

POSITION PAPER

Risk management in the era of data-centric engineering

Domenic Di Francesco^{1,2,3,4} 

¹The Alan Turing Institute, London, UK

²Computational Statistics and Machine Learning Research Group, Civil Engineering, Cambridge University, Cambridge, UK

³Science Division, Health and Safety Executive, Buxton, UK (Secondment)

⁴AQ Live, London, UK

Email: ddifrancesco@turing.ac.uk

Received: 28 August 2024; **Revised:** 23 January 2025; **Accepted:** 28 January 2025

Keywords: assurance; computational statistics; machine learning; risk management

Abstract

Novel methods of data collection and analysis can enhance traditional risk management practices that rely on expert engineering judgment and established safety records, specifically when key conditions are met: Analysis is linked to the decisions it is intended to support, standards and competencies remain up to date, and assurance and verification activities are performed. This article elaborates on these conditions. The reason engineers are required to perform calculations is to support decision-making. Since humans are famously weak natural statisticians, rather than ask stakeholders to *implicitly* assimilate data, and arrive at a decision, we can instead rely on subject matter experts to *explicitly* define risk management decision problems. The results of engineering calculation can then also communicate which interventions (if any) are considered to be risk-optimal. It is also proposed that the next generation of engineering standards should learn from the success of open source software development in community building. Interacting with open datasets and code can promote engagement, identification (and resolution) of errors, training and ultimately competence. Finally, the profession's tradition of independent verification should also be applied to the complex models that will increasingly contribute to the safety of the built environment. Model assurance will be required to keep pace with model development to identify suitable use cases as adequately safe. These are considered to be increasingly important components in ensuring that methods of data-centric engineering can be safely and appropriately adopted in industry.

Impact Statement

How should engineers engage with so-called artificial intelligence and data-driven technologies? Is skepticism warranted? Why intervene with established, safe ways of working? This article intends to highlight why practicing engineers should give some attention to novel methods of data collection and analysis, while ensuring that these approaches should be explicitly linked to how they supports decision making.

Emerging new challenges and constraints on our engineering systems will require new methods. As a profession, engineering has an established history of solving complex problems by investigating increasingly complex approaches, and pragmatically adopting them once they are considered trustworthy. In this article, this case is made, and three key propositions are introduced regarding: decision-making, the next generation of standards and competency frameworks, and technology assurance.

1. Introduction

1.1. What is risk?

When a future outcome is uncertain, rather than resign ourselves to an unknowable fate, we use the information and judgment that is available to us to consider what might happen. We must do this to avoid so-called *decision paralysis* and live our lives. This is either done explicitly (using documented calculations) or implicitly (by arriving at a decision that is consistent with some unspecified model of the problem).

Risk is sometimes not considered distinct from probability, in the context of an uncertain event occurring (World Economic Forum, 2023). Some definitions extend this and consider risk to be somehow related to both the probability and consequences of uncertain events (Health and Safety Executive [HSE], n.d.), but many engineering standards explicitly define risk as the product of probability and consequence (American Petroleum Institute, 2016; DNV, 2019), that is, risk is the expected consequences of uncertain events. This definition provides a rationale to help rank-order decision alternatives, by selecting the option that is expected to benefit us the most (or cost us the least) (Von Neumann and Morgenstern, 1953). What is not conveyed in this metric is the workflow of checking diagnostics, performing evaluation, and verification of the model(s), or the sensitivity analysis required to define a representative utility function that adequately describes the decision makers subjective judgment. (It is therefore challenging to provide specific guidance on how to choose among decision alternatives, and trust that if an organization is at least considering probabilities and outcomes in some capacity, a sensible decision will be made. However, even attempting to formalize this process encourages us to critically evaluate the justification for decisions, and provides a framework for proposing specific improvements to models and utility functions. Most importantly, the decision will be fully documented as a result, so that it can be reevaluated at any time.)

