

RESEARCH ARTICLE

Open Access



How to tell the difference between a model and a digital twin

Louise Wright* and Stuart Davidson

*Correspondence:
louise.wright@npl.co.uk
National Physical Laboratory,
Hampton Road, Teddington,
UK


Abstract

"When I use a word, it means whatever I want it to mean": Humpty Dumpty in Alice's Adventures Through The Looking Glass, Lewis Carroll. "Digital twin" is currently a term applied in a wide variety of ways. Some differences are variations from sector to sector, but definitions within a sector can also vary significantly. Within engineering, claims are made regarding the benefits of using digital twinning for design, optimisation, process control, virtual testing, predictive maintenance, and lifetime estimation. In many of its usages, the distinction between a model and a digital twin is not made clear. The danger of this variety and vagueness is that a poor or inconsistent definition and explanation of a digital twin may lead people to reject it as just hype, so that once the hype and the inevitable backlash are over the final level of interest and use (the "plateau of productivity") may fall well below the maximum potential of the technology. The basic components of a digital twin (essentially a model and some data) are generally comparatively mature and well-understood. Many of the aspects of using data in models are similarly well-understood, from long experience in model validation and verification and from development of boundary, initial and loading conditions from measured values. However, many interesting open questions exist, some connected with the volume and speed of data, some connected with reliability and uncertainty, and some to do with dynamic model updating. In this paper we highlight the essential differences between a model and a digital twin, outline some of the key benefits of using digital twins, and suggest directions for further research to fully exploit the potential of the approach.

Keywords: Digital twin, Data science, Uncertainty evaluation, Model updating

Introduction

Digital twin is a term that is being used for a wide range of things across a wide range of applications, from high value manufacturing and personalised medicines to oil refinery management and risk identification and mitigation for city planning. For some of the definitions, the reason why "twin" is used has been lost. The danger of this variety and vagueness is that a poor definition and explanation of a digital twin may lead people to reject it as just hype, so that once the hype and the inevitable backlash are over the final level of interest and use (the "plateau of productivity", see Fig. 1) may fall well below the maximum potential of the technology.



The Defense Acquisition University definition of digital twin, commonly used in defence, aerospace and related industries, (quoted in [1]) is:

“an integrated multiphysics, multiscale, probabilistic simulation of an as-built system, enabled by Digital Thread, that uses the best available models, sensor information, and input data to mirror and predict activities/performance over the life of its corresponding physical twin.”

and NAFEMS council member Rod Dreisbach recently defined it as:

“a physics-based dynamic computer representation of a physical object that exploits distributed information management and virtual-to-augmented reality technologies to monitor the object, and to share and update discrete data dynamically between the virtual and real products”

in the April 2018 issue of Benchmark Magazine [2].

From these definitions it is clear that there are three important parts in the digital twin of an object:

- a model of the object,
- an evolving set of data relating to the object, and
- a means of dynamically updating or adjusting the model in accordance with the data.

Not all practitioners use this definition. For instance in their e-book “Forging the digital twin in discrete manufacturing” [3], LMS research state:

“While we see many definitions of “Digital Twin,” LMS Research keeps it simple: A Digital Twin is an executable virtual model of a physical thing or system.”

This definition merely renames technology that has existed for many years, leading many engineers wondering why they are being told that a practice they have been successfully employing for decades is a new and vibrant thing.

The model used in a digital twin need not be a data-driven model, but it should produce results that are directly equivalent to a measured quantity (so that the model updating process is data-driven), and it is likely that the model will take in other measured quantities as boundary conditions, loads, or material properties. The key features of a model within a digital twin are discussed in more detail in “[What sort of model should a digital twin use?](#)” section of this paper.

The use of evolving data means that a key strength of the digital twin approach is that it provides an accurate description of objects that change over time. A validated model can provide a snapshot of the behaviour of an object at a specific moment, but using that model within a digital twin can extend the use of that model to timescales over which the object and its behaviour will change significantly.

