of the charity sector. Climate change focussed charities are likely to be a relatively recent phenomenon. It is also perhaps worth noting that issues that are quite prominent within the academic and public discourse, such as 'social investment' (0.01 per cent), or social entrepreneurship (0.2 per cent), are also very rarely mentioned.

Figure 2: number of charities matched against each UK-CAT group (not mutually exclusive).

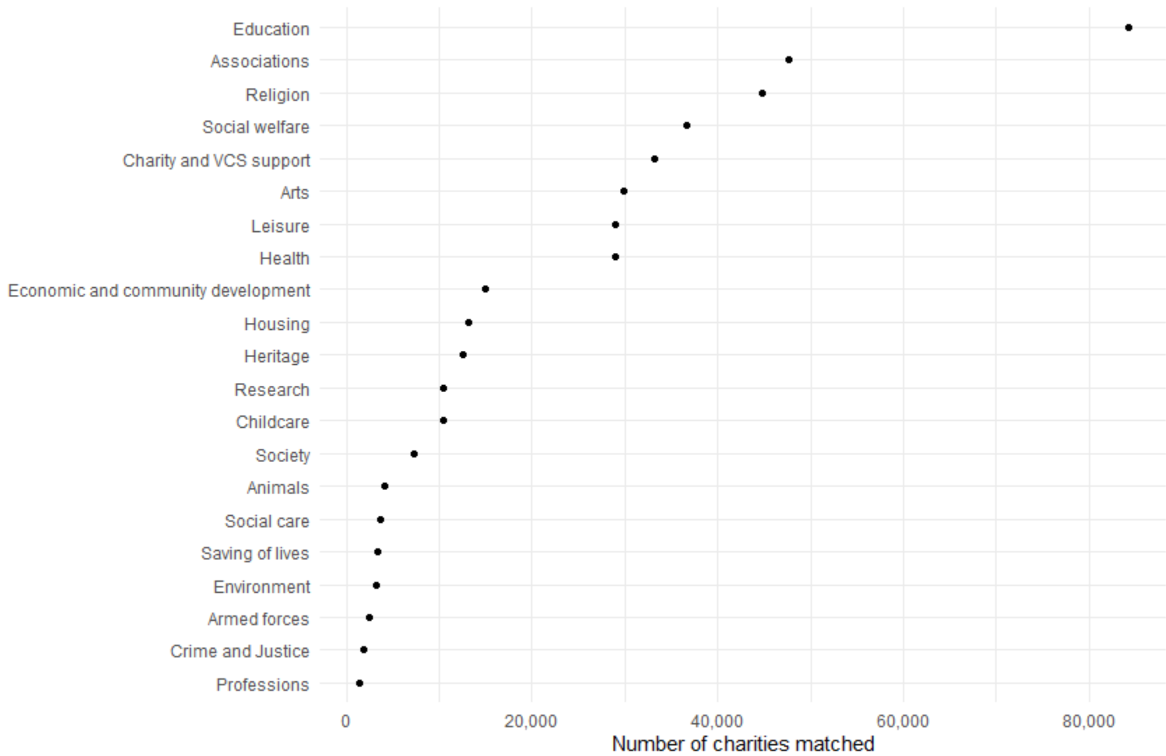

Income by tag group

The number of charities, of course, provides only one way of measuring charitable activity. Others include staff, volunteer and membership numbers, different measures of income or expenditure, assets, or output measures (Pennerstorfer and Rutherford, 2019). While there is not space to explore all of these, it is worth looking briefly at how income levels are distributed between groups, to see whether the overall picture changes.

The aggregate income, for all charities matched against each group, is shown in Figure 3. These groups are not mutually exclusive, so the same income may be counted in more than one category. The pattern is in many respects similar to the number of charity matches. Groups matching a higher number of charities, unsurprisingly, tend to have a higher aggregate level of income. Some groups, however, such as 'Health' and 'Research' rank higher in terms of income than they do for the number of charities matched. This may partly reflect the flow of large amounts of funding into large medical research charities such as Cancer Research UK (income of £656,107,415 in 2020), Welcome Trust (£463,593,136), and several others found within the largest 20 charities (NCVO, 2018). On the other hand, groups which match against a large number of smaller charities, such as 'associations' are lower on the income ranking than their number of charities would suggest.