Since we need to act under uncertainty so frequently, there are many examples of personal risk management. For instance, people choose to pay for insurance because of the possibility of large losses, and although few of us do any formal calculations to justify making this purchase (though decision-theorist Howard Raiffa did claim to do full analysis to identify which job offer he should accept; Raiffa, 2002), there would be a certain threshold above which we would not be happy to pay. Similarly, in a professional context, each time an organization makes a decision to invest (or not) in a safety measure, they are implicitly assigning utilities to the possible outcomes. These decisions are generally never interrogated or the subjective judgments back-calculated, except in instances where it is clear that an error has been made. This point is similarly, but more forcefully, made in the study by Hopkins (2002), regarding pipeline engineering:

you don't need anything at all! You don't need qualified engineers, you don't need quality systems, you don't need risk management, you don't need safety audits, you don't need inspections, you don't need training. You don't need anything! Until something happens ... then you need everything.

Risk analysis has an inherent link between probabilities (uncertainty quantification) and outcomes (causal inference, optimization, and decision analysis). The fundamental risk management problem can be represented graphically, as shown in Figure 1, which is sometimes referred to as an *influence diagram*. Here, the decision of what, if any, investment is required is shown as a square, the associated implementation costs are diamond-shaped, and the results of a relevant analysis are a circle. The arrow from the decision to the implementation costs indicates that these costs directly depend on the risk management plan. To facilitate quantitative risk management, we would need an understanding of what is at stake if any aspects of our system stop functioning (the arrow from the system modeling to outcome costs) and the effect of our actions on the probability of these outcomes occurring (the arrow from the decision to the modeling), so that we can coherently manage this trade-off. However, the relationship between engineering calculations and decisions is not always evident. Given that the need to support a decision is generally the reason the analysis is commissioned, the results should have clear implications for which actions we should then take.

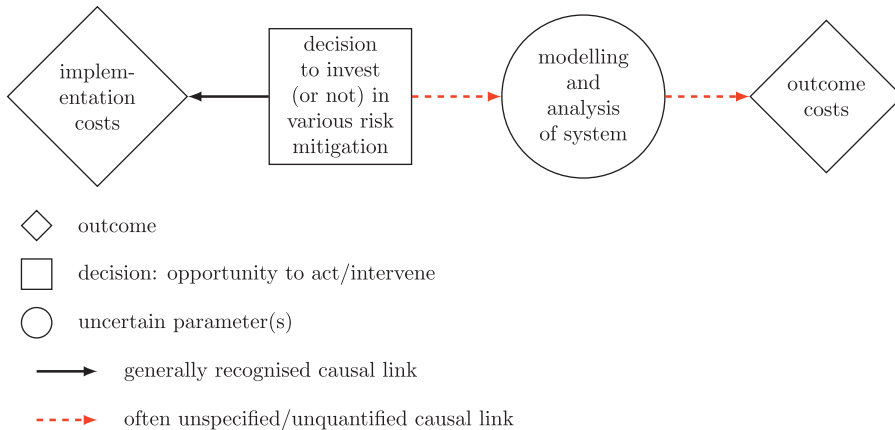


Figure 1. Influence diagram representation of risk management challenge: these components need to be connected.

There are various benefits to tying analysis and decisions, such as maintaining consistent objectives, improved transparency in outcomes, providing a framework for collecting (structural) subject matter expertise, and finding risk-optimal strategies in a quantitative and replicable way. This principle is also advocated for in Professor Jordaan’s engineering textbook (Jordaan, 2005):

It is important that all problems of inference be visualized as problems of decision.

In the UK, the Health & Safety Executive (HSE) regulate many engineering industries but tend not to provide prescriptive guidance. Rather, they require risk to be kept *as low as reasonably practicable* (ALARP; HSE, 2001). This term was introduced in the ruling of a 1949 court case, when, after an individual died after the collapse of a coal mine road, it was concluded that a risk assessment could be used to identify which roads required strengthening and which did not. This outcome continues to imply that there are risks we should be willing to tolerate and that some calculations are required to demonstrate that risk mitigation is being implemented where and when it is worthwhile. The lack of a coherent connection between these calculations and interventions (as proposed in Figure 1) makes it more challenging for courts, regulators, and internal audits to evaluate and feedback.

Proposition 1. Connecting engineering analysis and decisions can allow for more coherent, transparent, and auditable risk management.