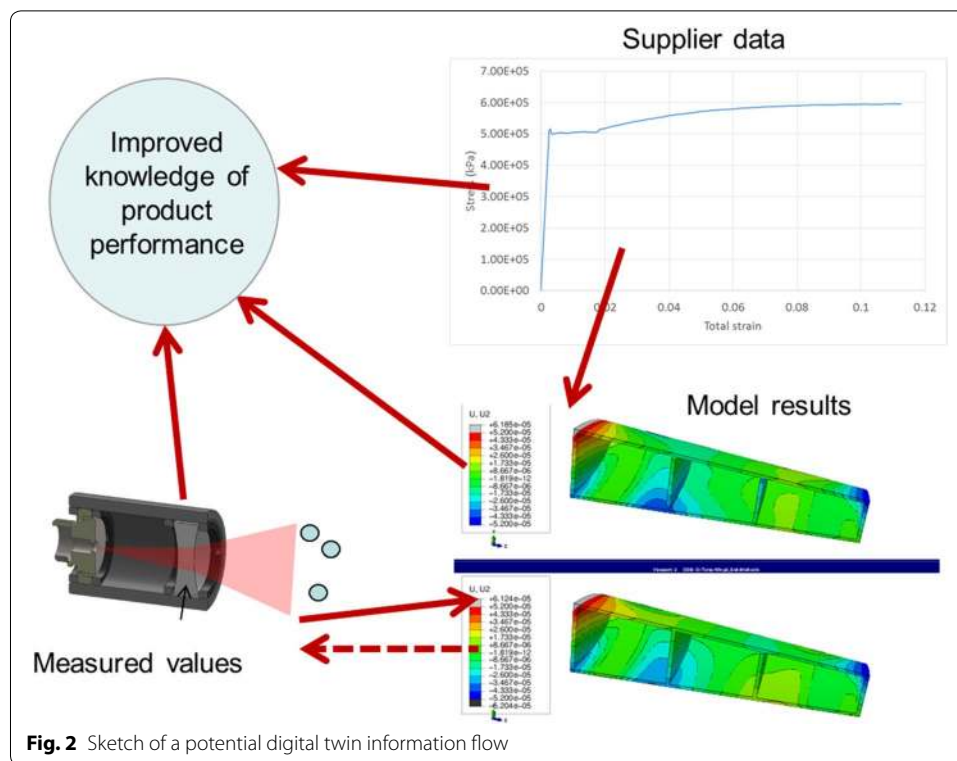
One of the key aspects of the parts listed above is that a digital twin has to be associated with an object that actually exists: *a digital twin without a physical twin is a model*. Biological identical twins are created at the same time, are the same when they are created, and (for the sake of this argument) they continue to be the same as they develop and age. This concept of similarity of two things throughout development and evolution is key to the true digital twin, and for this similarity to be possible, the physical twin must also exist.

The corollary of the requirement for the physical object to exist means that “digital twins for design” is only meaningful when the prototyping stage is reached. At this stage, the usual procedure is to test the prototype and update the design based on the results, engineering understanding, and previous experience. A digital twin approach would use the test data from the prototype to update parameters in the model of the prototype, use the updated model to predict performance in use, and then update the design. In many cases this adds extra steps (model updating and prediction) to the design process without bringing additional knowledge.

The exceptions are cases where the model and the prototype are insufficiently similar for model predictions to be reliable, so that updating the model of the prototype brings new knowledge. Prototypes are generally manufactured in small numbers and are therefore less likely to be subject to the variability associated with large-scale manufacture and so are more likely to be similar to the modelled versions. Hence it is likely that the application of digital twins to design will be fairly limited in impact.

For most high value engineering applications, the grand vision into which the digital twin fits is of a digitally enabled supply chain that can feed supplier data, in-house testing results, and on-line and off-line measurement results into a digital twin of products to obtain rapid performance predictions based on the latest data, as sketched in Fig. 2. For other applications the real-time aspect is less critical (or real-time means hours not seconds), but the ability to evolve the model as reality changes and to make predictions with a high level of confidence that the model is accurate is still of benefit.

