

Models of everywhere revisited: a technological perspective

BLAIR, Gordon, BEVEN, Keith, LAMB, Rob, BASSETT, Richard, CAUWENBERGHS, Kris, HANKIN, Barry, DEAN, Graham, HUNTER, Neil, EDWARD, Liz, NUNDLLOL, Vatsala, SAMREEN, Faiza http://orcid.org/0000-0002-9522-0713, SIMM, Will and TOWE, Ross

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

https://shura.shu.ac.uk/26223/

This document is the Published Version [VoR]

Citation:

BLAIR, Gordon, BEVEN, Keith, LAMB, Rob, BASSETT, Richard, CAUWENBERGHS, Kris, HANKIN, Barry, DEAN, Graham, HUNTER, Neil, EDWARD, Liz, NUNDLLOL, Vatsala, SAMREEN, Faiza, SIMM, Will and TOWE, Ross (2019). Models of everywhere revisited: a technological perspective. Environmental Modelling and Software, 122, p. 104521. [Article]

Copyright and re-use policy

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

Contents lists available at ScienceDirect



Environmental Modelling and Software

journal homepage: http://www.elsevier.com/locate/envsoft



Models of everywhere revisited: A technological perspective

Gordon S. Blair^{a,*}, Keith Beven^b, Rob Lamb^{c,b}, Richard Bassett^a, Kris Cauwenberghs^d, Barry Hankin^{e,b}, Graham Dean^a, Neil Hunter^e, Liz Edwards^a, Vatsala Nundloll^a, Faiza Samreen^a, Will Simm^a, Ross Towe^a

^a Ensemble Team, Lancaster University, UK

^b Lancaster Environment Centre, Lancaster, UK

^c JBA Trust, UK

^d Flanders Environment Agency (VMM), Belgium

^e JBA Consulting, UK

ARTICLE INFO

Keywords: Environmental modelling Models of everywhere Grid computing Cloud computing Data science The concept 'models of everywhere' was first introduced in the mid 2000s as a means of reasoning about the environmental science of a place, changing the nature of the underlying modelling process, from one in which general model structures are used to one in which modelling becomes a learning process about specific places, in particular capturing the idiosyncrasies of that place. At one level, this is a straightforward concept, but at another it is a rich multi-dimensional conceptual framework involving the following key dimensions: models of everywhere, models of everything and models at all times, being constantly re-evaluated against the most current evidence. This is a compelling approach with the potential to deal with epistemic uncertainties and non-linearities. However, the approach has, as yet, not been fully utilised or explored. This paper examines the concept of models of everywhere in the light of recent advances in technology. The paper argues that, when first proposed, technology was a limiting factor but now, with advances in areas such as Internet of Things, cloud computing and data analytics, many of the barriers have been alleviated. Consequently, it is timely to look again at the concept of models of everywhere in practical conditions as part of a trans-disciplinary effort to tackle the remaining research questions. The paper concludes by identifying the key elements of a research agenda that should underpin such experimentation and deployment.

1. Introduction

The concept of 'models of everywhere' was introduced by Beven in 2007 (Beven, 2007), and revised in a follow up paper (Beven and Alcock, 2012). The concept is fundamentally about having a stronger association between a given environmental model and the *place* that it represents. In the 2012 paper, they argue that it is "useful, and even necessary, to think in terms of models of everywhere ... [and this] ... will change the nature of the modelling process, from one in which general model structures are used in particular catchment applications to one in which modelling becomes a learning process about places". The 'necessity' stems from the need to constrain uncertainty in the modelling process in order to support policy setting and decision-making, particularly around water management (e.g. flooding and water quality), although the principles can also potentially apply to other areas of environmental modelling and management. (In the rest of the paper, we will tend to use examples and illustrations from hydrology and flood modelling although we stress this

potential generality of the approach.)

This is a compelling argument, and a reaction against the view that there can be generic environmental models capable of representing processes and behaviours across multiple places and indeed across multiple scales. Such general models are "expensive to develop, difficult to maintain and to apply because of their data demands and need for parameter estimation or calibration" (Beven, 2007). They also have problems in dealing with local epistemic uncertainties and non-stationarities, for example caused by change in local characteristics and climate drivers (e.g. Prudhomme et al., 2010; Beven, 2009, 2016).

Some examples of models of everywhere have been deployed but for relatively constrained applications and at specific scales. However, the concept has not been developed to the extent that the authors envisaged, where everywhere is represented across all scales in a coherent way. This paper re-examines the concept of models of everywhere from a technological perspective arguing that, at the time, the underlying technology was not sufficiently advanced to support the concept. Now,

https://doi.org/10.1016/j.envsoft.2019.104521

Received 27 April 2018; Received in revised form 6 August 2019; Accepted 24 September 2019 Available online 27 September 2019

1364-8152/© 2019 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

^{*} Corresponding author. E-mail address: g.blair@lancaster.ac.uk (G.S. Blair).

however, this has changed, with significant developments in areas such as data acquisition techniques, data storage and processing technologies, and data analytics capabilities, alongside a move towards a more open science supported by these developments.

Note that in developing ideas of models of everywhere we mean something quite different to the "hyperresolution" models that are starting to be used in Earth Systems Science (e.g. Wood et al., 2011; Beven and Cloke, 2012; Bierkens et al., 2015; Gilbert and Maxwell, 2017). With resolutions of the order of 1 km, the latter do not (as yet) provide simulations and visualisations at scales that local stakeholders can relate to directly (see the discussion of Beven et al., 2015). This is a critical aspect of how the models of everywhere concept has the potential to change the way that modelling is done. Both approaches do, however, focus attention on a requirement for scale dependent parameterisations that has proven difficult to resolve (e.g. McDonnell and Beven, 2014).

The overall aim of the paper is to determine the *current feasibility of* models of everywhere, particularly in the area of hydrological modelling, given the state-of-the-art in underlying technology. This breaks down into the following objectives:

- To carry out a detailed examination of the concept of models of everywhere to determine key underlying technological requirements;
- 2. To compare the state-of-the-art in technology in the period 2007–2012 and the present day to evaluate whether the time is now right for a widespread deployment of models of everywhere;
- 3. To provide a research roadmap to support such deployment in terms of outstanding research questions and challenges.

Note that there are other issues related to models of everywhere that also should be addressed, most notably human and societal issues. Such issues include the need to move towards open science and open data, and the role of communities in improving models in representing local places. These are alluded to in the paper but a full treatment of this important dimension is beyond the scope of the paper. We elect instead to focus on technological readiness.

The work is being carried out in the context of a significant reevaluation of approaches to flood modelling and associated risk management. For example, the UK Government's National Flood Resilience Review, published in September 2016¹ included important recommendations around improvements to long-term modelling capabilities. The review also encouraged the use of natural flood mitigation methods or "working with natural processes". This concept involves the use of many distributed in-channel and off-channel storage features, coupled with changes of land use to try to retain more flood runoff in catchment headwaters, or at least slow down its arrival to areas at risk of flooding (see Dadson et al., 2017). There are many current projects in the UK that are implementing natural flood management measures. Very few, however, have been associated with detailed monitoring of changes to the flow, or the operation of individual mitigation measures. Additionally, there are issues about whether the strategy will be effective under extreme flood events, which in the UK are often preceded by a period of prior catchment wetting (see Metcalfe et al., 2017; Hankin et al., 2017). In fact, by slowing the flow in some parts of the catchment, it is possible that the peak flow might increase elsewhere. This is called the synchronicity problem, the impact of which will vary from event to event (because of the different patterns of rainfall intensities) and with changes in the scale of catchment being considered. The distributed nature of this problem, and the potential for such mitigation effects to have impacts on other environmental factors, requires an integrated catchment modelling approach to evaluate possible implementation

scenarios. However, the outputs from such models will be associated with significant uncertainty, even after calibration on historical data. Thus, this is a prime example where the concept of models of everywhere and the local constraint of uncertainty using local information would be useful, especially when assessing the uncertainty in potential outcomes might make a difference to the decision that might be made. A re-evaluation of the concept of models of everywhere is therefore very timely.

The paper is structured as follows. Section 2 examines the concept of models of everywhere in more depth, highlighting the different dimensions behind this initial vision, and culminating in a set of technological requirements to support models of everywhere. Section 3 looks at the technological landscape, as it existed in the period 2007–2012, systematically reviewing the different technological requirements and concluding with an overall assessment of technology readiness at that time. Section 4 repeats this analysis, but looking at the state-of-the-art now. The paper then presents ongoing research in this area, including the identification of a research roadmap for the implementation of the concept of models of everywhere, (Section 5). Section 6 documents related work, including existing deployments of the concept of models of everywhere. Finally, Section 7 concludes the paper with some final reflections on models of everywhere from a technological perspective.

2. Models of everywhere unpicked

While models of everywhere at one level is quite a straightforward concept representing as association of models with particular places, at another level, it is a rich multi-dimensional conceptual framework. In particular, we discuss three (mutually supportive) dimensions with the goal of highlighting the technological requirements to support the overall vision:

Models of everywhere; Models of everything; Models at all times.

These are discussed in turn below.

2.1. Models of everywhere

2.1.1. Key characteristics

The starting point of models of everywhere is to move from generic models that can then be customised to particular locations, for example through appropriate parameterisation, to models that are specific to particular places. As such, they can be tailored to represent the behaviour at a specific place without the need to represent any other place. In particular, observations and inputs from local stakeholders can be used to constrain the uncertainty that is associated with environmental modelling (e.g. Beven, 2009).

Note that models of everywhere is often interpreted as models representing specific localised areas but the concept does not imply any particular scale; rather models of everywhere can represent local, regional, national and global scale with these models often co-existing. While there is an expectation that model parameterisations should be resolution dependent (e.g. McDonnell and Beven, 2014; Beven, 2019), in the absence of any adequate scaling theory for many environmental processes, particular models may need to be tailored for the scale at which they operate in terms of both process representations and effective parameter values. The approach is illustrated in Fig. 1: a generic model with a specific set of parameter values cannot represent flooding at all areas of the catchment and hence five localised models need to be developed.

A core motivation of models of everywhere is to *constrain uncertainty*, exploiting as much knowledge as is available about a particular place. This is developed further when we look at models of everything and models at all times in Sections 2.2 and 2.3 below.

¹ https://www.gov.uk/government/publications/national-flood-resilienc e-review.



Fig. 1. Models of everywhere: showing a range of models across a catchment (the Conwy) representing different places, in this case at the same scale (see also Fig. 2/3 which build on this).

2.1.2. Technological requirements

The main requirement of models of everywhere is very large-scale computational capacity. For example, if this approach were to be adopted for future flood prediction, there would be a need for the deployment of very large numbers of models all over the country at different scales. For illustration, consider models applied to support Flood and Coastal Risk Management (FCRM) in England.

For any given community (city, town, village or street), there may be an array of models developed, paid for and used by different organisations, for different purposes, and using different data resources (or, very often, common data sets, exploited in different ways). Most localities are included within national scale models, sometimes referred to as "strategic", with the outputs of the National Flood Risk Assessment (NaFRA² published as "Risk of Flooding from Rivers and Sea"³) being typically the most generic. Separate models have been applied to provide mapping of the flooding from "Surface Water"⁴ (ubiquitous in coverage, representing potential flooding from overland flows and ponding, rather than water overflowing from rivers or the sea), the "Risk of Flooding from Reservoirs" (predicting places at risk in the event of dams and impoundments being breached), and groundwater flooding.