1.2. What is data-centric engineering?

In principle, the idea of using data to support engineering decisions is nothing new. Engineers have historically had to rely on pragmatism to design, build, and manage physical assets long before the development of detailed simulation or digital twins. (The historical willingness of (in particular British) engineers to bypass what they considered to be overly theoretic approaches to characterizing stress and strain is discussed in Professor James E. Gordon’s popular books (Gordon, 1991a, 1991b).) Rather, domain knowledge has always been (suboptimally) integrated with data from experiments to prevent failures.

However, in practice, DCE is the term used to describe the novel methods of collecting and analyzing data that are now available to engineers and organizations. Analysis of *big data* is perhaps the most widely publicized example of this. The successes of machine learning (ML) and artificial intelligence (AI) algorithms in analyzing (generally unstructured) data is rightly considered a breakthrough in

computational statistics and machine learning (CSML), but engineering can benefit from various new approaches:

1.2.1. Predictive modeling

Intended to identify relationships between quantities for the purposes of understanding plausible outcomes. Models based on deep neural networks (neural networks are models with large numbers of [nonlinearly] interacting parameters, organized in layers) have been shown to be capable at learning nonlinear relationships from large amounts of historical data and generalize well to unseen data from the same data generating process (interpolation rather than extrapolation; Prince, 2022). They are seemingly uniquely effective in unstructured data settings, such as images, audio or text, for which traditional regression analysis appears to be of limited use. Another application is as *surrogate models*. For instance, neural operators (a type of neural network that works with functions as inputs/outputs, rather than traditional fixed-length data) can be used to solve partial differential equations, once adequately trained, orders of magnitude quicker than conventional numerical methods (Kovachki et al., 2022). Typical neural network approaches do not provide much insight on why a specific outcome was predicted, and without sophisticated verification and assurance methods, alternative modeling approaches may be more appropriate in settings where subsequent decisions must be clearly explained or justified.

1.2.2. Probabilistic inference (uncertainty quantification)

Quantifying uncertainty using probability allows us to describe the information content of data more accurately and understand where and when a model may be less informed/confident in prediction. As noted to in Section 1.1, this is an important component of risk management. Engineering data, in particular, often consist of an indirect measurement of a complex physical phenomena, in challenging conditions. Similarly, deviations of real systems from idealized models can also be modeled using this approach. A key aspect in probabilistic modeling is inference (sometimes referred to as inverse probability) where models can be updated so that unknown quantities can be estimated consistently with measurements. Describing the probabilistic relationship between quantities, so that it can inform predictive models, is readily achievable using probabilistic programming languages (Ghahramani, 2015), which have helped democratize scalable uncertainty quantification (Walia et al., 2019). Note that some methods of quantifying uncertainty are also available to neural networks (Abdar et al., 2021; Murphy, 2022).

1.2.3. Optimizing decisions

Concerned with finding the best sequence of actions/interventions within a dynamic environment to achieve a specific goal. Many decision problems subject to uncertainty can be naturally represented as Markov decision processes, including their various associated operational and financial constraints, so that expected optimal strategies can be identified. Reinforcement learning (Sutton and Barto, 2020) methods seek to learn a strategy or policy that maximizes the cumulative reward over time. Extensions which optimize decisions under partial information provide a powerful framework for reasoning under uncertainty (Kochenderfer, 2015).

1.2.4. Generative AI

Typically comprising large-scale neural networks, these models are able to generate highly sophisticated image, text, or other outputs based on prompts (inputs as instructions in natural language). These have been made available to the public in recent years, and several use cases have emerged. For instance, GitHub copilot provides suggested code as software developers are working, and can be interacted with, via a chatbot. An initial report (admittedly from GitHub) indicated that users of this generative model assistance saw increases in productivity and satisfaction (GitHub, 2022). It is conceivable that similar assistance could be utilized by engineers when working on calculations, checking compliance with standards, or report writing, provided they were aware of the limitations

of such a tool. Generative AI forms part of an emerging trend of foundation models which are trained on a large corpus of text or image data and which can then be fine-tuned for a wide range of different tasks. Unlike models trained for a single purpose, foundational models learn a broad representation of knowledge from their training data, which allows them to adapt to various applications with additional, task-specific training.