The National Physical Laboratory is the UK’s National Measurement Institute, responsible for realising and disseminating the SI units. Our expertise in measurement, modelling, uncertainty evaluation and data analysis means that we have experience in the key areas that contribute to a digital twin. Our aim is to ensure that measurements are traceable and reliable, and that models are verified and validated and are updated using



trustworthy methods that are suitable given the characteristics of the data, so that we support the confidence in the intelligent and effective use of data in digital twins.

Do I need a digital twin?

Digital twins are of most use when an object is changing over time, thus making the initial model of the object invalid, and when measurement data that can be correlated with this change can be captured. These changes could be undesirable, for instance wear in bearings or fatigue in metal components, or they could be neutral but important, for instance variations in supplied material properties. If an object does not change much over time, or if data associated with that change cannot be captured, then a digital twin is not likely to be useful.

One reason that the digital twin concept is so valuable in manufacturing is that it allows for development of individual models of individual objects within a unified framework that makes model development, validation, and updating simple. An individually tailored model can be used for many applications during manufacture and service, including:

- enhanced in-line measurement processes, using the model to identify the data that will most improve physical understanding,
- smart assembly, ensuring that individual parts are chosen for optimal performance and scrap rates are reduced,
- assembly verification, so that complex multi-component structures whose internal details are not easily accessible can use measurements of the outer surface of

an assembly to provide confidence that the internal structure has been correctly deployed,

- performance assurance, to check that any measured deviations from the product specification do not compromise performance to an unacceptable degree,
- maintenance scheduling and smart maintenance, by using the twin to update parameters related to known possible faults and thus identifying problems before they become catastrophic,
- lifetime prediction, including the ability to revise component or system lifetime estimate in service.

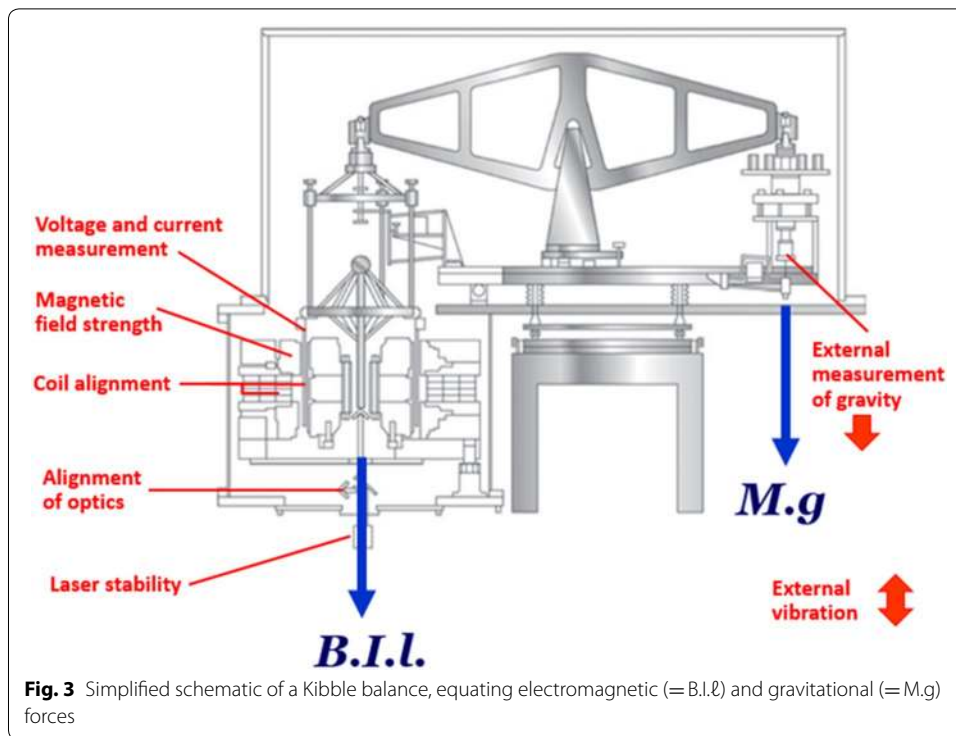
These applications are particularly relevant to complex multi-component assembled objects, and to products where some of the individual components or materials are sufficiently expensive that reducing scrap is an important driver. The approach is of less relevance at a product level where a degree of product variability in performance or appearance is acceptable, where well-established reliable quality control procedures are already in place, or where reworking and scrap are not significant costs.