Figure 3: total income of charities matched against each UK-CAT group (not mutually exclusive).

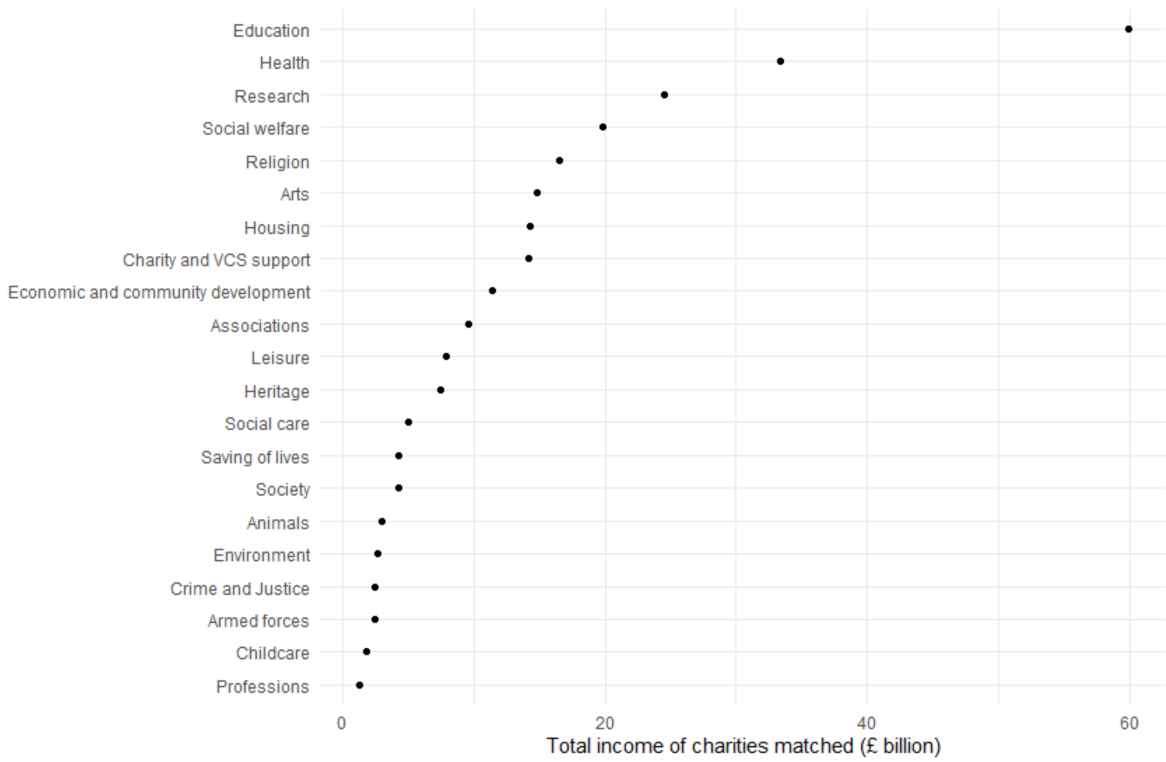

How the UK-CAT compares to existing classification systems

As Table 1 showed, the most obvious difference compared to existing schemes is that the UK-CAT contains a much larger number of categories compared to the alternatives. As already discussed, this was a conscious decision to try and address 'gaps' within the other schemes, and to include causes which were previously subsumed into broad 'general' or 'other' categories. The idea of gaps is perhaps misleading, however, as no classification scheme can ever be comprehensive. As with all choices relating to the schema design, the level of detail is a choice, showing one possible version of the social world.

Nevertheless, compared to the various national registers, the greater number of categories in the UK-CAT does provide the opportunity to explore some previously excluded categories. This can be demonstrated by a more in-depth comparison with the register of charities for England and Wales, the largest of the three national charity registers. This is arguably more comparable to the UK-CAT than the ICNPTSO, which is designed for international comparisons.