More detailed models also exist in many places to support activities such as the economic appraisal of proposed flood defence schemes, flood risk assessments for proposed floodplain developments, or the detailed design, construction and maintenance of drainage systems. These models usually capture further information about infrastructure (e.g. bridges, culverts, weirs, sluice gates) and river channel surveys.

Organisations commissioning and owning such models may include the Environment Agency, which leads on FCRM in England for local government, water companies supplying drainage services, private developers or other landowners. It is usual for all of the above models to evolve over time, incorporating both new data and technical

 ² https://www.gov.uk/government/publications/flooding-in-england-na tional-assessment-of-flood-risk.
³ https://data.gov.uk/dataset/risk-of-flooding-from-rivers-and-sea1.

⁴ https://www.gov.uk/government/publications/flood-maps-for-surface -water-how-they-were-produced.



Fig. 2. Models of everything: showing the rich diversity of data and models that may exist for a given place and how they may inform each other.

improvements (e.g. better numerical solution schemes). It is also common for multiple instances of each model to be executed, i.e. many individual "runs" or simulations, to support scenario or uncertainty analysis. The Environment Agency has over 1500 such detailed local models, and reported in its 2010–2015 modelling strategy⁵ an investment of approximately £17 million per year in modelling and mapping and an additional £15 million in gathering and processing data to support FCRM.

This has significant *resource requirements* in terms of the number of processors or virtual machines to run these models and data storage, as well as the human costs in terms of developing and tailoring the models for given places, and analysing and understanding the outputs. For example, production⁶ of the "Risk of Flooding from Surface Water" maps cited earlier involved more than 70,000 individual simulations of flood inundation on a mosaic of approximately 7100 36 km² tiles covering all of England, run on a $2m \times 2m$ resolution digital height map that included over 91,000 manually-determined corrections. This process needed around two months for data preparation and one month of computer processing time, fully utilising a grid of over 100 GPU-accelerated PCs.

Where possible, technological and operational support would need to be provided for such development. This also asks important questions over the underlying distributed systems architecture to support such massive deployment, e.g. centralised, distributed or decentralized (or indeed combinations of different approaches).

The approach also asks fundamental questions over the relationship and consistency of models at different scales and how to support *reasoning across scales* in terms of supporting a deeper understanding of the science and all its complexities and inter-dependencies.

2.2. Models of everything

2.2.1. Key characteristics

The second dimension is concerned with exploiting information about a place, in particular, coupling a model of that place with as much local data as can be collected, thus embracing the heterogeneity of available data sources. The availability of such data is increasing significantly and now includes (Blair et al., 2019):

Remote sensing data collected by satellites or aircraft-borne instruments (including drones);

Other monitoring technologies that consist of a range of sensor technologies typically in close proximity with the observed phenomena, including the use of Internet of Things (IoT) technologies to provide real-time streaming and multi-faceted data about the natural environment;

Historical records held in a variety of locations and scales; The increasing amount of data available from national/local government and other open data portals often increasingly offering APIs, e.g. data.gov.uk;

⁵ https://www.gov.uk/government/uploads/system/uploads/attachment _data/file/292949/geho0310bsbt-e-e.pdf.

⁶ https://www.whatdotheyknow.com/request/208078/response/ 520170/attach/3/National%20modelling%20and%20mapping%20method% 20statement%20May%202013.pdf.

Data mining provides additional information from the web or social media;

Data collected from citizen science, with the potential to direct citizen science to areas of data scarcity.

Together, this adds up to the potential for having environmental data at an unprecedented scale (Blair et al., 2019). For example, if focusing on flood prediction, it is possible to use a variety of data sources such as historical Parish records and flood marks, satellite imagery, local sensors, photographs from social media and citizen science to help steer process-based hydrological models. Indeed many researchers are advocating such approaches in hydrology, e.g. (Di Baldassarre et al., 2009; Smith et al., 2009, 2015), and more generally in disaster risk reduction (McCallum et al., 2016).

The concept also naturally extends to other aspects of environmental science, for example collecting and analysing data around water quality issues, biodiversity, or soils and indeed the inter-dependencies between them.

The additional dimension of models of everything is shown in Fig. 2. The concept of models of everything has the potential to further reduce the uncertainty around predictions for a variety of different variables of environmental interest in a coherent way. This is particularly important where the relevant processes are intrinsically coupled, for example the water flows that drive the transport of nutrients from farmland and households into rivers and lakes. This constraint of uncertainty is important because many of the sources of uncertainty are epistemic in nature. Epistemic uncertainties are those are those that arise from lack of knowledge, in contrast to aleatory uncertainties that represent random variability that derives from 'irreducible natural variability' (see Beven, 2009, 2016; Rougier et al., 2013; Beven and Hall, 2014; Di Baldassarre and Beven, 2016). By definition, it is not possible to deal with epistemic uncertainties in process models without breakthroughs or deepening of knowledge about a given place and its states and behaviours. Beven et al. (2015) also talk about the role of models of everything in overcoming what they refer to as hyperresolution ignorance in modelling, that is evaluating the hyperresolution information produced by simulation to overcome the local lack of data and unknowns in scientific understanding (e.g. the understanding of subsurface structures in hydrology).

2.2.2. Technological requirements

Models of everything adds significantly to the requirements imposed on the computational infrastructure around five key areas:

- 1. *How to store the 'big data'*? With models of everywhere, there is a need to capture and store significant quantities of data about a given place, and then repeat this across all places. This therefore very quickly becomes a 'big data' problem. In many ways, though, this is more demanding than many areas of big data given the high level of heterogeneity in the data-sets with some of the data being structured and other elements being unstructured, and inevitably captured in a wide variety of formats (Blair et al., 2019).
- 2. *How to represent and manage the collected data?* Given the variety and heterogeneity discussed above, there is a need to represent, evaluate and manage the overall collection of data, and this must include support for interoperability, data discovery and also the association of appropriate meta-data and ontologies, including provenance information.
- 3. *How to ensure open access to data*? The concept of models of everything implies a move towards open data, where data is openly available for use and stored in a way that allows such open access (also important to support a more collaborative and crossdisciplinary science as required to interpret this data). As mentioned in the introduction, while this is technically straightforward to achieve, this requirement is more concerned with cultural issues, for example around the perceived value of data. Note that

ideally this open philosophy would also extend to models, with models available as open source.

- 4. How to make sense of the heterogeneous data elements? It is one thing to have access to this rich underlying data, but it is another thing to be able to make sense of this data and therefore data analysis techniques are also required to build this higher-level of understanding from the underlying data (cf. environmental data science (Blair et al., 2019)). This will inevitably imply the construction of data models using a rich array of statistical and machine learning techniques.
- 5. *How to combine process models with data models?* Once the data models are constructed, there is a need to couple the data model or models with process models to build a complete understanding of a given place. Techniques are therefore required to support model coupling between process and data models. In hydrology, it has been shown that even hydrological observations may not always be informative in model calibration and validation (e.g. Beven and Smith, 2015).

There are also important human and societal issues around privacy and security but, as mentioned in the introduction, this area is not considered in this paper (but is an important area of future investigation). Note that the execution of additional data models, and the need to allow for observational uncertainties also increases the computational requirements of the system.

2.3. Models at all times

2.3.1. Key characteristics

The final dimension is that models should be active at all times, constantly re-evaluating what is known about particular places and adapting accordingly. This does not necessarily mean that models are always executing as this would consume significant computational resources without any real gain. Rather, models should execute periodically and frequently, for example when new data becomes available, to understand the "idiosyncrasies of particular places" (Beven, 2007) and how this might change over time. At other times, the model would be in a quiescent state, but otherwise ready to re-execute at any time. This contrasts significantly with existing practice when distinct model runs are carried out infrequently under the auspices of a scientific experiment, with perhaps more runs carried out an at a pre-deployment phase to understand the sensitivities and uncertainties of a given (generic) hydrological model. This iterative approach is nicely aligned with the work of Box (1980) on the iterative relationship between practice and theory (Box's loop), recently extended by Blei (2014) in the context of latent variable models applied to complex data-sets. This may also introduce more consistency between models used for short-term forecasting, often applied in an adaptive framework where on-line data can be used for updating, and simulation models that are rarely updated.

This leads to a new perspective of "modelling as a learning process", as discussed in depth in the 2007 paper (Beven, 2007). The 2012 paper (Beven and Alcock, 2012) develops this further talking about models as hypotheses to be tested against current and historical observations with some models being rejected in favour of others and indeed this changing over time, so the current chosen model structure and associated assumptions best reflect the full idiosyncrasies of a given place as represented by numerous additional data observations. This in turn leads to an *adaptive* approach to modelling.

The final aspect to consider is what can be adapted about the model. There are various possibilities here, increasing in level of sophistication and ambition:

 The outcomes from a model can be adapted for the purposes of realtime forecasting when data can be made available for assimilation, and post-event analysis can then be used to inform local improvements to the model, including adaptation of parameter values to best represent behaviour at the current time;



Fig. 3. Models at all times: showing the meta-level reasoning framework associated with models as a learning process, extracting meaning from diverse data about a place, applying learning techniques to extract meaning from this data and making appropriate adaptations around model selection and parameterisation.

- 2. A number of models can co-exist in an ensemble approach, with model selection applied to identify the best models for that given place/time;
- 3. The internal structure and behaviour of a given model can be adapted, for example, by changing fine-grained elements of the underlying hydrology to best reflect the current place/time;
- 4. The representation of residual uncertainty can be adapted as more information is obtained locally.

Clearly, these approaches can also be combined in different ways. Indeed, a combination of all four offers a new and radical approach to models of everywhere.

The concept of models at all times is illustrated in Fig. 3.

The key motivation of models at all times is to offer a modelling framework that supports explicit reasoning about uncertainty, with the explicit goal of reducing uncertainty for a given place. More specifically, the approach also has the potential to deal with epistemic uncertainties, as argued in (Beven and Alcock, 2012). In this paper, following Beven (2006), the authors argue for an approach based on *limits of acceptability*, whereby models that perform well according to such limits are acceptable (and perhaps reinforced), while others are rejected, with this driven by the collected set of observational data (cf. models of everything). Finally, the approach has significant potential to deal with non-linearities and fundamental changes over time, for example related to climate change, with its emphasis on ongoing adaptation to the current context.

2.3.2. Technological requirements

Models at all times is crucial to the overall vision of models of

everywhere, but adds a whole new level of complexity in terms of the underlying technology requirements. In particular, the approach amplifies the underlying *resource requirements* and also the underlying *distributed systems architecture* as discussed in Section 2.1. For example, the approach requires the frequent execution of potentially ensemble models at large numbers of places and different scales.