Proposition 2. The adoption of novel methods of data collection and analysis will benefit from a new generation of engineering standards and training materials.

At their core, these methods typically require the calculation of gradients of highly complex functions (algorithmically; Baydin et al., 2018) and the adjustment of probability distributions or parameter estimates in a direction that is increasingly consistent with the evidence provided. Modern computer software and hardware has allowed these approaches to be applied at new scales, meaning that the corresponding matrix and vector operations become too large to track conceptually, and ML models are therefore often considered to be black box (for the purposes of this article, a *black box* model is considered to mean the specifics of the internal operations, which transform inputs to outputs, are too complex (or otherwise obfuscated) to be interpreted and understood by humans) in nature. Understanding where such complex models can safely adopted in industry therefore requires careful thought, but this may not be an entirely new challenge for engineers.

1.3. Comparison with finite element analysis

Engineers routinely use complex computational methods that were once considered novel approaches but ultimately demonstrated value and became standard practice. The finite element method is one such example (of many) that has become an established method with multiple software packages available for running analysis. Much of the below discussion may be equally valid for alternative existing complex methods and tools.

Finite element analysis discretizes complex geometries so that problems can be represented as systems of linear equations. This has many parallels with the way that large numbers of interacting neurons in a neural network are also essentially related by linear algebra operations. In finite element analysis, displacements at nodes are calculated by (finding an approximate solution, (often using gradient-based methods) to inverting a stiffness matrix, and then performing some matrix–vector multiplication. In a neural network gradient-based methods are also iteratively used to find a solution, and at any stage the outputs are found by performing some matrix–vector multiplications.

Although finite element models are applying physical laws, and neural networks are applying statistical functions, the inner workings of both methods can be considered as black boxes, particularly for large problems. Both typically produce interpretable results, but their reliability should be critically considered. One (of many) example of a round-robin trial, in which an identical stress analysis is posed to multiple organizations is documented in the study by Kluess et al. (2019). Here, seven independent analyses of a human femur bone subject to a point load were found to differ greatly from each other, and from experimental data. The largest average (across 10 different strain gauge locations) deviation in predicted strain from the test data was 354%. The magnitude of these errors did not necessarily reduce with more elements, or longer computation in this trial.

Looking forward, more can be done to understand and quantify the sources of variation, so finite element analysis can more directly inform risk management, and condition on measurement data, and this work has begun (Girolami et al., 2021). One such application would be to understand the information content simulated data from such a model, to better understand where supplementary physical testing is required in the design or certification of a new component. Decision-theoretic approaches have been demonstrated for quantifying the expected value of engineering data (Di Francesco et al., 2021), which represents the formal statistical solution to a challenge, about which discussion remains vague in industry. For instance, as stated in the study by Johnston (2017):

There is often a balance to be struck between the number of tests that would give the result with highest level of statistical confidence and the cost implications of producing many full-scale test specimens.

Despite an appreciation the potential challenges and uncertainties associated with this computational modeling technique, finite element analysis remains commonplace in engineering projects. Quality assurance procedures, model verification, and subject matter knowledge are used to ensure it is used safely (DNV, 2021), and a similar approach can be adopted for the safe adoption of methods of data-centric engineering. Recently, the first independent verification of a digital twin in the maritime industry was completed (Lloyds Regsiter, n.d.), and the evidence from this analysis was sufficient to gain regulatory approval. As standardized guidance on the assurance of such models is developed (existing (early) guidance in this area is encouraging; DNV, 2023, but is generally confined to higher-level recommendations, rather than methods to quantify and evaluate model reliability in the context of functional safety), such projects will become more acceptable. There may be an opportunity for initially only permitting the use of complex models in lower consequence domains, where learning and testing will be more tolerable. However, these may not be the applications that are hoped to benefit the most in safety and economically.

Adopting a complex modeling approach introduces a risk (*model risk*) to a business, namely the expected consequences of an incorrect or otherwise unhelpful output. To demonstrate that this model risk is ALARP, it needs to be weighed against the expected benefit(s) of using the model, accounting for the non-zero risks associated with current ways of working.

Proposition 3. There is some (valid) skepticism of new approaches in established engineering disciplines. Methods of verification and assurance should be developed (and eventually standardized) to help ensure complex models are adopted safely.