The UK-CAT includes tags such as 'food banks' (935 matches[7]), 'domestic abuse' (528), 'palliative care' (360), 'bereavement' (799), 'mental health' (2,536), 'addiction and dependency' (769), or the 'climate emergency' (187). All of these relate to important areas of social policy and voluntary sector activity, but at the time of writing none of these categories are included in the English and Welsh register of charities. It is a substantive finding to know how many charities are working in these fields, amongst others, which was not possible using previous schemes. A full mapping of the UK-CAT against the categories covered in the English and Welsh register of charities, and the ICNPTSO, is available on the project website[8].

Working in the opposite direction, most of the categories included in the current version of the English and Welsh register are replicable using one or more tags (or tag groups) from the UK-CAT. For the purposes of space, Table 3 focuses just on the 17 options for 'what does your charity do?' in the register of charities. All except 'general charitable purposes' and 'other charitable purposes' can be replicated using the UK-CAT.

In terms of the number of charities per tag, some of the categories match up reasonably closely. For example, roughly the same number of charities self-select religion under both systems. This may reflect a relative lack of ambiguity around the concepts and key words involved. Charities that mention 'church' or 'the gospel' are matched under UK-CAT, but are also highly likely to self-identify as a religious organisation. Overall, it is worth noting that in proportional terms the two schemes do not appear radically different, where comparable categories are available. Education, for example, is the most frequently matched in both, while Animals, or the Armed Forces, are relatively infrequent. This would suggest that while the UK-CAT expands our picture of the charity sector, it does not necessarily upend our existing understanding.

In a few cases, such as 'Armed forces' / emergency services efficiency' and 'recreation', the number of UK-CAT matches is lower than the number self-selecting in their annual returns. In the case of recreation, this may be partly because the charity register categories and the UK-CAT do not match up particularly well in this case. The 'Sport' tag is a lower-level tag within the UK-CAT 'leisure' category, whereas it remains a separate category in the English and Welsh charity register.

In the case of the armed forces and emergency services, we explored in more depth some of the charities matched against the UK-CAT, but not within the English and Welsh charity register. Outside of a very small number of false positives, most of these charities mentioned the armed forces very

[7] Charities registered in England and Wales only
[8] https://charityclassification.org.uk/

explicitly. For example, 'The McCandless ex-service men memorial homes' provides 'homes for ex-servicemen'. Within its annual return, however, it has self-selected only 'Accommodation/housing'. There were many similar examples of charities limiting their selections in this way, potentially obscuring other important aspects of their work.

For most groups, however, more charities have self-selected than have been allocated by the UK-CAT keyword search rules. For example, over 2.5 times as many charities selected 'Disability' as the UK-CAT would suggest. In some cases, there is nothing at all in the name or activities descriptions of these charities to suggest a link to disability. In some cases, they have quite general descriptions. For example, one describes simply 'supporting local charities', another 'Religious, Educational and Charitable Activities'. This highlights how keyword searching struggles when charities provide relatively sparse details in their textual documentation. Some charities refer more vaguely to relieving 'poverty and sickness', which is used as something of a catch-all category.

Finally, some charities appear to work with people with long term health conditions, elderly people, or armed forces veterans. In these cases, disability may indeed feature quite heavily in their work, but it is not explicitly mentioned directly in their activity descriptions. This perhaps points towards both the advantage and disadvantage of self-selection. Charities are best placed to judge the extent to which their work is relevant to a particular category, but we have no means of verifying their reasons for a selection, or the degree of importance to their work. Again, it is worth stressing that there is no 'correct' number of charities to allocate, with much subjectivity inherent in the process. An advantage of the UK-CAT system, however, is that because our keywords are publicly available, it is possible to see exactly *why* a category is included or not.