The approach also introduces significant additional (mutually supportive) requirements:

- 1. *How to support adaptive reasoning?* As discussed above, 'models at all times' is fundamentally a learning process and implies that models are constantly adapted in response to new knowledge extracted from available data-sets. There is therefore the need to support this adaptive reasoning and ideally this should involve a strong element of automation as provided, for example, by autonomic computing (supporting self-adaptive systems) (Kephart and Chess, 2003; McKinley et al., 2004).
- 2. How to incorporate reasoning about uncertainty? Building on the above, it is important that adaptation decisions incorporate reasoning about uncertainty, and this implies making uncertainty explicit in the modelling process, and also incorporating approaches to deal with epistemic uncertainties and non-linearities as inevitably encountered in such complex systems.
- 3. *How to support adaptation*? A truly adaptive system requires ready access to a range of elements that can be changed. Supporting more coarse-grained adaptation is relatively straightforward, and implemented in terms of selecting from different models in model ensembles, or changing model parameters. Supporting fine-grained strategies is however more challenging as this requires intimate

Table 1

Technological requirements for models of everywhere, everything and at all times.

Requirement	Primary motivation(s)
R1: Massive resource requirements in terms of underlying computation and storage (cf. big data)	Models of everywhere, everything and at all times
R2: Appropriate underlying distributed systems	Models of everywhere,
architecture for models of everywhere	everything and at all times
R3: Support for reasoning across scales	Models of everywhere
R4: Appropriate data representation architecture,	Models of everything
recognising the heterogeneity of underlying data	
(structured and unstructured) and its complexities	
R5: Support for data discovery and navigation	Models of everything
R6: Supporting open access to data	Models of everything
R7: Providing rich data analytics methods to make sense of this data	Models of everything
R8: Supporting the integration of process and data	Models of everything, at all
models	times
R9: Support for adaptive reasoning	Models at all times
R10: Support for reasoning about uncertainty,	Models at all times
including epistemic uncertainties and non-linearities	
R11: Ability to support coarse-grain and fine-gran adaptation of environmental models	Models at all times
R12: Supporting deployment of models of everywhere	Models of everywhere,
at scale	everything and at all times
models R9: Support for adaptive reasoning R10: Support for reasoning about uncertainty, including epistemic uncertainties and non-linearities R11: Ability to support coarse-grain and fine-gran adaptation of environmental models R12: Supporting deployment of models of everywhere at scale	times Models at all times Models at all times Models at all times Models of everywhere, everything and at all times

access to the structure and behaviour of individual models in terms of, for example, alternative hydrological equations at the heart of the model. Most existing models will not provide such access, i.e. black box implementations. To fully realise the vision however, we need to go further than this and provide more white-box access to internal software architectures of environmental models, as provided by, for example, reflective architectures (Maes, 1987; Kon et al., 2002).

'Models at all times' also places additional emphasis on the need to integrate process and data models (discussed in Section 2.2) to support adaptive reasoning.

There is also an over-arching requirement emanating from this analysis, and that is the ability to support deployment at *scale*, and this implies the ready deployment of individual models (of everywhere) and also of large numbers of models at different scales. There are many dimensions to this scalability involving, for example, making it easier to deploy models in underlying computational infrastructure, whether provided by HPC or cloud facilities, offering software frameworks that can support the deployment of models or ensembles of models ready to be tailored for the idiosyncrasies of places, and automating the subsequent adaptation/learning process (hence the importance of selfadaptive approaches).

2.4. Overall analysis

'Models of everywhere' is an important and potentially crucial approach to environmental modelling, particularly in terms of managing uncertainty. A full implementation of the concept however imposes very significant requirements in terms of the technological infrastructure alongside other fundamentals, most notably cultural elements around a move to open science (incorporating more open approaches to data and modelling). The approach is best understood as a combination of models of everywhere, everything and at all times, with this trichotomy used to analyse the overall requirements in more depth in the discussions above. The resultant requirements are shown in Table 1 below.

These requirements can usefully be clustered as follows:

- The capacity and level of sophistication of the underlying *technological infrastructure* in terms of both computation and data (R1, R2, R4, R5);
- 2. The availability of rich *data analytics* capability to make sense of complex and highly heterogeneous data-sets (R3, R6, R7, R8);

- 3. The ability to support *modelling as a learning process,* including reasoning about uncertainties (R9, R10, R11);
- Practical issues around *deployment at scale*, including availability of open data and approaches to support large-scale deployment (R3, R12).

This clustering will be used in the assessment of the changing technological landscape as discussed in Sections 3 and 4 below.

3. Technological landscape (2007-2012)

3.1. Overview of the landscape

In the period 2007 to 2012, the landscape was dominated by *grid computing*. The concept of grid computing was first introduced in in the 1990s and became prominent with the publication of the seminal paper by Foster and Kesselman (1998), introducing the grid as a "blueprint for a new computing infrastructure". The term was introduced as a metaphor for the electricity grid, with the goal of making computational power as accessible and ubiquitous as electricity. Software platforms were developed to support the deployment of applications and services in the grid, most notably the Globus toolkit, with various versions released starting in 1997 with the last major release (Globus toolkit version 5) in 2009.⁷ Around this time, the grid was being superseded by cloud computing (for example, the first version of Amazon Web Services⁸ was introduced in 2006 with rapid growth since); this growth in cloud computing is discussed further in Section 4.1 below.

In parallel, researchers were becoming interested in the use of such computational power to support a range of application domains including, for example, eCommerce. Most notably, in the context of this paper, there was also great interest in *eScience*, that was, the use of technological infrastructure including the grid to support a new kind of computationally intensive and data-rich science (Hey et al., 2009). For example, in the UK, the national eScience programme ran from 2001 to 2010, supporting a range of infrastructure projects and application projects in areas as diverse as bioinformatics, neuroinformatics and medical informatics. In the environmental area, the most prominent project was climateprediction.net.⁹ Similar, large-scale initiatives were launched in other countries, for example in the States the National Science Foundation (NSF) funded a series of cyberinfrastructure initiatives starting around 2003, including the Open Science Grid¹⁰ developed by the Open Science Grid Consortium (OSGC).

3.2. Addressing the requirements

3.2.1. Underlying technological infrastructure

The emergence of the grid and also the eScience community that coalesced around the grid provided important expertise, experience and also facilities to support the development of models of everywhere. However, in practice (and this is clear in retrospect), the grid did not meet the full set of requirements to support the broader vision of models of everywhere.

Although the vision of the grid was to provide plentiful resources on demand, the reality was somewhat different at that time. The availability of resources varied greatly and depended on access to one of the experimental grid facilities that were introduced in different global centres. The overall distributed systems architecture was therefore one of centres at given fixed locations offering (by definition) relatively centralised services with partial access and limited control of these services. A number of researchers explored more decentralized

⁷ http://toolkit.globus.org/toolkit/.

⁸ https://aws.amazon.com/.

⁹ http://www.climateprediction.net/.

¹⁰ https://www.opensciencegrid.org/.





sophistication to meet the technology infrastructure requirements as identified in Table 1.

The main middleware technology used at the time was the Globus Toolkit, with the overall architecture of the Toolkit (v5) shown in Fig. 4.

This was a large and complex architecture with many dimensions but, as can be seen, the emphasis is on supporting resource sharing, and at a fairly low level of abstraction. As stated in the seminal paper on the "anatomy of the grid", Foster et al. (2001) argue that the grid was fundamentally about coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organisations". They go on to argue that this implies "direct access to computers, software, data, and other resources" and this sharing should be "highly controlled". Hence, the emphasis was very much on meta-level concerns such as standardised APIs to ensure interoperability, service discovery, access control and resource management. There was also more emphasis in practice on computational resources rather than data management, for example GRAM (Grid Resource Allocation and Management offered an architecture to submit and monitor (batch) jobs in the Grid (see Fig. 5). This is quite different though from the execution style required for models of everywhere.



Fig. 5. The architecture of GRAM (http://toolkit.globus.org/toolkit/docs/4.0/execution/wsgram/WS_GRAM_Approach.html).

architectures, for example climateprediction.net (mentioned above) and SETI@home¹¹ utilising BOINC¹² – a more peer-to-peer volunteer computing platform, but such initiatives were not mainstream and not integrated into other grid initiatives.

It is also important to emphasise that this was a research programme and hence the underlying platforms were not stable, with frequent changes over time in terms of services and facilities on offer. As will be seen, this contrasts significantly with what is available now in terms of both capacity and stability of services (see Section 4.2).

More fundamentally, the services on offer did not have the level of

As mentioned above, the data side was quite primitive with an emphasis on low-level facilities for access to data remotely (GridFTP¹³) and to assist in replication of data. While the grid was used successfully for a number eScience experiments involving elements of 'big data', the level of sophistication of data management was insufficient for the rich and heterogeneous data required for models of everywhere (and associated needs in terms of discovery and navigation), and indeed for environmental data more generally. As will be seen below, this is one area that has advanced significantly in the last few years.

There was also a general lack of experience of using grid computing

¹¹ https://setiathome.berkeley.edu/.

¹² https://boinc.berkeley.edu/.

¹³ http://toolkit.globus.org/toolkit/docs/latest-stable/gridftp/.

for the earth and environmental sciences. Other scientific communities were more advanced in terms of their use of grid computing and embracing eScience. This led to a lack of services specific to this field, e. g. to support environmental modelling in the grid, although parallel developments such as the OpenMI (Open Modelling Interface) standard¹⁴ as adopted by the Open Geospatial Consortium, provided important building blocks to support model deployment and (most crucially) interoperability across models.

In summary, grid computing was important in terms of establishing a community working together on distributed architectures and infrastructure and, in particular, for building a strong dialogue with the science community in terms of a new open, computational and data-rich style of science. However, there are a number of limitations that impacted on the feasibility of models of everywhere, most notably the difficulties of access to computational and data resources, the lack of sophistication of the distributed infrastructure particularly in terms of data management, and the primitive nature of many of the services on offer (also revisited below under 'deployment at scale').

3.2.2. Data analytics

The ability to represent and access very-large scale and highly heterogeneous data is important. Equally, it is crucial to have a range of techniques to make sense of this data. As can be seen above, this requires: a move towards open data as a prerequisite for open science; the availability of a rich set of techniques to analyse data; the ability to extend this reasoning across scales; and an integration of process modelling with data models produced to analyse the complex data (over and above the baseline requirement for storing, accessing and managing large and complex data-sets). Open data was in its infancy in the period 2007–2012 with data often regarded as core intellectual property with many institutions seeking ways to commercialise their rich data-sets. There was however, a growing recognition with the complexity of modern science, that a new, more open approach to data was necessary. For example, the Royal Society published "Science as an Open Enterprise" in 2012¹⁵ with a core recommendation:

"Scientists should communicate the data they collect and the models they create, to allow free and open access, and in ways that are intelligible, assessable and usable for other specialists in the same or linked fields wherever they are in the world"

This built on the emergence of Science 2.0 (Waldrop, 2008), seeking an open approach to science based on emerging Web standards (particularly Web 2.0 technologies offering user generated content and a move towards a more social web). In practice, however, at that time there were many cultural and technological barriers to a world where data-sets were available for open access in common repositories.