2. In practice: a duty to do better versus if it is not broken ...

Thankfully, buildings very rarely collapse, planes very rarely crash, pressure vessels very rarely explode (Melchers and Beck, 2018). However, these same standards of safety will increasingly be subject to new constraints and pressures, for instance automation and climate change are anticipated to shape the future of maritime trade in the coming decades (Economist, 2023). It is nevertheless reasonable to ask why we should consider changing existing practices that generally keep us safe. In fact, more complex analysis may be even worse than *not required*, it may actually introduce new risks.

Some skepticism is helpful when evaluating new solutions, and this should drive the requirements to demonstrate safety. However, excessive resistance to change can prevent progress, and data-centric engineering offers the promise of improved risk management. Our current reliance on heuristics does not require engineers to quantify risk, as noted in a review paper (Hadley and Wu, 2018) on a widely used standard for assessing damage (BSI, 2019):

the principle behind the (BS 7910) procedure has long been avoidance, not prediction, of failure.

Without consistent quantification of risk and reliability, we are unable to justify investments in resource allocation (safety!) in a principled way, and this leaves both private and public spending open to being influenced by undesirable factors. For instance, spending on safety measures in air travel following the attacks in New York on 11 September 2001 is evaluated in the study by Stewart and Mueller (2018). The enormous increase in spending could not be justified by any of the quantitative evaluation methods that the authors considered, suggesting that much of this policy was introduced for political benefit.

Even in such cases where a failure to quantify risk has resulted in an overspend on safety, this is still considered suboptimal. Investment in safety measures that do not improve safety is money that could have been better spent. In such cases, a utility function may describe the various negative outcomes (financial costs, risks to personnel associated with the activity, greenhouse gas emissions associated with new

construction) which outweigh subsequent safety benefits in a formal analysis. Of course, conversely, political distortions that tend towards underspending in safety measures are more directly problematic.

If an expert disagrees with an output from such a tool, this could indicate that the underlying model may be failing to account for some important factor, or domain knowledge. This should lead to a review and possible targeted improvements to the influence diagram, since there is a baseline structure to amend, as opposed to a less structured discussion. Alternatively, disagreement with a result may instead indicate that a new strategy has been correctly identified, which may have previously been considered somehow counter-intuitive.

Would UK water companies recent (BBC News, 2017) (and continued!; NewScientist, 2023) use of divining rods (dowsing) to locate buried water pipelines stand up to the scrutiny of a formal decision analysis? Certainly, the factors driving their perceived utility would have to be made explicit and, therefore, be more straightforward to challenge. This is not to be considered a particularly consequential example of the absence of DCE. Much greater costs are involved in the HS2 rail project, which has been suggested to have lacked (among other things) transparency, sufficient analysis at early stages, and consideration of alternative transport plans (Tetlow and Shearer, 2021). Current projects do not only suffer from outdated or absent methods, but also from the software that is used. Too much engineering still takes place in spreadsheets. Earlier this year, a miss-typed date in a spreadsheet resulted in Norway's sovereign wealth fund losing approximately £ 70,000,000 (Financial Times, 2024). Tens of thousands of Covid-19 cases were not reported in the UK because an outdated file format was selected, which did not make it clear that the total number of rows of data had been exceeded. Following this investigation, Professor Jon Crowcroft suggested that spreadsheets were not suitable for analysis with high-contour applications (BBC News, 2020):

Excel was always meant for people mucking around with a bunch of data for their small company to see what it looked like ...

... and then when you need to do something more serious, you build something bespoke that works —there's dozens of other things you could do.

Perhaps most famously, economics professors from Harvard University were found to have not selected all relevant cells in their calculation of average GDP growth. This was only identified when a PhD student was attempting to reproduce the analysis, and yet this influential paper supported austerity policies being introduced in the US and Europe (BBC News, 2013). There are many more examples on the *Horror Stories* page on the website of the European Spreadsheet Risk Interests Group (EuSpRiG) (EuSpRiG, n.d.), and Monash University (Caprani, n.d.), which promotes useful practices from modern software development, recorded a discussion on the risks of using spreadsheets in an engineering context. The Alan Turing Institute has developed a best-practice guide for performing research and analysis and is freely available (The Turing Way Community, 2022).