Table 3: comparison of categories under 'what does your charity do?' in the English and Welsh register or charities, versus their UK-CAT equivalents

Register of charity classification	Charities self-selected	Relevant UK-CAT classifications	Charities tagged
General charitable purposes	56,426	N/A	-
Education / training	86,550	Education (20 tags)	68,725
The advancement of health or saving of lives	28,782	Health (33 tags) OR Saving of lives (contains humanitarian relief) (4 tags)	23,365
Disability	25,293	People with disabilities (1 tag) OR People with learning disabilities (1 tag)	9,987
The prevention or relief of poverty	33,573	Individual poverty (1 tag)	17,203
Overseas aid / famine relief	10,421	Humanitarian relief (1 tag)	1,527
Accommodation/ housing	8,439	Housing (6 tags)	10,887
Religious activities	35,855	Religion (22 tags)	36,543
Arts / culture / heritage / science	30,282	Arts (20 tags) OR Heritage (7 tags) OR Research (4 tags)	35,188

Amateur sport	27,247	Sports (1 tag)	10,265
Animals	4,454	Animals (5 tags)	3,358
Environment / conservation / heritage	19,087	Environment (5 tags)	2,256
Economic / community development /employment	21,768	Economic and community development (12 tags)	9,804
Armed Forces / emergency service efficiency	1,065	Armed forces (5 tags) OR Emergency Services (1 tag)	2,750
Human rights / religious or racial harmony/equality or diversity	6,244	Society (8 tags)	2,464
Recreation	16,032	Leisure (includes sports) (7 tags)	21,628
Other charitable purposes	13,320	N/A	-

Discussion

This project has established a new, detailed way to classify the activities of charities in the UK and to automatically apply this classification using keyword matching. As outlined in the methodology, all the code and classification results are available online via GitHub for others to use or replicate. We can take away several key reflections, both substantive and methodological.

First, the results are a testament to the great variety amongst UK charities (Kendal and Knapp, 1995). The fact that we arrived at over 250 tags, despites efforts to rationalise, points to the sheer number of different causes. Most charities also match more than one tag, reflecting their multi-purpose nature. Despite this diversity, many charities use the language of education, association, religion, and poverty when describing their activities. Whilst this may be partly due to the choice of keywords, the results provide a reminder that many charities continue to be small scale, associational, and operate relatively far from state or market influence (Rochester, 2013).

We have shown that the UK-CAT is considerably more detailed than main alternative classification schemes used within the UK, providing an insight into areas of charitable activity that were previously not as visible. In some cases, the keyword search method is matching fewer charities than when charities can self-select similar categories themselves from the English and Welsh register of charities. On the other hand, whilst there is no 'correct' number of matches, the UK-CAT does have considerable advantages in terms of transparency.

It is also clear from the results that there is no single, overarching rationale, or guiding hand, determining which causes attract the most charities. The voluntary sector is an aggregation of many different decisions to volunteer, donate, purchase or commission, at different times, for different reasons. As such, its distribution provides a fascinating window into our collective concerns and interests, but it does not operate according to any obvious central theory of need. It is not hard to spot issues, such as climate change, where the number of matched charities is relatively low compared to the size of the challenge (see Salamon, 1987).

From a methodological perspective, we have shown that the textual fields in the various registers of charities are a rich source of data on charitable activities and causes. This data does, however, present some challenges. Importantly, charitable activities and objects data is rarely updated. This may partly explain why many of the most common groups of tags would have looked equally at home a hundred years ago, while more recent areas of public concern receive fewer matches. Further analysis might usefully explore which categories the newest charities match against, to gain a sense of how the picture presented may change over time. In addition, the absence of alternative organisational forms, such as Community Interest Companies, may partly explain the lack of more market-based language, and a focus on registered charities may also miss more radical, informal groups.

We have shown that the keyword matching approach is able to perform the task of automatic classification relatively well, with few obvious false positives and a respectable rate of false negatives (if we assume that the manually coded classifications are 'correct'). It also has the great advantage of being relatively transparent, replicable, and easy to understand. This method too, however, has its limitations. Much of the contextual detail, which humans process to interpret text, is missed. It is possible, that building on this foundation, and the training dataset created, more advanced machine learning models may be able to improve upon the results. In addition, machine learning models may be able to apply a ‘relevance’ score for each tag against every charity. This may provide the best of all worlds, allowing us to apply multiple tags, whilst also ranking their relevance to pick a single ‘primary’ tag. It should be noted, however, that our own training dataset is relatively small compared to some machine learning based projects (Ma, 2020), and that machine learning can bring its own costs in terms of transparency and flexibility.