In terms of making sense of data, the environmental sciences, including hydrology, make extensive use of process models to understand fundamental processes of nature and then use these models to make future predictions. A wide range of process models have been developed, for example in hydrology, where there have been recent attempts to incorporate multiple process components into a common framework (e.g. Fenicia et al., 2011; Clark et al., 2015). In applications for flood risk assessment, there have been many codes routinely used both by industry and researchers to model the flow of water through the landscape, including interactions with physical infrastructure systems. These codes can be categorised reasonably precisely in terms of the approximations made to a set of physical governing equations (in the case of flood models this means simplifications of the fundamental Navier-Stokes equations for fluid dynamics). Even so, there remain differences in the interpretation of the prototypical physical equations, the numerical schemes that are used to solve them, the discretisations involved in applying those schemes to real data sets, and in the very many "edge cases" for which special solutions are required. Benchmark comparisons¹⁶ have shown how important these differences can be in controlling the results of flood simulations in various situations. There is also a strong body of research on training models based on historical data, and current observations can be used to steer future states of the model (data assimilation) (Lahoz et al., 2010; Park and Xu, 2017).

More generally, in the time period under consideration, there was a deep concern that process models alone are not sufficient, and that fundamental issues remain, for example, reasoning about uncertainty and dealing with epistemic uncertainties and non-linearities in complex systems. Indeed, this is the prime driver for models of everything. This reflects a sense that it is necessary to integrate the process model view of science with one that recognises the importance of data and associated data analytic techniques (effectively data models). This is a significant cross-disciplinary challenge requiring input from environmental, computer and mathematical scientists. At that time, this dialogue was not happening (discussed further in Section 4.1). Scientists also tended to focus on specific experiments and studies to understand phenomena at a given scale, so reasoning across scales was in its infancy.

Overall, even by 2012, there were major barriers around data analytics that made it very hard to support the realisation of models of everywhere.

3.2.3. Modelling as a learning process

As discussed above, the perspective of models as a learning process is the most important but also most demanding aspect of implementing models of everywhere requiring a new, adaptive approach to learning. From our analysis, this breaks down into support for adaptive reasoning, explicitly representing consideration of uncertainty in this reasoning, and also being able to carry out both coarse-grained and (importantly) fine-grained adaptations.

In the field of computer science, in the time period under consideration, a deep understanding of adaptive computing developed. For example, IBM launched a new initiative examining autonomic systems in 2001, that is systems that can self-manage (mirroring the autonomic functioning of nervous system in the human body), in terms of a range of self-* properties (e.g. self-awareness, self-configuration, self-healing and self-optimisation) (Kephart and Chess, 2003). More generally, there was a large literature around software architectures to support self-adaptation (including reflective architectures), the use of control loops in decision making, and the use of more advanced machine learning techniques to support higher levels of autonomic self-management, for example dealing with unknowns (Oreizy et al., 1999; McKinley et al., 2004). However, there has been little or no consideration about how such techniques can be used in terms of adaptive environmental modelling. Existing environmental models are also often written in older programming languages, most notably Fortran, and tend to be monolithic, black box implementations, hence do not lend themselves to the implementation of fine-grained adaptation strategies.

There is also the important requirement to represent and reason about uncertainty explicitly as part of the adaptation process. At that time, there was growing recognition of the need to represent uncertainty in modelling and reasoning about uncertainty across scientific experiments. For example, the UncertWeb project introduced techniques to capture uncertainties as meta-data in web-based environments, with details of the resultant UncertWeb framework published in early 2013 (Bastin et al., 2013). Researchers had also developed a number of frameworks to reason about uncertainty in scientific experiments, including seminal work by Beven and Binley (1992, 2014) and others

¹⁴ http://www.openmi.org/.

¹⁵ https://royalsociety.org/~/media/Royal_Society_Content/policy/projects/ sape/2012-06-20-SAOE.pdf.

¹⁶ https://www.gov.uk/government/publications/benchmarking-the-latest -generation-of-2d-hydraulic-flood-modelling-packages.

Table 2

Assessment of technological readiness (2007–2012).

Requirements cluster	Technological readiness	Most significant barriers
Technological infrastructure	**	Insufficient level of resources offered by the grid; lack of stability of grid platforms; lack of sophistication of services offered; lack of support for complex and highly heterogeneous data.
Data analytics	*	Lack of progress towards open data; immaturity and lack of cross- disciplinary dialogue on data analytics; lack of sophistication in dealing with uncertainty in process models; lack of research on process and data model integration; lack of research on
Modelling as a learning process	**	reasoning across scales. Lack of cross-disciplinary research looking at adaptive techniques in environmental modelling; little support for fine-grained adaptation due to existing model structures; major issues around representing and reasoning about uncertainties; lack of support for epistemic uncertainties and dealing with non-linearities
Deployment at scale	*	Low level of abstraction in grid environments; lack of programming models or frameworks to support deployment at scale.

Key (readiness level): **** = very high (no significant barriers); *** = high (some significant barriers); ** = medium (a number of important barriers); * = low (major barriers remain).

(see Renard et al., 2010; Vrugt and Sadegh, 2013; Nearing et al., 2016). As discussed in the 2012 models of everywhere paper, Beven and Alcock (2012) were just starting to think about reasoning about uncertainties in model selection or rejection (as a key part of models as a learning process).

In summary, most of the building blocks were there by 2012, but the work was fragmented and split across many communities, and key issues remained over how to support more advanced reasoning of uncertainties, including dealing with epistemic uncertainties.

3.2.4. Deployment at scale

Finally, and importantly, there is the key question of whether there was sufficient advancement at that time to support deployment of the kind of scale that makes models of everywhere a reality. As discussed above, there are several key dimensions to support such large-scale deployment, including how easy it is to deploy individual models, what support there is to then repeat this across many places (at different scales) and also whether the learning (and hence tailoring process) can be automated. The latter issue is intrinsically inked to the support for self-adaptive modelling and hence we focus more on the first two issues.

One of the key problems with deploying in the grid environments, or indeed to other HPC facilities, is the low level of abstraction offered by software platforms. This was discussed in the consideration of the Globus Toolkit above. Given this, the development and deployment of even an individual model is a tedious, expensive and error prone process and this in itself is a barrier (Simm et al., 2018) to the deployment of models of everywhere. This is a barrier to more general deployment across a range of places where the individual models need to be specific to this place, both initially and also with the model or models refined over time to reflect the particular idiosyncrasies of this place. This implies some form of software framework coupled with models as a learning process and, at that time, this was significantly beyond the state-of-the-art for model development.

It is interesting to note that the initial models of everywhere paper (Beven, 2007) discusses an object-oriented approach to programming

models of everywhere, mapping individual active spatial objects to places and also explicitly representing the relationship between places (mainly in terms of fluxes). This is an attempt to seek a higher level of abstraction to support the deployment of models of everywhere. At the time of writing, object-oriented computing and indeed distributed objects were an important area of research reflected in the importance of technologies such as CORBA (Common Object Request Broker Architecture).¹⁷ This approach is now largely superseded by alternative programming models, reflecting (most principally) difficulties in realising distributed objects in Internet-scale developments.

3.3. Overall assessment and technological readiness

It is clear from the assessment above that, even by the end of this period (2012), there were major technological barriers in terms of the deployment of models of everywhere. Our overall assessment is summarised in Table 2, which shows an overall rating against each of the requirements together with the identification of the most important barriers.

As can be seen from Table 2, the overall readiness level is generally low to medium, with important barriers remaining across all categories. It is interesting to note that quite a number of the barriers are due to a silo-ed approach to research and can be addressed by more crossdisciplinary collaboration in this area. Overall, we would argue that, in the period 2007–2012, the vision of models of everywhere was right but the technology was not ready. We continue our discussion by considering how things have advanced to date, noting important developments that make an implementation of the concept more realistic.

4. Technological landscape (current)

4.1. Overview of the landscape

The technological landscape has changed enormously since 2012 and indeed this is one of the key drivers to revisit the concept of models of everywhere in terms of technological readiness. In particular, there have been three mutually supportive areas of significant innovation, namely cloud computing, data science and IoT. We look at each in turn below.

The concept of *cloud computing* first came to prominence in the last decade. For example, Amazon introduced Amazon Web Services, as an early cloud offering, in 2006. It has really been in the last five years though that the area has exploded in terms of scale and sophistication of the underlying services on offer. The cloud is defined as "a set of Internet-based application, storage and computing services sufficient to support most users' needs, thus enabling them to largely or totally dispense with local data storage and application software" (Coulouris et al., 2011). Cloud computing further promotes the view of everything as a service, from low-level services such as data storage or virtualised machines, through intermediary middleware services supporting parallel/distributed computing or database facilities, through to a plethora of applications (referred to as Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS)). Cloud computing may be offered by companies and made available to others as services, i. e. public clouds such as those offered by Amazon, Google, IBM, Microsoft and Yahoo, or private clouds that can be established within an organisation or associated community (e.g. using open source software such as OpenStack or CloudStack). Hybrid solutions are also possible where an organisation may have their own private cloud but extended with extra capacity from public clouds. There is also a move in cloud computing from owning resources to a more *elastic* use, where resources can be requested (and paid for in the case of public clouds) only when required.

The growth of cloud computing over the last five years in particular

¹⁷ http://www.omg.org/spec/CORBA/.



Fig. 6. Expected growth of IoT (https://www.i-scoop.eu/internet-of-things-guide/).

has been phenomenal. For example, a report by Cisco indicates that in 2015 total data storage capacity in data centres is 382 EB, with this projected to grow to 1.8 ZB by 2020¹⁸. There has also a corresponding growth in processing capabilities, and innovation around cloud services, most notably for the purposes of this paper in the area PaaS, with a wide range of new services introduced to storage and process massive datasets, e.g. BigTable, Cassandra and HBase in terms of 'big' data storage and MapReduce and Apache Spark in terms of distributed computation (we return to this innovation in Section 4.2 below).

The developments in cloud computing have also stimulated interest in 'big data' or more generally *data science*, that is the science of analysing and making sense of very large and/or highly complex data-sets. This is a fundamentally cross-disciplinary area of study involving, for example mathematical sciences, computational sciences and areas of application. To support this, a number of such cross-disciplinary institutes have been set up worldwide, including the Alan Turing Institute in the UK, and Data Science Institutes at Berkeley and Columbia in the United States, and Imperial College London, UCL, Warwick and Lancaster also in the UK, amongst many others. With the huge investments in data science, there is a growing body of literature on techniques to extract meaning for large and complex data-sets, including techniques that embrace unstructured data. More importantly, there is a dialogue across disciplines to understand how different techniques can work together to resolve major challenges around big data.

A lot of the research in data science is targeted at underlying algorithms and their scalability and efficiency. There is also an emphasis on more applied research, most notably in the areas of eCommerce and marketing, smart cities, logistics and transport, and also health and wellbeing (Blair et al., 2019). There is also huge potential in data science for the natural environment although, perhaps surprisingly, this is an

¹⁸ https://www.cisco.com/c/dam/en/us/solutions/collateral/service-pro vider/global-cloud-index-gci/white-paper-c11-738085.pdf area that is relatively under-developed (it is though one of the major themes of the Data Science Institute¹⁹ at Lancaster University, UK).