Engineering remains a profession trusted by the public (The Royal Academy of Engineering, 2022), in part due to the focus on ethics in our institutions (The Engineering Council, 2017). These principles include a commitment to keep skills and knowledge up-to-date. This will require engineers to engage with topics of computational statistics and ML, to identify where the associated methods and tools can contribute to improved risk management in our organizations.

3. In summary

The following key propositions detail how some anticipated industrial challenges could be addressed.

- **Connecting analysis and decisions**

If a stress analysis is being completed because a design decision needs to be made, then the optimal decision (conditional on the model) should be identified, and explained as part of the result. More generally, our calculations should not be disconnected from the reasons we have been asked to do them. Furthermore, risk management decisions should be documented so that during future evaluations or

audits, it is clear why a decision was arrived at. This can be achieved by using methods of CSML to transition from implicit data assimilation (that effectively happens when we ask stakeholders to arrive at a decision based on various reports and calculations) to explicit, reproducible analysis.

- **DCE standards and competencies**

Engineers are generally willing to maintain their technical competencies, and this is reflected in societies continued trust in the profession. The development of specific industry guidance and new standards in the topic of DCE, such as is underway at the AI standards hub (The Alan Turing Institute, n.d.) may need to be accompanied by accessible training content. Such standards may also benefit from adopting ideas from open-source software development, so that any datasets and underlying analysis is made public. Building communities with access to the calculations that justify the guidance may help identify any errors more quickly and improve the competence of the users.

- **Model assurance and verification needs to keep up**

Many engineering industries have long histories of independent verification of mechanical/structural design documents and calculations. As the safety of our assets becomes tied to increasingly complex models, these too should be subject to same practices. Regulation could initially direct the use (and assurance/verification) of complex models such as digital twins towards lower consequence applications. This may counteract the understandable resistance to methods of DCE in various fields of engineering.

Transition to methods of DCE can have a transformative, (net) positive impact on risk management. This could allow us to make quantitative business cases for infrastructure investment which may be less susceptible to misinterpretation and organizational/political corruption. The historic achievements of engineers are self-evident, but they were operating with what will soon be considered to be great limitations. As new methods of data collection and analysis become increasingly available, we will be capable of quantifying uncertainty over large engineering systems to coherently support risk management decisions. Performing this analysis using well-documented and reproducible workflows can help direct further innovation. As we progress along this path, reliance on outdated tools (such as spreadsheets) and methods (such as simplified rules and heuristics) will become increasingly indefensible in cases where alternative approaches are risk-optimal.

Data availability statement. Not applicable, since no new data has been generated as part of this position paper.

Acknowledgments. I would like to thank the following experts for kindly agreeing to review this article and provide feedback, in no particular order: Professor Adam Sobey, Professor Chris Dent, Dr. Andrew Duncan, Dr. Rozelle Kane, Justin Bunker, and Tom Savage. I would also like to thank the journal reviewers for their helpful suggestions.

Author contribution. Conceptualization, Resources, and Writing original draft: D.D.F.

Funding statement. Domenic Di Francesco is supported by the Ecosystem Leadership Award under the EPSRC Grant EP/X03870X/1, and The Alan Turing Institute, particularly the Turing Research Fellowship scheme under that grant.

Competing interest. None.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Abdar M, Pourpanah F, Hussain S, Rezadegan D, Liu L, Ghavamzadeh M, Fieguth P, Cao X, Khosravi A, Acharya UR, Makarenkov V and Nahavandi S (2021) A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion* 76, 243–297. <https://doi.org/10.1016/j.inffus.2021.05.008>.
- American Petroleum Institute (2016) *Risk-Based Inspection, API RP 580*, 3rd edn. Washington, DC: American Petroleum Institute.
- Baydin AG, Pearlmutter BA, Radul AA and Siskind JM (2018) Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research* 18, 1–43.
- BBC News (2013) Reinhart, Rogoff ... and Herndon: The student who caught out the profs. Available at <https://www.bbc.com/news/magazine-22223190>.