Our hope is that this project has provided a means to view the charity sector in a new light, helping researchers and their audiences to better understand what thousands of charities do day-in, day-out across the UK. As such, we also hope that the UK-CAT is only the start of many further improvements to charity classifications and charity data in years to come.

References

Appe, S. (2012) ‘What about who is mapping and its implications? Comments on Brent Never’s ‘‘The case for better maps of social service provision.’’’ *Voluntas*, **23** (1), pp.204–212.

Barman, E. (2013) ‘Classificatory Struggles in the Nonprofit Sector: The Formation of the National Taxonomy of Exempt Entities, 1969—1987’, *Social Science History*, **37** (1), pp.103-141.

Bowker, G.C. and Star, S. (1999) *Sorting Things Out: Classification and Its Consequences*, MIT press: London.

Carmel, E. and Harlock, J. (2008) ‘Instituting the “third sector” as a governable terrain: partnership, procurement and performance in the UK’, *Policy & Politics*, **36** (2), pp.155–71.

DiMaggio, P. J. (1983), ‘State Expansion and Organizational Fields’ *in* R. Hall and R. E. Quinn (eds.), *Organization Theory and Public Policy*, Sage, Beverly Hills, pp.147-161.

Eikenberry, A. and Kluver, J.D. (2004), ‘The Marketization of the Nonprofit Sector: Civil Society At Risk?’, *Public Administration Review*, **64** (2), pp.132-140

Fligstein, N. and McAdam, D. (2011) ‘Toward a General Theory of Strategic Action Fields’, *Sociological Theory*, **29** (1), pp.1-26

Fyall, R., Moore, M. K. and Gugerty, M. K. (2018) ‘Beyond NTEE codes: Opportunities to understand nonprofit activity through mission statement content coding’, *Nonprofit and Voluntary Sector Quarterly*, **47** (4), pp.677–701.

Kendall, J. and Knapp, M. (1995) ‘A loose and baggy monster: boundaries, definitions and types’, In: J. Davis Smith, C. Rochester and R. Hedley, *An Introduction to the Voluntary Sector*, Routledge: London, pp. 66-95.

Kendall (2003) *The Voluntary Sector: Comparative Perspectives in the UK*, London: Routledge

Kenley, A. and Wilding, K. (2021), *Better data, bigger impact A review of social sector data,* October, Pro Bono Economics: London.

Lampkin, L., Romeo, S. and Finnin, E. (2001) ‘Introducing the nonprofit program classification system: The taxonomy we’ve been waiting for’, *Nonprofit and Voluntary Sector Quarterly*, **30** (4), pp. 781–793.

Lantz, B. (2015) *Machine Learning with R,* Packt Publishing: Birmingham.

Leung, R. (2020) *Charities speak Mapping arts and culture charities in England and Wales using data science*, March, Creative Industries Policy and Evidence Centre, Nesta: London.

Litofcenko, J., Karner, D. and Maier, F. (2020) ‘Methods for classifying nonprofit organizations according to their field of activity: A report on semi-automated methods based on text’, *Voluntas*, **31** (1), pp.227-237.

National Council for Voluntary Organisations (NCVO) (2012) *The size and scope of the voluntary addiction sector in England and Wales*, July, Action on Addiction: London.

NCVO (2021) *UK Civil Society Almanac 2021*, NCVO Almanac Series, NCVO: London, Available at: https://beta.ncvo.org.uk/ncvo-publications/uk-civil-society-almanac-2021/

NCVO (2018) *Britain's biggest charities: key features*, NCVO research briefing, November, NCVO: London.

Nickel, P.M. and Eikenberry, A.M. (2016) 'Knowing and Governing: The Mapping of the Nonprofit and Voluntary Sector as Statecraft', *Voluntas*, **27** (1), pp.392–408.

Newbigging, K., Mohan, J., Rees, J., Harlock, J. and Davis, A. (2017) 'Contribution of the voluntary sector to mental health crisis care in England: protocol for a multimethod study', *BMJ Open*, **7**, pp.1-9.