Looking beyond manufacturing, the digital twin concept offers significant benefits in medicine. The ability to tailor drug characteristics, implant and prosthetic geometries, and treatment planning to match the needs of individual patients will lead to more efficient and effective treatment with fewer side effects, and associated improved health outcomes. Examples include tailoring drug dosage or composition based on the patient's response, monitoring of prosthetics to detect damage and wear, and adaptive radiotherapy where the treatment can be adapted to account for internal anatomical changes and for temporal changes in patient response.

Digital twins can also be of benefit to fundamental science. Scientific equipment can have characteristics that affect the experimental results. These characteristics may change over time and may be difficult to evaluate directly. A digital twin approach can ensure that the uncertainty associated with the results of the experiment can include justifiable contributions from these characteristics. Reliable interpretation of scientific results requires an understanding of the uncertainties associated with experimental results, and use of a digital twin can provide accurate estimates of, and identify methods of minimising, these uncertainties.

One example of this type of equipment is the Kibble balance. The Kibble balance is one of the experiments that will be used to realise the SI unit of mass, the kilogram, following the redefinition of some of the SI base units on 20th May 2019. Previously the kilogram was defined as the mass of a platinum–iridium cylinder kept in Paris, but after the 2019 redefinition it is now defined in terms of a fixed value of the Planck constant and can be realised with any suitable experiment [4]. The realisation typically occurs at a National Measurement Institute (NMI) level, and is cascaded to end users through a traceability chain. In simple terms, the Kibble balance equates electromagnetic and gravitational forces to determine mass in terms of very accurate quantum electrical standards [5].

A sketch of a Kibble balance with key components labelled is shown in Fig. 3. There are several features of the Kibble balance that will affect its performance. These include:



- stability of the voltmeter and resistor used to measure the voltage and current in the coil,
- changes of the field strength of the magnet with temperature and time,
- alignment of the electromagnetic coil within the field generated by the permanent magnet,
- alignment and stability of the optics used in the interferometer measuring the position of the coil,
- stability of the laser frequency used in the interferometer,
- local acceleration due to gravity,
- airborne and ground vibration.

By using all of the data gathered during the measurement process, these parameters and their variation with time can be better understood and will allow optimisation of the experiment and also lead to more accurate uncertainty estimates, thus benefitting worldwide mass measurements throughout the traceability chain. One of the major benefits of a digital twin of the Kibble balance will be dynamic evaluation of uncertainties, particularly with relation to their correlation. Because there are about 90 uncertainty components, correlation between them is largely ignored or at best estimated very approximately.

More widely, the digital twin idea can be applied at a higher level than individual products. Industries whose specific products are too high volume to merit individual twins can benefit from a twin at the factory level. This level of twin enables management of inventory, maintenance, shift patterns and scheduling to optimise production efficiency by minimising down time and waste. The models used in such applications are

generally empirical rather than physics-based. The use of real-time data to inform these models allows for improved determination of the parameters of these models, and for identifying when the parameters change which can be an early warning of the need for maintenance.

A further possibility for manufacturing industry is to use the digital twin to monitor the supply chain. If a model linking, for instance, mechanical properties of the raw materials used in manufacture to end product quality can be developed then the quality assessment data of the product can be used to estimate these properties. The estimates can then potentially be used to adjust process settings to improve quality and to provide feedback to the supplier of the material.