- BBC News** (2017) Scientist finds UK water companies use ‘magic’ to find leaks. Available at <https://www.bbc.co.uk/news/uk-england-oxfordshire-42070719> (accessed 15 February 2024)
- BBC News** (2020) Excel: Why using Microsoft’s tool caused Covid-19 results to be lost. Available at <http://bbc.com/news/technology-54423988>.
- BSI** (2019) *BS 7910:2019 Guide to Methods for Assessing the Acceptability of Flaws in Metallic Structures*. London: British Standards Institution.
- Caprani C** (n.d.) Excel is evil – Why it has no place in research. Available at <https://www.youtube.com/watch?v=NuTlcZV72Q> (accessed 16 February 2024).
- Di Francesco D, Chryssanthopoulos M, Faber MH and Bharadwaj U** (2021) Decision-theoretic inspection planning using imperfect and incomplete data. *Data Centric Engineering* 2(18). <https://doi.org/10.1017/dce.2021.18>.
- DNV** (2019) Risk Management in Marine and Subsea Operations. Tech. Report. Available at <https://standards.dnv.com>.
- DNV** (2021) DNV-SE-0474 Risk-Based Verification. Tech. Report. Available at <https://standards.dnv.com/>.
- DNV** (2023) DNV-RP-A204 Assurance of Digital Twins. Tech. Report. Available at <https://standards.dnv.com/>.
- Economist** (2023) Global Maritime Trends 2050. Tech. Report. Available at <https://impact.economist.com/ocean/global-maritime-trends-2050/>.
- EuSpRiG** (n.d.) European Spreadsheet Risks Interest Group. Available at <https://eusprig.org> (accessed 16 February 2024).
- Financial Times** (2024) The Norwegian sovereign wealth fund’s 92mn dollar Excel error. Available at [The Norwegian sovereign wealth fund’s \\$92mn Excel error](https://www.ft.com/content/2024/02/16/norwegian-sovereign-wealth-fund-excel-error).
- Ghahramani Z** (2015) Probabilistic machine learning and artificial intelligence. *Nature* 521(7553), 452–459. <https://doi.org/10.1038/nature14541>.
- Girolami M, Febrianto E, Yin G and Cirak F** (2021) The statistical finite element method (statFEM) for coherent synthesis of observation data and model predictions. *Computer Methods in Applied Mechanics and Engineering* 375. <https://doi.org/10.1016/j.cma.2020.113533>.
- GitHub** (2022) Research: quantifying GitHub Copilot’s impact on developer productivity and happiness. Available at <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivityand-happiness/> (accessed 1 November 2023).
- Gordon JE** (1991a) *Structures or Why Things Don’t Fall Down*. London: Penguin Books.
- Gordon JE** (1991b) *The New Science of Strong Materials or Why You Don’t Fall Through the Floor*. London: Penguin Books.
- Hadley I and Wu G** (2018) Treatment of reliability in fracture assessments using BS 7910 Annex K. *International Journal of Pressure Vessels and Piping* 168(November), 310–322. <https://doi.org/10.1016/j.ijpvp.2018.11.006>.
- Health and Safety Executive (HSE)** (2001) Reducing Risks, Protecting People. Tech. Report. Available at <https://www.hse.gov.uk/enforce/expert/r2p2.htm>.
- Health and Safety Executive (HSE)** (n.d.) ALARP “at a glance.” Available at <https://www.hse.gov.uk/enforce/expert/alarplance.htm>.
- Hopkins P** (2002) Training Engineers in Pipeline Integrity. Paper presented at Western Regional Gas Conference, Tempe, AZ, USA.
- Johnston C** (2017) Statistical Analysis of Fatigue Test Data. ASME 2017 36th International Conference on Ocean, Offshore and Arctic Engineering, OMAE2017. <https://doi.org/10.2472/jms.29.24>.
- Jordaán I** (2005) *Decisions under Uncertainty*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511804861>.
- Klues D, Soodmand E, Lorenz A, Pahr D, Schwarze M, Cichon R, Varady PA, Herrmann S, Buchmeier B, Schröder C, Lehner S and Keibach M** (2019) A round-robin finite element analysis of human femur mechanics between seven participating laboratories with experimental validation. *Computer Methods in Biomechanics and Biomedical Engineering* 22(12), 1020–1031. <https://doi.org/10.1080/10255842.2019.1615481>.
- Kochenderfer MJ** (2015) *Decision Making Under Uncertainty*. Cambridge: The MIT Press. Available at <https://mykel.kochenderfer.com/textbooks/>.
- Kovachki N, Li Z, Liu B, Azizadenesheli K, Bhattacharya K, Stuart A and Anandkumar A** (2022) Neural Operator: Learning maps between function spaces. *Journal of Machine Learning Research* 23, 1–97. <https://doi.org/10.5555/3648699.3648788>.
- Lloyds Register** (n.d.) LR award ‘Digital Twin Approved Certification’ to Furuno for HermAce VDR. Available at [https://www.lr.org/en/about-us/press-room/press-listing/press-release/lr-award-digital-twin-approved-certification-to-furuno-for-hermace-vdr/#:\\$%5Csim\\$:text=%E2%80%99CLR%20is%20delighted%20to%20award,technology%20across%20the%20maritime%20industry](https://www.lr.org/en/about-us/press-room/press-listing/press-release/lr-award-digital-twin-approved-certification-to-furuno-for-hermace-vdr/#:$%5Csim$:text=%E2%80%99CLR%20is%20delighted%20to%20award,technology%20across%20the%20maritime%20industry) (accessed 14 February 2024).
- Melchers RE and Beck AT** (2018) *Structural Reliability Analysis and Prediction*, 3rd edn. Hoboken: John Wiley Sons.
- Murphy KP** (2022) *Probabilistic Machine Learning: Fundamental Topics*. Available at <https://github.com/probml/pml-book>.
- NewScientist** (2023) Two of the UK’s water companies are still using dowsing to find leaks. Available at <https://www.newscientist.com/article/2356376-two-of-the-uks-water-companies-are-still-using-dowsing-to-findleaks/> (accessed February 15, 2024).
- Prince SJD** (2022) *Understanding Deep Learning*. Cambridge: The MIT Press. Available at <http://udlbook.com>.
- Raiffa H** (2002) Decision analysis: A personal account of how it got started and evolved. *Operations Research* 50(1), 179–185. <https://doi.org/10.1287/opre.50.1.179.17797>.
- Stewart MG and Mueller J** (2018) *Are We Safe Enough?: Measuring and Assessing Aviation Security*. Amsterdam: Elsevier. <https://doi.org/10.1016/C2016-0-01215-9>.
- Sutton RS and Barto AG** (2020) *Reinforcement Learning An Introduction*, vol. 01. Cambridge: The MIT Press.