Our Community (2020), *CLASSIE: A classification system for social sector initiatives and entities* [online], available at: https://www.ourcommunity.com.au/classie

Pennerstorfer, A. and Rutherford, A.C. (2019), 'Measuring Growth of the Nonprofit Sector: The Choice of Indicator Matters', *Nonprofit and Voluntary Sector Quarterly,* **48** (2), pp.440–456

Ma, J. (2021) 'Automated Coding Using Machine Learning and Remapping the U.S Nonprofit Sector: A Guide and Benchmark', *Nonprofit and Voluntary Sector Quarterly*, **50** (3), pp.1-26.

Rochester, C. (2013) *Rediscovering Voluntary Action: The Beat of a Different Drum,* Palgrave Macmillan: London.

Salamon, L.M. and Anheier, H. (1992a) 'In search of the nonprofit sector. I: The question of definitions', *Voluntas*, **3** (1), pp.125-51.

Salamon, L.M. and Anheier, H.K. (1992b) 'In search of the non-profit sector II: The problem of classification'. *Voluntas*, **3** (2), pp.267–309.

Salamon, L.M. (1987) 'Of Market Failure, Voluntary Failure, and Third-Party Government: Toward a Theory of Government-Nonprofit Relations in the Modern Welfare State', *Nonprofit and Voluntary Sector Quarterly,* **16** (1), pp.29–49.

Salamon, L.M. (2010) 'Putting the civil society sector on the economic map of the world', *Annals of Public and Cooperative Economics*, **81** (2), pp.167-210.

Python libraries:

Van Rossum, G. & Drake, F.L. (2009), *Python 3 Reference Manual*, Scotts Valley, CA: CreateSpace.

The pandas development team (2022), *pandas-dev/pandas: Pandas (v1.5.0),* Zenodo, Available at: https://doi.org/10.5281/zenodo.7093122

Harris, C.R., Millman, K.J., van der Walt, S.J. et al. (2020), Array programming with NumPy, *Nature,* **585**, pp.357–362.

Pedregosa et al., (2011), Scikit-learn: Machine Learning in Python, *JMLR,* **12** (18), pp.2825-2830, http://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html

Bird, Steven, Edward Loper and Ewan Klein (2009), *Natural Language Processing with Python,* O'Reilly Media Inc: Sebastopol.

Appendix A

Group	Sub-group	UK-CAT tag	Matching UK charities
Armed forces		Armed forces	2,017
Armed forces		Army	474
Armed forces		Navy	409
Armed forces		RAF	380
Armed forces		Veterans	207
Animals		Animals	3,358
Animals		Cats	359
Animals		Dogs	792
Animals		Donkeys	53
Animals		Horses	549
Arts		Arts	24,053
Arts		Festival	2,030
Arts		Languages	422
Arts		Visual arts	1,761
Arts	Media and publishing	Media and publishing	7,810
Arts	Media and publishing	Film	896
Arts	Media and publishing	Literature	6,282
Arts	Media and publishing	Media	306
Arts	Media and publishing	Print media	498
Arts	Media and publishing	Radio	338
Arts	Media and publishing	Television	273
Arts	Performing art	Performing art	11,882
Arts	Performing art	Choirs	2,287

Arts	Performing art	Dance	2,109
Arts	Performing art	Music	7,499
Arts	Performing art	Musical theatre	238
Arts	Performing art	Opera	563
Arts	Performing art	Orchestra	633
Arts	Performing art	Theatre	3,429
Associations		Associations	39,157
Associations		Community association	2,952
Associations		Fraternal societies	635
Associations		Inner Wheel	293
Associations		League of Friends	7,718
Associations		Social club	1,752
Associations		Townswomen's Guild	214
Associations		Women's Institute	2,129
Associations		YWCA / YMCA	158
Associations	Service clubs	Service clubs	2,167
Associations	Service clubs	Lions club	600
Associations	Service clubs	Rotary club	1,262
Associations	Youth Groups	Youth Groups	6,229
Associations	Youth Groups	Cadets	453
Associations	Youth Groups	Girlguiding	517
Associations	Youth Groups	Scouting	3,779
Beneficiary group		Beneficiary group	63,990
Beneficiary group		Asylum seekers and refugees	1,073
Beneficiary group		Children	22,858