Finally, there have been significant developments in the area of *IoT*, with the Internet evolving from being an Internet of computers to one that is an Internet of 'Things', with the Things being everyday objects with embedded intelligence (Atzori et al, 2010). Experts predict that IoT will embrace over 50 billion devices by 2020 (see Fig. 6).

As with data science, the main growth areas are expected to be around smart cities, logistics and transport and health and wellbeing. There is also significant potential for IoT deployments in the natural environment; for example, Nundloll et al. (2019) describe an experiment in deploying an environmental IoT in a catchment in Wales.

It is clear however that this is an area in its infancy. The real significance of IoT technology in this area is when data can be combined with other sources including remote sensing, earth monitoring technologies, historical records and other data mined from the web in support of models of everything (as discussed in Section 2.2).

These technologies are mutually supportive in that cloud computing provides the underlying very large-scale and elastic pool of resources and associated services to store, process and present very large data-sets. Data science provides a range of methods to make sense of complex data and extract meaning from this data, and IoT technology provides access to real-time observations on a very large scale. This symbiotic relationship is illustrated in Fig. 7.

4.2. Addressing the requirements

4.2.1. Underlying technological infrastructure

The underlying technological infrastructure has changed significantly in terms of both the availability of large-scale computational resources and also the stability of the associated platforms. There has also

¹⁹ http://www.lancaster.ac.uk/dsi/.



Fig. 7. The symbiotic relationship between cloud computing, data science and IoT.

been significant innovation in this area with an explosion of new services now available.

The developments in cloud computing, as documented above, are particularly significant in this regard. Whereas grid computing was a rather niche and immature technology, cloud computing provides access to an abundance of underlying resources (in terms of both computation and storage) and also an ever increasing set of associated services. The services most relevant for models of everywhere include:

A rich underlying set of programming constructs to support distributed programming, including service-oriented architecture, containers and microservices/serverless computing, e.g. Docker²⁰ and Rocket²¹ (for containers) and OpenWhisk²² and AWS Lambda²³ (for microservices/serverless approaches);

Services to support the subsequent deployment and execution of complex distributed executions, e.g. Kubernetes²⁴ and ZooKeeper²⁵; A range of underlying storage architectures that cater for very large scale and highly heterogeneous data-sets, including unstructured data, e.g. Cassandra²⁶ HBase²⁷ and MongoDB²⁸;

Parallel and distributed programming paradigms to process and manipulate such data-sets, including historical and streaming datasets, e.g. the Hadoop framework²⁹ MapReduce (Dean and Ghemawat, 2008), Spark³⁰ and Pig³¹;

Techniques to semantically enrich and subsequently navigate very large scale and highly heterogeneous data-sets, e.g. building on technologies such as OWL, SPARQL and RDF³² and also graph databases such as GraphDB³³ AllegroGraph³⁴ or Neo4j³⁵;

Software frameworks and associated libraries to support data analytics, e.g. Mahout³⁶ or RStudio³⁷;

³⁴ https://allegrograph.com/.

³⁶ http://mahout.apache.org/.

Services to support scientific workflow in the cloud, e.g. Taverna³⁸ and Kepler.³⁹

There is also strong interest in achieving integration between cloud computing and IoT technology, although this work is at early stages of development. Most significantly, there is a rapidly growing body of research around edge computing (sometimes also referred to as fog computing) to provide intermediary storage and processing capabilities closer to end devices (Lopez et al., 2015). For example, edge devices can be used to carry out initial analyses of real-time streaming data from IoT devices, with only aggregate or significant data then sent to the cloud environment. Edge computing can also support the integration of mobile devices (Ahmed and Ahmed, 2016).

The technological landscape has therefore changed dramatically with many of the technologies now in place to support models of everywhere. A number of significant barriers though still remain, most notably the lack of standardisation in cloud computing, with different providers offering quite distinct programming paradigms and APIs. This leads to problems of vendor lock-in and also difficulties in managing computations that span multiple providers (including hybrid cloud environments embracing public and private providers). There are also difficulties in programming and managing the underlying technological infrastructure especially when combining cloud computing with IoT technology, the result being a rather sophisticated but highly complex system in itself (more accurately described as a system of systems (Jamshidi, 2011)). We return to this point below (under deployment at scale).

4.2.2. Data analytics

There have been similar advances in terms of data analytics. There is now much more awareness of the need to move to open science, including the need for open data policies. Governments and research funding bodies are also moving towards the need for open data, and there is a similar move towards more open, reproducible or repeatable science.⁴⁰ It is fair to say though that important barriers remain and these tend to be cultural rather than technological.⁴¹

In terms of making sense of data, the emergence of data science as a discipline is strongly encouraging albeit with a need to attract more data scientists to work on environmental challenges and problems (Blair et al., 2019). The most important development has been the cross-disciplinary dialogue that is now happening within the data science community involving statisticians, computer scientists and domain experts (amongst others). This is very significant and is leading to breakthroughs in terms of efficient algorithms and their application in important societal problems. In the Data Science Institute at Lancaster, for example, we are interested in how contemporary techniques such as extreme value theory, changepoint analysis, time-series analyses and statistical/machine learning can be applied to complex environmental data. We are also particularly interested in how resultant data models can co-exist and inform process models, combining stochastic and deterministic understanding of complex environmental phenomena. While there is increasing awareness of the potential of such approaches, this is a relatively immature area; solutions tend to be ad hoc and a more principled understanding of how such techniques can work together has yet to emerge. There is also a similar narrative around reasoning across scales; while there is more experience of this in the earth and environmental sciences, the solutions are also quite ad hoc and often not shared across different areas of study.

²⁰ https://www.docker.com/.

²¹ https://coreos.com/blog/rocket.html.

²² https://openwhisk.incubator.apache.org/.

²³ https://aws.amazon.com/lambda/.

²⁴ https://kubernetes.io/.

²⁵ https://zookeeper.apache.org/.

²⁶ http://cassandra.apache.org/.

²⁷ https://hbase.apache.org/.

²⁸ https://www.mongodb.com/.

²⁹ http://hadoop.apache.org/.

³⁰ https://spark.apache.org/.

³¹ https://pig.apache.org/.

³² https://www.w3.org/standards/semanticweb/.

³³ http://graphdb.ontotext.com/.

³⁵ https://neo4j.com/.

³⁷ https://www.rstudio.com/.

³⁸ http://www.taverna.org.uk/.

³⁹ https://kepler-project.org/.

⁴⁰ http://royalsociety.org/uploadedFiles/Royal_Society_Content/policy/proje cts/sape/2012-06-20-SAOE.pdf.

⁴¹ http://crc.nottingham.ac.uk/projects/rcs/OpenScience_Report-Sarah_Curri er.pdf.

In conclusion, there have been significant developments since 2012, particularly in terms of the required cross-disciplinary dialogue around data science for the natural environment. Nevertheless, this work is still at a relatively early stage of maturity with important (and fairly unique) challenges of this area still to be addressed (Blair et al., 2019).

4.2.3. Modelling as a learning process

As discussed above, many of the building blocks for models as a learning process were already in place by 2012, albeit fragmented across different communities. The state-of-the-art now is quite similar and there remains a need for stronger cross-disciplinary dialogue between researchers working on environmental modelling, data science and adaptive/autonomic computing. The most significant changes in this time have been: i) advances in areas such as statistical and machine learning that directly supports meta-reasoning about model selection and rejection, and ii) the computational capacity offered by the cloud, which supports both the execution of complex environmental models in the cloud, and the execution of associated reasoning algorithms.

There has been little progress on the crucial area of uncertainty – in terms of representing uncertainty explicitly in computations, and also reasoning about uncertainty as part of the decision making process. More generally, one of the most significant developments over this time period in the environmental sciences has been the recognition of the unavoidable uncertainties associated with predictive models, whether used for simulation or forecasting purposes (e.g. Beven, 2009). As noted earlier, a primary driver for the models of everywhere concepts was the potential for using local information to constrain local uncertainties in predicting local variables. This is not just a problem of assessing the statistics of model residuals (though many studies have approached the problem in this way). This is because many sources of uncertainty are the result of lack of knowledge about processes, variables or forcings (particularly into the future) that are not necessarily easily represented in simple statistical forms. In particular, input uncertainties will be processed through the nonlinear dynamics of a model to produce complex nonstationary residual structures, that will then interact with uncertainties in the observational data used in model evaluation, which might also have associated epistemic uncertainties (e.g. in hydrology, arising from the rating curves used in the estimation of river flows, (see Westerberg et al., 2011; Westerberg and McMillan, 2015; Coxon et al., 2015)). These issues underlay the development of the Generalised Likelihood Uncertainty Estimation (GLUE) methodology (e.g. Beven and Binley, 1992, 2014; Beven, 2006, 2016), which includes some statistical methods as special cases.

As new data become available, it should be possible to learn more about the characteristics of the uncertainties associated with different predictands, at least where the new data are informative (that this may not always be the case has been shown by Beven et al., 2011 and Beven and Smith, 2015). In doing so, it will be possible to combine prior information with the new information to update the estimates. This leads naturally to a form of Bayesian reasoning, where uncertainties can be represented as probabilities, but much more research is needed in environmental models to understand how best to define the likelihoods used in the Bayesian methodology. Simple statistical likelihood functions used with multiplicative Bayesian updating appear to lead to overconditioning of model parameters because they do not take any account of the epistemic nature of sources of uncertainty (e.g. Beven et al., 2011; Beven, 2016, 2019). There are also issues of whether even the best models might be fit-for-purpose for the type of decisions that they might be used for (see the discussion of Beven and Lane, 2019).

A critical aspect of the models of everywhere concept is the potential for using local knowledge within this learning process to improve the representations of places. This is where information from local stakeholders and the Internet of Things might be used in local model evaluations to reject potential model structures and constrain uncertainties in parameterisations and outcomes. This can be considered as an extension of the collaborative and participatory learning that has already been used in a number of local flood risk assessments and water resource management projects (e.g. Lane et al., 2011; Landström et al., 2011; Evers et al., 2012; Maskrey et al., 2016; Ferré, 2017; Basco-Carrera et al., 2017; see also Voinov et al., 2016). An important component of this learning process is the potential to visualise model outcomes at scales that allows consideration of local detail by local stakeholders so that different scenarios (and their uncertainties) can be explored in collaborative ways (Hankin et al., 2017, see below).

4.2.4. Deployment at scale

There have been several important developments in terms of deploying at scale, with containers in particular making it far easier to deploy and subsequently manage executing models in the cloud in a platform-independent manner. The availability of cloud-based workflow engines is also significant, although there are questions over whether workflow offers the right abstraction for all elements of environmental modelling (Blair et al., 2019).

More generally, there is still a problem-implementation gap (France and Rumpe, 2007) between what scientists would like to do in the cloud, and the level of support offered by existing technologies and services, with a prior knowledge of the underlying technical details required. This makes it very time consuming and also error prone to execute environmental models or ensemble models in the cloud, and also requires access to computing expertise, which may be a scarce resource in many environmental research labs. In the context of models of everywhere, the models may themselves be quite complex, involving different ensembles of process models or the integration of process and data models for example. This makes the cost quite prohibitive, especially when this would entail the deployment of many instances of these models at many different places and scales.