- Tetlow G and Shearer E** (2021) *HS2: Lessons for Future Infrastructure Projects*. Tech. Report. Available at <https://www.instituteforgovernment.org.uk/sites/default/files/publications/lessons-hs2.pdf>.
- The Alan Turing Institute** (n.d.) AI Standards Hub. Available at <https://aistandardshub.org> (accessed 16 February 2024).
- The Engineering Council** (2017) Statement of Ethical Principles. Tech. Report. Available at <https://www.engc.org.uk/ethics>.
- The Royal Academy of Engineering** (2022) Engineering Ethics: Maintaining Society's Trust in the Engineering Profession. Tech. Report. Available at https://www.rae-engineering-ethics-full-report_v7.pdf.
- The Turing Way Community** (2022) *The Turing Way: A Handbook for Reproducible, Ethical and Collaborative Research*. Tech. Report. <https://doi.org/10.5281/zenodo.3233853>.
- Von Neumann J and Morgenstern O** (1953) *Theory of Games and Economic Behaviour*, 3rd edn. Princeton: Princeton University Press.
- Walia R, Narayanan P, Carette J, Tobin-Hochstadt S and Shan CC** (2019) From high-level inference algorithms to efficient code. *Proceedings of the ACM on Programming Languages* 3, 1–30. <https://doi.org/10.1145/3341702>.
- World Economic Forum** (2023) *The Global Risks Report 18th Edition Insight Report*. Available at <https://www.weforum.org/publications/global-risks-report-2023/>.