Beneficiary group		Families	9,461
Beneficiary group		Girls	2,240
Beneficiary group		LGBTQ+	234
Beneficiary group		Men	954
Beneficiary group		Migrants	444
Beneficiary group		Older people	7,067
Beneficiary group		Parents and guardians	8,620
Beneficiary group		People with learning disabilities	1,850
Beneficiary group		Racial; ethnic or national communities	433
Beneficiary group		Widows; widowers and orphans	1,787
Beneficiary group		Women	7,304
Beneficiary group		Young children	7,705
Beneficiary group		Young people	17,701
Beneficiary group	People with disabilities	People with disabilities	9,956
Beneficiary group	People with disabilities	Riding for the disabled	328
Charitable activities		Charitable activities	24,811
Charitable activities		Advice and individual advocacy	2,739
Charitable activities		Charity shops	761
Charitable activities		Policy campaigning and advocacy	4,694
Charitable activities		Social Investment	18
Charitable activities	Grant making	Grant making	17,405
Charitable activities	Grant making	Grants to individuals	1,132
Charitable activities	Grant making	Grants to organisations	8,240
Childcare		Childcare	8,821
Childcare		Nursery	6,145

Childcare		Out of school club	1,235
Childcare		Playground	788
Childcare		Playgroup	1,880
Crime and Justice		Crime and Justice	1,508
Crime and Justice		Offender support and rehabilitation	663
Crime and Justice		Prevention and safety	125
Crime and Justice		Road safety	219
Crime and Justice		Trafficking and modern slavery	135
Crime and Justice		Victim support	82
Charity and VCS support		Charity and VCS support	28,712
Charity and VCS support		Financial investment	51
Charity and VCS support		Fundraising	15,268
Charity and VCS support		Umbrella bodies	582
Charity and VCS support		Volunteering	4,079
Economic and community development		Economic and community development	9,804
Economic and community development		Community development	2,230
Economic and community development		Economic development	519
Economic and community development		International development	908
Economic and community development		Planning and architecture	143
Economic and community development		Rural and farming areas	2,749
Economic and community development		Social enterprise	283

Economic and community development		Unemployment	3,108
Economic and community development		Urban areas	428
Economic and community development	Infrastructure	Infrastructure	368
Economic and community development	Infrastructure	Energy	106
Economic and community development	Infrastructure	Water	272
Education		Education	68,752
Education		Adult education	295
Education		Further education	2,356
Education		Higher education	1,859
Education		Primary education	4,947
Education		Student support	1,935
Education		Schools	28,164
Education		Secondary education	1,325
Education		Student union	126
Education		University of the Third Age	788
Education	School support	School support	10,283
Education	School support	Parent teacher	5,029
Education	School support	School fundraising	4,658
Education	Training	Training	16,050
Education	Training	Basic skills	108
Education	Training	Employability training	3,271
Education	Training	ESOL	235
Education	Training	IT and digital	257

Education	Training	Mentoring	1,005
Education	Training	Vocational training	335
Environment		Environment	2,256
Environment		Climate Emergency	187
Environment		Conservation and sustainability	1,188
Environment		Recycling	269
Environment		Wildlife	906
Facilities		Facilities	17,032
Facilities		Cemetery	544
Facilities		Community cafe	715
Facilities		Community centre	2,867
Facilities		Green space	2,849
Facilities		Open spaces	935
Facilities		Playing fields	2,537
Facilities		Village hall	8,179
Facilities		Youth centre	217
Health		Health	21,659
Health	Health condition	Health condition	8,022
Health	Health condition	Addiction and dependency	769
Health	Health condition	Cancer	1,525
Health	Health condition	Cerebral palsy	74
Health	Health condition	Chronic Fatigue Syndrome	25
Health	Health condition	Dementia	466
Health	Health condition	Fibromyalgia	7
Health	Health condition	Hearing loss	502