Software frameworks offer a promising technology to support the more rapid deployment of recurrent software architectures (Johnson, 1997). Software frameworks are tailored towards particular domains of application, abstracting over the lower level details and capturing the commonalities within that domain, while allowing some degree of specialisation. They are heavily used in cloud computing, for example MapReduce abstracts over the complexities of managing a large and complex underlying cloud infrastructure and supports the execution of distributed algorithms in the cloud, allowing the user to plug-in and specialise the computation through providing application specific map and reduce operations (Dean and Ghemawat, 2008). At present though such frameworks tend to be relatively generic, e.g. dealing with distributed computation, and are not specific enough to support something as domain dependent as environmental models.

In terms of programming models, distributed objects have now been replaced by alternative paradigms supported in the cloud, around service-oriented architecture enhanced by concepts such as deployment in containers and optionally support for microservices. This approach overcomes the problems associated with distributed object technology, being much better suited to large-scale Internet wide deployment. Some research is required though in terms of how to map models of everywhere on to such programming concepts.

4.3. Overall assessment and technological readiness

It is apparent from the discussion above that there have been significant advances in the underlying technology to support the vision of models of everywhere. Equally, a number of barriers remain. Our overall assessment is summarised in Table 3, repeating the style of analysis carried out for the period 2007–2012 (in Table 2).

From this analysis, we can see that there have been significant shifts in readiness around the underlying technological infrastructure and in data analytics and also (partially) around supporting deployment at scale. Support for modelling as a learning process has not changed much although the developments in cloud computing and data analytics does offer the potential (as yet unrealised) of significant advances in this area.

Table 3

Assessment of technological readiness now.

Requirements cluster	Technological readiness	Most significant barriers
Technological infrastructure	***	Lack of standardisation in cloud computing; difficulties in managing complex underlying distribute systems infrastructure (or systems of systems).
Data analytics	***	Cultural impediments to open data; need to address particular data science challenges related to the environment, including around process and data model integration and on reasoning
Modelling as a learning process	**	across scales. Lack of cross-disciplinary research looking at adaptive techniques in environmental modelling; little support for fine-grained adaptation due to existing model structures; major issues
		around representing and reasoning about uncertainties; lack of support for epistemic uncertainties and dealing with non-linearities.
Deployment at scale	***	Problem-implementation gap and the need to raise the level of abstraction in terms of supporting execution in the cloud; lack of experience of using cloud programming paradigms in this area.

Key (readiness level): **** = very high (no significant barriers); *** = high (some significant barriers); ** = medium (a number of important barriers); * = low (major barriers remain).

The need for cross-disciplinary dialogue is a common theme across all these areas and is crucial in terms of addressing the remaining barriers.

Overall, we conclude that, in terms of technological readiness, the time is right to carry out large-scale experiments of the concept of models of everywhere. The next section explores ongoing research in this area.

5. Initial experiments and research roadmap

Ongoing research at Lancaster is looking at an experimental deployment of the concept of modes of everywhere in the area of hydrology, supported by recent developments in cloud computing, data science and new sources of data (including but not limited to IoT technology). The initial deployment is targeting a specific place with the intention of having a modelling framework that is able to capture and indeed learn the idiosyncrasies of that place. The overarching goal of this work is to identify software architectural principles for implementing models of everywhere in the cloud with a view to supporting more widespread deployment of models of everywhere at different places and at different scales (discussed further below). We are also strongly interested in supporting decision making at different scales, for example over the potential effectiveness of different natural flood management strategies and also over how to use constrained national or regional budgets most effectively.

The high-level systems architecture is as shown in Fig. 8 below.

This recognises the existence of multiple sources of data and the importance of integrating this data and, in turn, looking at model integration on top of this, which includes both data and process models coupled together. The top layer then supports interrogation and querying of the information about that particular place. This maps on to a more detailed cloud-based systems architecture exploiting the range of services supported by the cloud in each of these areas. This is shown in Fig. 9.

Through this work, we intend to overcome the remaining technological barriers around models of everywhere and therefore open the door to the desired, more widespread deployment of the concept in hydrology and beyond.

Further details of this research can be found in Towe et al. (2019) with Edwards et al. (2017) also discussing human and societal dimensions of the research around decision making.

Our overall research roadmap is summarised in Table 4, which shows the key research questions and associated areas of investigation.

Having deployed the concept of models of everywhere at a given place, we then hope to consider how to generalise the approach to model other environmental facets at that place, to model other places (including places at other scales), and to support coherent reasoning across scales. This also involves key questions over discretisation, especially given the fact that data may exist at different scales for a given place. In the longer term, we will also be interested in how the concept can be applied to other areas of environmental science, including biodiversity and soil management, and also how models of everywhere can help us in understanding the inter-dependencies across such areas (a key motivation of models of everything as discussed above). This quickly becomes a large research agenda that goes beyond the scope of our research study, and we hope to stimulate other research to address these



Fig. 8. High-level systems architecture.



Fig. 9. More detailed cloud-based systems architecture.

Table 4

Overall research roadmaps.

Research questions	Potential solutions
How to effectively and efficiently map the concept of models of everywhere	Identifying appropriate software architectures for models of everywhere
on to contemporary cloud programming paradigms and	and examining the mapping on such architectures to service-oriented
associated services?	architecture, containers and microservices.
How to deploy models of everywhere at	Investigating the role of specialised
scale, with a view to supporting the rapid deployment of new instances?	software frameworks for models of
	techniques from the model-driven
	engineering community, especially
	around domain specific languages.
How to achieve data integration given	Investigating underlying cloud storage
highly heterogeneous sources of data	technologies such as Cassandra or HBase,
(including unstructured and more	and associated technologies for semantic
structured data)?	integration (especially ontologies and
Here is a second s	linked data).
How to make sense of this complex data?	Explore a range of appropriate data
	combination
How to achieve integration between process	Seek underlying principles related to
models and data models?	process and data model integration;
	investigate how this can support a
	reduction of uncertainty and also how it
	can deal with epistemic uncertainty and
	non-linearities.
How to realise the concept of modelling as	Seek to bring together expertise in
a learning process?	adaptive/autonomic computing and
	environmental modelling; seek ways to
	annotate computations with uncertainty
	and use this in the reasoning/adaptation
	coarse and fine-grained adaptation: seek
	approaches to accept, refine or reject
	models, including consideration of the

limits of acceptability approach.

key issues.

6. Related work

There are, of course, already models of everywhere (and to some sense everything) in the sense of global earth system science models that have been developed from global atmosphere and ocean dynamic circulation models. Examples are the Japanese Earth Simulator (Habata et al., 2003); EC-Earth (Hazeleger et al., 2010) and the Community Earth System Model (Hurrell et al., 2013). While these are still limited to grid resolutions of kilometres for global applications, these systems commonly include the possibility of nesting finer grid domains, with boundary conditions provided by global simulations. The philosophy of such approaches, however, has been quite different from that presented here. The model structure and parameterisations are generally fixed, so that application everywhere has been a matter of finding appropriate effective parameter values for different grid locations using whatever data might be available.

There have also been some attempts to produce distributed modelling systems that could be applied widely at finer grid scales, allowing for a more flexible choice of structures. In hydrology, for example, there was the inter-agency Object Modelling System (OMS) of Leavesley et al. (2002) that developed into a more general modelling system (Lloyd et al., 2011; David et al., 2013). More recently, Clark et al. (2015) have proposed the Structure for Unifying Multiple Modeling Alternatives (SUMMA) framework. In both cases, several different model representations were provided for the user to choose from in producing a models structure for a particular catchment area. Within these systems, the expertise of users can be elicited to define appropriate model structures, although identification of appropriate model parameters and hypothesis testing of competing model structures are still major issues (e.g. Weiler and Beven, 2015). The type of approaches presented here could be used with such systems.

In flood risk management, there are inherent motivations to view



Fig. 10. (a) Co-production and local improvements of a flood risk model by Flanders Environment Agency (Vlaamse Milieumaatschappij, VMM), illustrating placebased dialogue about local model errors and data; (b) Difference viewer allowing cooperating parties to assess model updates, supported by contextual information.

modelling as a process of learning about places, stemming from two distinctive features of the problem. Firstly, the likelihood and impacts of flooding, although driven ultimately by weather and climate, are strongly influenced by local features of landscapes, land use, and human activities. In some cases, even very small topographic features or infrastructure assets can have a significant control on flood risk, for example by directing the flow of flood waters towards or away from buildings. This information is not always captured well (if at all) in generic model structures and data sets. Secondly, the assessment of flood risk involves gathering information about extreme events, which tends to place an emphasis on historical knowledge, often reliant upon detailed knowledge of the locality for interpretation, and on the updating of risk assessments as new observations become available.

For these reasons, some flood risk management applications already implement frameworks for iterative co-production of modelling, based on the incorporation of knowledge about specific localities from multiple stakeholders. One such system⁴² has been developed for the Flanders Environment Agency (Vlaamse Milieumaatschappij, VMM) for mapping areas at risk of flooding from surface runoff. Here, a web-based interface creates a shared collaboration space enabling local partners, such as town councils or local water- and sewer managers, to engage in a dialogue about model improvements (Fig. 10).

Important features that have been incorporated within the modelling through this process include areas where flood water can be held back by embankments, or drained by pumps and control structures (e.g. gates, sluices, weirs, culverts) that are known to local staff and may not be represented adequately without that detailed local knowledge, such as

⁴² https://www.youtube.com/watch?v=lEWtELDWTsU&feature=youtu. be&t=74.



Fig. 11. Engagement, capture of local data and knowledge on the iTable.



Fig. 12. FiGIS and model outputs on the iTable following one of the workshops.

the flood retention storage and flow control structures represented in Fig. 10.

The co-production website is shared with professional partners and within its first three months of operation enabled more than 9000 detailed improvements to data and modelling to be implemented together with nearly 300 re-simulations for 103 sub-models involving 150 organisations, and resulting in positive evaluations of model improvements at 500 locations across the whole of Flanders.

The VMM example discussed here explicitly exposes modelling as part of a process of learning about place through knowledge sharing, supported by digital technology. A more constrained example is the flood hazard mapping of the FEMA National Flood Insurance Program⁴³

where model-based flood risk assessments are published and can be challenged by individuals, on the basis of detailed local knowledge, via a web-based facility.

Stakeholder interaction with the outputs of models of everywhere can also made more direct and local, to demonstrate and explain assumptions, and to alter inputs or outputs as a way of iterating to a coproduced model of a place. In this way local models will be better constrained by the local stakeholder information, and more trusted if done well. As an example of this way of working, a series of workshops, sponsored by Natural England, were used to engage local farming communities on the potential benefits of 'working with natural processes' (WWNP) to mitigate flooding, often called Natural Flood Management (see Hankin et al., 2017 for modelling concepts).

Two engagement devices were used. An Augmented Reality Sandbox was used to provide real-time feedback on the response of flow pathways to a user sculpting channels in sand. Virtual inputs to the sandbox are

⁴³ https://www.fema.gov/national-flood-insurance-program-flood-hazard -mapping.

controlled by waving a hand over the sandbox. Water flow pathways and storages are then shown by projection of blue onto the sand. This was used as a precursor to the demonstration of more quantitative modelling results, with visualisations being projected onto a large interactive iTable and shown in Fig. 11.