Health	Health condition	HIV / Aids	346
Health	Health condition	Maternity	394
Health	Health condition	Mental health	2,536
Health	Health condition	Motor Neurone Disease	23
Health	Health condition	Multiple Sclerosis	112
Health	Health condition	Sickle Cell	41
Health	Health condition	Strokes	213
Health	Health condition	Visual impairment	969
Health	Health services	Health services	5,687
Health	Health services	Alternative medicine	42
Health	Health services	Ambulance service	318
Health	Health services	Complementary therapies	88
Health	Health services	Counselling and therapy	2,549
Health	Health services	Health and wellbeing	1,238
Health	Health services	Nursing	687
Health	Health services	Palliative care	360
Health	Health services	Physiotherapy	109
Health	Health services	Surgery	581
Health	Healthcare provider	Healthcare provider	3,890
Health	Healthcare provider	Hospice	546
Health	Healthcare provider	Hospital	2,801
Health	Healthcare provider support	Healthcare provider support	762
Health	Healthcare provider support	Friends of healthcare provider	558
Housing		Housing	10,887
Housing		Accommodation	6,735

Housing		Almshouse	1,291
Housing		Homelessness	1,975
Housing		Housing association	75
Housing		Temporary or emergency housing	448
Heritage		Heritage	8,720
Heritage		Archaeology	376
Heritage		Historical conservation and restoration	2,618
Heritage		History	3,236
Heritage		Monuments; statues and memorials	532
Heritage		Museum	1,988
Heritage		Natural history	178
Leisure		Leisure	21,628
Leisure		Exercise and fitness	2,000
Leisure		Gardening	916
Leisure		Hobbies	337
Leisure		Outdoor pursuits	897
Leisure		Recreation	12,402
Leisure		Sports	10,265
Professions		Professions	1,129
Professions		Clergy	496
Professions		Emergency service workers	49
Professions		Healthcare workers	304
Professions		Miners	284
Religion		Religion	36,543
Religion		Baha'i	72

Religion		Buddhism	321
Religion		Hinduism	349
Religion		Islam	2,264
Religion		Jainism	35
Religion		Judaism	1,667
Religion		Sikhism	262
Religion		Spiritualism	241
Religion	Christianity	Christianity	27,432
Religion	Christianity	Church of England	2,240
Religion	Christianity	Church of Ireland	1
Religion	Christianity	Church of Scotland	3
Religion	Christianity	Jehovah's Witnesses	1,253
Religion	Christianity	Roman Catholic	420
Religion	Christianity	Society of Friends (Quakers)	180
Religion	Religious activities	Religious activities	24,552
Religion	Religious activities	Chaplaincy	212
Religion	Religious activities	Church or place of worship	18,464
Religion	Religious activities	Parochial Church Council	2,937
Religion	Religious activities	Religious education	1,761
Religion	Religious activities	Religious ministry	6,177
Research		Research	7,447
Research		Medical research	749
Research		Philosophy	254
Research		Science	2,363
Social care		Social care	3,132

Social care		Adult day care	678
Social care		Carer support	257
Social care		Children in care	120
Social care		Children's homes	567
Social care		Domiciliary care	179
Social care		Residential care	714
Social care		Residential care with nursing	43
Social care		Respite	558
Saving of lives		Saving of lives	2,683
Saving of lives		Emergency services	755
Saving of lives		Humanitarian relief	1,527
Saving of lives		Search and rescue	349
Society		Society	2,464
Society		Citizenship	584
Society		Conflict resolution	71
Society		Democracy	43
Society		Equality and diversity	508
Society		Human rights	421
Society		Racial justice	53
Society		Religious; racial or cross-border harmony	592
Social welfare		Social welfare	29,182
Social welfare		Benevolent Society	1,486
Social welfare		Bereavement	799
Social welfare		Clothes	1,134
Social welfare		Community transport	697

Social welfare		Individual poverty	17,203
Social welfare		Loneliness	1,918
Social welfare		Social activities	4,424
Social welfare	Abuse	Abuse	1,303
Social welfare	Abuse	Child abuse	73
Social welfare	Abuse	Domestic abuse	528
Social welfare	Abuse	Refuge or shelter	206
Social welfare	Abuse	Sexual abuse	229
Social welfare	Food	Food	4,875
Social welfare	Food	Food banks	935