Engagement using the iTable followed 4 key steps:

- 1) The model, and modelling assumptions, is explained following a general discussion of NFM. A baseline run of the model under flood conditions is discussed for acceptability in terms of local knowledge of patterns of flooding.
- 2) A GIS package is used to show different layers for the local catchment and bring in layers of potential opportunities that might be based on national strategic layers.⁴⁴ These are discussed with the participants and options and be switched on and off according to where the catchment partners identify where they would be happy to try different sorts of NFM.
- 3) The measures are then plugged into the model and the model is run to predict the outcome of the changed configuration.
- 4) The distributed changes to the hydrological responses are explored with the partners to understand model behaviour and effectiveness. Fig. 12 shows the outputs following one of the workshops.

Interestingly, in some discussions of significant measures in front of their peers, landowners came up with interventions that were very significant, for example sacrificing some summer irrigation storage to act as flood storage areas in the winter season. This process of standing around the tables and discussing the catchment with peers appeared to make people more forthcoming, and the process more effective. The process does, however, require a different approach to modelling since some of the feedback from local stakeholders might not be positive. It is therefore important that the modellers involved should not be too protective about their model, but should recognise and explain the assumptions and uncertainties inherent in the modelling process and be prepared to incorporate new knowledge as far as possible. Finding ways of conveying (and if necessary recalculating) prediction uncertainties within this context is the subject of on-going work.

7. Conclusions

This paper has carried out a systematic analysis of the technological readiness for the concept of models of everywhere. In particular, the paper has examined the various dimensions associated with models of everywhere and determined a set of technological requirements that must be met for the successful large-scale deployment of the concept. This set of requirements was then used to compare technological readiness when models of everywhere was first proposed against the readiness levels now, showing that the time is right for widespread experimentation and deployment of the concept. Although many of the technological barriers have been removed, key research issues remain and the paper has highlighted a set of open research questions that must be addressed before progress can be made. Importantly, this research agenda represents a shift in environmental modelling from an approach centred on process understanding (through deterministic models) to one that embraces a more data-centric perspective, whereby the two approaches can work in tandem to achieve a deeper understanding of specific places and hence to support more nuanced decision making about a given place.

There are also key limitations of the models of everywhere approach including the computational requirements if rolled out on a large scale and also the advances needed in environmental science to develop scale dependent parameterisations and embrace an underlying science of everything, everywhere and at all times. In conclusion, the concept of models of everywhere is even more important than when it was first proposed given the environmental challenges we face, and this paper has demonstrated that the time is right for more large-scale experimentation with the concept. Further research though is clearly needed to deliver against this vision and this research has to be fundamentally trans-disciplinary in nature bringing together environmental scientists, data scientists and computer scientists to reach a common understanding of representing complex environmental data, making sense of the resultant highly heterogeneous data, integrating knowledge from process and data models, and rolling out the concept of scale. Equally importantly, there is a need to work closely with social scientists to understand the human and societal issues related to models of everywhere, including the necessary cultural shift to open data, treatments of security and privacy and the role of communities in ensuring models represent the peculiarity of places.

Acknowledgements

This work is supported by EPSRC Grants: EP/P002285/1 and EP/ N027736/1 (Blair's Senior Fellowship on the Role of Digital Technology in Understanding, Mitigating and Adapting to Environmental Change, and Models in the Cloud: Generative Software Frameworks to Support the Execution of Environmental Models in the Cloud, respectively). We would also like to thank our colleagues in the Data Science Institute and the Centre for Ecology and Hydrology (CEH) for input and for providing such an inspiring environment for cross-disciplinary dialogue around data science for the natural environment. Finally, we thank the terrific inputs from a wide range of researchers and practitioners in partner organisations, including from JBA Trust, JBA Consulting, the Environment Agency, United Utilities, the European Centre For Medium Range Weather Forecasts, the Oasis Loss Modelling Framework, and the Environmental Change Institute (Oxford University).

References

- Ahmed, A., Ahmed, E., 2016. A survey on mobile edge computing. In: Proc. 10th IEEE International Conference on Intelligent Systems and Control (ISCO'16). https://doi. org/10.1109/ISCO.2016.7727082.
- Atzori, L., Iera, A., Morabito, G., 2010. The Internet of things: a survey. Comput. Network. 54 (15), 2787–2805. https://doi.org/10.1016/j.comnet.2010.05.010.
- Basco-Carrera, L., Warren, A., van Beek, E., Jonoski, A., Giardino, A., 2017. Collaborative modelling or participatory modelling? A framework for water resources management. Environ. Model. Softw 91, 95–110.
- Bastin, L., Cornford, D., Jones, R., Heuvelink, G.B.M., Pebesma, E., Stasch, C., Nativi, S., Mazzetti, P., Williams, M., 2013. Managing uncertainty in integrated environmental modelling: the UncertWeb framework. Environ. Model. Softw 39, 116–134. https:// doi.org/10.1016/j.envsoft.2012.02.008.
- Beven, K.J., 2006. A manifesto for the equifinality thesis. J. Hydrol. (Wellingt. North) 320, 18–36.
- Beven, K.J., 2007. Working towards integrated environmental models of everywhere: uncertainty, data, and modelling as a learning process. Hydrol. Earth Syst. Sci. 11 (1), 460–467.
- Beven, K.J., 2009. Environmental Modelling: an Uncertain Future? Routledge, London. Hb: 978-0-415-46302-7; Pb: 978-0-415-45759-0). http://www.uncertain-future.org.uk.
- Beven, K.J., 2016. EGU Leonardo Lecture: facets of Hydrology epistemic error, nonstationarity, likelihood, hypothesis testing, and communication. Hydrol. Sci. J. 61 (9), 1652–1665. https://doi.org/10.1080/02626667.2015.1031761.
- Beven, K.J., 2019. Towards a Methodology for Testing Models as Hypotheses in the Inexact Sciences. Proceedings Royal Society A (in press).
- Beven, K.J., Alcock, R., 2012. Modelling everything everythere: a new approach to decision making for water management under uncertainty. Freshw. Biol. 56, 124–132. https://doi.org/10.1111/j.1365-2427.2011.02592.x.
- Beven, K.J., Binley, A.M., 1992. The future of distributed models: model calibration and uncertainty prediction. Hydrol. Process. 6, 279–298.
- Beven, K.J., Binley, A.M., 2014. GLUE, 20 years on. Hydrol. Process. 28 (24), 5897–5918. https://doi.org/10.1002/hyp.10082.
- Beven, K.J., Cloke, H.L., 2012. Comment on "Hyperresolution global land surface modeling: Meeting a grand challenge for monitoring Earth's terrestrial water" by Eric F. Wood et al. Water Resour. Res. 48 (1).
- Beven, K.J., Cloke, H., Pappenberger, F., Lamb, R., Hunter, N., 2015. Hyperresolution information and hyperresolution ignorance in modelling the hydrology of the land surface. Sci. China Earth Sci. 58 (1), 25–35.
- Beven, K.J., Hall, J.W. (Eds.), 2014. Applied Uncertainty Estimation in Flood Risk Management. Imperial College Press, London.

⁴⁴ http://naturalprocesses.jbahosting.com/.

Beven, K.J., Lane, S., 2019. Invalidation of models and fitness-for-purpose: a rejectionist approach. In: Beisbart, C., Saam, N.J. (Eds.), Computer Simulation Validation -Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives. Springer, Cham, pp. 145–171 (Chapter 6) in.

- Beven, K.J., Smith, P.J., 2015. Concepts of information content and likelihood in parameter calibration for hydrological simulation models. ASCE J. Hydrol. Eng. https://doi.org/10.1061/(ASCE)HE.1943-5584.0000991.
- Beven, K., Smith, P.J., Wood, A., 2011. On the colour and spin of epistemic error (and what we might do about it). Hydrol. Earth Syst. Sci. 15, 3123–3133. https://doi.org/ 10.5194/hess-15-3123-2011.
- Bierkens, M.F., Bell, V.A., Burek, P., Chaney, N., Condon, L.E., David, C.H., de Roo, A., Döll, P., Drost, N., Famiglietti, J.S., Flörke, M., 2015. Hyper-resolution global hydrological modelling: what is next? "Everywhere and locally relevant". Hydrol. Process. 29 (2), 310–320.
- Blair, G.S., Henrys, P., Leeson, A., Watkins, J., Eastoe, E., Jarvis, S., Young, P.J., 2019. Data Science of the Natural Environment: a Research Roadmap (In preparation).
- Blei, D.M., 2014. Build, compute, critique, repeat: data analysis with latent variable models. Ann. Rev. Stat. Appl. 1, 203–232.
- Box, G., 1980. Sampling and Bayes' inference in scientific modeling and robustness. J. R. Stat. Soc. Ser. A 143 (4), 383–430.
- Clark, M.P., Nijssen, B., Lundquist, J.D., Kavetski, D., Rupp, D.E., Woods, R.A., Freer, J. E., Gutmann, E.D., Wood, A.W., Gochis, D.J., Rasmussen, R.M., 2015. A unified approach for process-based hydrologic modeling: 2. Model implementation and case studies. Water Resour. Res. 51 (4), 2515–2542.
- Coulouris, G., Dollimore, J., Kindberg, T., 2011. Distributed Systems: Concepts and Design, fifth ed. Pearson.
- Coxon, G., Freer, J., Westerberg, I.K., Wagener, T., Woods, R., Smith, P.J., 2015. A novel framework for discharge uncertainty quantification applied to 500 UK gauging stations. Water Resour. Res. 51 (7), 5531–5546.
- Dadson, S.J., Hall, J.W., Murgatroyd, A., Acreman, M., Bates, P., Beven, K.J., Heathwaite, L., Holden, J., Holman, I.P., Lane, S.N., O'Connell, P.E., Penning-Rowsell, E., Reynard, N., Sear, D., Thorne, C., Wilby, R., 2017. A restatement of the natural science evidence concerning catchment-based "natural" flood management in the United Kingdom. Proc. R. Soc. A 473, 20160706. https://doi.org/10. 1098/rspa.2016.0706.
- David, O., Ascough, J.C., Lloyd, W., Green, T.R., Rojas, K.W., Leavesley, G.H., Ahuja, L. R., 2013. A software engineering perspective on environmental modeling framework design: the object modeling system. Environ. Model. Softw 39, 201–213.
- Dean, J., Ghemawat, S., 2008. MapReduce: simplified data processing on large clusters. Commun. ACM 51 (1), 107–113. https://doi.org/10.1145/1327452.1327492.
- Di Baldassarre, G., Schumann, G., Bates, P.D., 2009. A technique for the calibration of hydraulic models using uncertain satellite observations of flood extent. J. Hydrol. 367 (3), 276–282.
- Di Baldassarre, G., Brandimarte, L., Beven, K.J., 2016. The seventh facet of uncertainty: wrong assumptions, unknowns and surprises in the dynamics of human-water systems. Hydrol. Sci. J. 61 (9), 1748–1758. https://doi.org/10.1080/ 02626667.2015.1091460.
- Edwards, L., Mullagh, L., Towe, R., Nundloll, V., Dean, C., Dean, G., Simm, W., Samreen, F., Bassett, R., Blair, G.S., 2017. Data-driven decisions for flood risk management. Proc. Data Policy 2017. http://doi.org/10.5281/zenodo.884180. htt p://dataforpolicy.org/papers/.
- Evers, M., Jonoski, A., Maksimovič, Č., Lange, L., Ochoa Rodriguez, S., Teklesadik, A., Cortes Arevalo, J., Almoradie, A., Eduardo Simões, N., Wang, L., Makropoulos, C., 2012. Collaborative modelling for active involvement of stakeholders in urban flood risk management. Nat. Hazards Earth Syst. Sci. 12 (9), 2821–2842.
- Fenicia, F., Kavetski, D., Savenije, H.H., 2011. Elements of a flexible approach for conceptual hydrological modeling: 1. Motivation and theoretical development. Water Resour. Res. 47, 11.
- Ferré, T.P., 2017. Revisiting the relationship between data, models, and decision-making. Gr. Water 55 (5), 604–614.
- Foster, I., Kesselman, C. (Eds.), 1998. The Grid: Blueprint for a New Computing
- Infrastructure. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. Foster, I., Kesselman, C., Tuecke, S., 2001. The Anatomy of the Grid: enabling scalable
- virtual organizations. Int. J. Supercomput. Appl. 15, 3. France, R., Rumpe, B., 2007. Model-driven development of complex software: a research
- roaden in: 2007 Future Of Software Engineering (FOSE '07). IEEE Computer Society, Washington, DC, USA, pp. 37–54. https://doi.org/10.1109/FOSE.2007.14.
- Gilbert, J.M., Maxwell, R.M., 2017. Examining regional groundwater-surface water dynamics using an integrated hydrologic model of the San Joaquin River basin. Hydrol. Earth Syst. Sci. 21 (2), 923.
- Habata, S., Yokokawa, M., Kitawaki, S., 2003. The earth simulator system. NEC Res. Dev. 44 (1), 21–26.
- Hankin, B., Metcalfe, P., Johnson, D., Chappell, N.A., Page, T., Craigen, I., Lamb, R., Beven, K.J., 2017. Strategies for testing the impact of natural flood risk management measures. In: Flood Risk Management. InTech Publishers. https://doi.org/10.5772/i ntechopen.68677.
- Hazeleger, W., Severijns, C., Semmler, T., Ştefănescu, S., Yang, S., Wang, X., Wyser, K., Dutra, E., Baldasano, J.M., Bintanja, R., Bougeault, P., 2010. EC-Earth: a seamless earth-system prediction approach in action. Bull. Am. Meteorol. Soc. 91 (10), 1357–1363.
- Hey, T., Tansley, S., Troll, K. (Eds.), 2009. The Fourth Paradigm. Data-Intensive Scientific Discovery, Microsoft Research.

- Hurrell, J.W., Holland, M.M., Gent, P.R., Ghan, S., Kay, J.E., Kushner, P.J., Lamarque, J. F., Large, W.G., Lawrence, D., Lindsay, K., Lipscomb, W.H., 2013. The community earth system model: a framework for collaborative research. Bull. Am. Meteorol. Soc. 94 (9), 1339–1360.
- Jamshidi, M. (Ed.), 2011. System of Systems Engineering: Innovations for the Twenty-First Century, vol. 58. John Wiley & Sons.
- Johnson, R.E., 1997. Frameworks = components+patterns. Commun. ACM 40 (10), 39-42. https://doi.org/10.1145/262793.262799.
- Kephart, J.O., Chess, D.M., 2003. The vision of autonomic computing. Computer 36 (1), 41–50. https://doi.org/10.1109/MC.2003.1160055.
- Kon, F., Costa, F., Blair, G.S., Campbell, R.H., 2002. The case for reflective middleware. Commun. ACM 45 (6), 33–38.
- Landström, C., Whatmore, S.J., Lane, S.N., Odoni, N.A., Ward, N., Bradley, S., 2011. Coproducing flood risk knowledge: redistributing expertise in critical 'participatory modelling'. Environ. Plan. 43 (7), 1617–1633.
- Lane, S.N., Odoni, N., Landström, C., Whatmore, S.J., Ward, N., Bradley, S., 2011. Doing flood risk science differently: an experiment in radical scientific method. Trans. Inst. Br. Geogr. 36 (1), 15–36.
- Lahoz, W., Khattatov, B., Menard, R. (Eds.), 2010. Data Assimilation: Making Sense of Observations. Springer Science & Business Media.
- Leavesley, G.H., Markstrom, S.L., Restrepo, P.J., Viger, R.J., 2002. A modular approach for addressing model design, scale, and parameters. Hydrol. Process. 16, 173–187.
- Lopez, P.G., Montresor, A., Epema, D., Datta, A., Higashino, T., Iamnitchi, A., Barcellos, M., Felber, P., Riviere, E., 2015. Edge-centric computing: vision and challenges. SIGCOMM Comput. Commun. Rev. 45 (5), 37–42. https://doi.org/1 0.1145/2831347.2831354.
- Lloyd, W., David, O., Ascough, J.C., Rojas, K.W., Carlson, J.R., Leavesley, G.H., Krause, P., Green, T.R., Ahuja, L.R., 2011. Environmental modeling framework invasiveness: analysis and implications. Environ. Model. Softw 26 (10), 1240–1250.
- McCallum, I., Liu, W., See, L., et al., 2016. Technologies to support community flood disaster risk reduction. Int. J. Disaster Risk Sci. 7, 198. https://doi.org/10.100 7/s13753-016-0086-5.
- McDonnell, J.J., Beven, K.J., 2014. Debates—the future of hydrological sciences: a (common) path forward? A call to action aimed at understanding velocities, celerities, and residence time distributions of the headwater hydrograph. Water Resour. Res. 50 https://doi.org/10.1002/2013WR015141.
- McKinley, P.K., Sadjadi, S.M., Kasten, E.P., Cheng, B.H.C., 2004. Composing adaptive software. Computer 37 (7), 56–64. https://doi.org/10.1109/MC.2004.48.
- Maskrey, S.A., Mount, N.J., Thorne, C.R., Dryden, I., 2016. Participatory modelling for stakeholder involvement in the development of flood risk management intervention options. Environ. Model. Softw 82, 275–294.
- Maes, P., 1987. Concepts and experiments in computational reflection. In: Meyrowitz, Norman (Ed.), Conference Proceedings on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '87). ACM, New York, NY, USA, pp. 147–155.
- Metcalfe, P., Beven, K.J., Hankin, B., Lamb, R., 2017. A modelling framework for evaluation of the hydrological impacts of nature-based approaches to flood risk management, with application to in-channel interventions across a 29-km² scale catchment in the United Kingdom. Hydrol. Process. https://doi.org/10.1002/ hyp.11140.
- Nearing, G.S., Tian, Y., Gupta, H.V., Clark, M.P., Harrison, K.W., Weijs, S.V., 2016. A philosophical basis for hydrological uncertainty. Hydrol. Sci. J. 61 (9), 1666–1678.
- Nundloll, V., Porter, B., Blair, G.S., Emmett, B., Cosby, J., Jones, D., Chadwick, D., Winterbourn, B., Beattie, P., Shaw, R., Shelley, W., Ullah, I., Brown, M., Dean, G., 2019. The design and deployment of an end-to-end IoT infrastructure for the natural environment. Future Internet 11 (6), 129. https://doi.org/10.3390/fi11060129.
- Oreizy, P., Gorlick, M.M., Taylor, R.N., Heimbigner, D., Johnson, H., Medvidovic, N., Quilici, A., Rosenblum, D.S., Wolf, A.L., 1999. An architecture-based approach to self-adaptive software. IEEE Intell. Syst. 14 (3), 54–62. https://doi.org/10.1109/52 54.769885.
- Park, S.K., Xu, L. (Eds.), 2017. Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications, vol. 3. Springer Science & Business Media.
- Prudhomme, C., Wilby, R.L., Crooks, S., Kay, A.L., Reynard, N.S., 2010. Scenario-neutral approach to climate change impact studies: application to flood risk. J. Hydrol. 390 (3), 198–209.
- Renard, B., Kavetski, D., Kuczera, G., Thyer, M., Franks, S.W., 2010. Understanding predictive uncertainty in hydrologic modeling: the challenge of identifying input and structural errors. Water Resour. Res. 46, 5.
- Rougier, J., Sparks, S., Hill, L., 2013. Risk and Uncertainty Assessment for Natural Hazards. Cambridge University Press, Cambridge, UK.
- Simm, W.A., Samreen, F., Bassett, R., Ferrario, M.A.F.C., Blair, G.S., Young, P.J., Whittle, J.N., 2018. SE in ES: opportunities for software engineering and cloud computing in environmental science. In: ICSE-SEIS '18. ACM 10, New York, ISBN 9781450356619.
- Smith, L., Liang, Q., James, P., Lin, W., 2015. Assessing the utility of social media as a data source for flood risk management using a real-time modelling framework. J. Flood Risk Manag. 10, 370–380. https://doi.org/10.1111/jfr3.12154.
- Smith, P.J., Hughes, D., Beven, K.J., Cross, P., Tych, W., Coulson, G., Blair, G., 2009. Towards the provision of site specific flood warnings using wireless sensor networks. Meteorol. Appl. 16, 57–64.
- Towe, R., Dean, G., Edwards, L., Nundloll, V., Blair, G.S., Lamb, R., Hankin, B., Manson, So, 2019. Rethinking Data-Driven Decision Support in Flood Risk Management for a Big Data Age (submitted for publication).

G.S. Blair et al.

- Voinov, A., Kolagani, N., McCall, M.K., Glynn, P.D., Kragt, M.E., Ostermann, F.O., Pierce, S.A., Ramu, P., 2016. Modelling with stakeholders–next generation. Environ. Model. Softw 77, 196–220.
- Vrugt, J.A., Sadegh, M., 2013. Toward diagnostic model calibration and evaluation: approximate Bayesian computation. Water Resour. Res. 49 (7), 4335–4345.
- Waldrop, M.M., 2008. Science 2.0 is open access science the future? Sci. Am. 298.Weiler, M., Beven, K.J., 2015. Should there be a community model? Water Resour. Res. 51 (9), 7777–7784.
- Westerberg, I.K., McMillan, H.K., 2015. Uncertainty in hydrological signatures. Hydrol. Earth Syst. Sci. 19 (9), 3951–3968.
- Westerberg, I., Guerrero, J.-L., Seibert, J., Beven, K.J., Halldin, S., 2011. Stage-discharge uncertainty derived with a non-stationary rating curve in the Choluteca River, Honduras. Hydrol. Process. 25, 603–613. https://doi.org/10.1002/hyp.7848.
- Wood, E.F., Roundy, J.K., Troy, T.J., Van Beek, L.P.H., Bierkens, M.F., Blyth, E., de Roo, A., Döll, P., Ek, M., Famiglietti, J., Gochis, D., 2011. Hyperresolution global land surface modeling: meeting a grand challenge for monitoring Earth's terrestrial water. Water Resour. Res. 47 (5).