Similar approaches can be used to manage larger-scale assets such as energy networks, and even towns and cities [6, 7], in real-time. Energy networks are becoming more complex due to the combination of traditional and green generation technologies, including domestic generation, and energy storage facilities. Demand profiles are changing as transportation and heating methods rely more on electricity. Many countries have ageing infrastructure that was not designed for these operational changes. The ability to balance supply and demand and to avoid overloading of assets known to be fragile based on real-time data and network models will ensure that these challenges are met. For urban environments, the increased availability of cheap sensors for air quality and noise could enable reactive traffic management to reduce pollution, and the ability to produce data-driven reactive management for events such as floods based on infrastructure management data could lead to significant savings.

Another application of city-scale digital twins is the testing of autonomous systems, and particularly testing of the decision-making algorithms of autonomous vehicles (AVs). The number of potentially dangerous scenarios that an AV could encounter is extremely large. Many of the most dangerous scenarios are (fortunately) rare occurrences, and so are unlikely to be encountered during normal driving. Staging real-world tests for all of these events would be costly, complicated to implement reliably, and potentially dangerous. It is therefore far more cost efficient and safe to simulate these tests. A digital twin for this application would link:

- an accurate and updatable model of a real-world driving environment, including other vehicles and pedestrians and some description of weather and other atmospheric and environmental effects,
- models that simulate the response of the sensors deployed on the vehicle based on a careful choice of data from tests of these sensors,
- a model of the vehicle response to driving commands (e.g. steering changes, braking, etc.) that takes road surface conditions into account,
- and the AV control algorithms.

What sort of model should a digital twin use?

Digital twins can use any sort of model that is a sufficiently accurate representation of the physical object that is being twinned. In an ideal world, where computation would be instantaneous and accuracy would be perfect, digital twins would use models

derived directly from physics that took all phenomena likely to affect the quantities being measured and updated into account. For instance, a digital twin of a machine tool would be able to simulate the thermal and mechanical processes involved in milling of metal in real time and update knowledge about tool wear based on real-time measurements of part temperature and shape, so that plant maintenance could become more proactive and efficient.

The barrier of computational cost at high accuracy does not mean that physics-based approaches such as finite element modelling and computational fluid dynamics should be discarded altogether. Some applications of digital twins do not require high-speed computation, because the time frame over which the twin is to be updated is hours rather than seconds. One example is a digital twin of a wind turbine in an offshore wind farm. If the twin is being used to schedule preventative maintenance, it is likely that the time scale over which the decision is made is hours or days rather than seconds, making physics-based modelling feasible. Similarly a digital twin of a jet engine may be able to be updated using take-off data whilst the plane is in flight, making maintenance decisions within a suitable timeframe possible.

Some applications of a digital twin can use local models of key parts of a structure or an object rather than considering the complete system. These models can be defined to include the region directly affected by the parameter to be updated and little else of the surrounding structure, replacing parts of the computational domain with appropriate boundary conditions or lumped element approximations instead. One example is a digital twin of an aeroplane landing gear. The model does not need to include a complete aerodynamic model of the entire aircraft.

In other cases, where the problem cannot easily be reduced to a submodel, a high accuracy physics-based model can be used to generate a set of reliable results within the known operating parameter envelope of the physical object, and a surrogate model or metamodel can be constructed based on those results. A surrogate model is a simplified model, typically data-driven rather than physics-based, that runs more quickly than a full physics-based model and so can be used to generate updated parameter estimates and associated uncertainties more quickly. The surrogate model will be less accurate than the physics-based model, but if the level of accuracy is known, and ideally re-evaluated for cases where the operating parameters are approaching the edge of their envelope, then the loss in accuracy can be taken into account when making decisions based on the digital twin.

It would also be possible to construct a purely data-driven model to sit at the heart of a digital twin. This approach is often not advisable for several reasons. The most obvious is that a data driven model is only reliable within the region of input parameter space from which the data used to construct the model was taken. Using data-driven models for extrapolation without imposing any constraints based on physical knowledge is a dangerous approach.

In general, a model for a digital twin should be:

- sufficiently physics-based that updating parameters within the model based on measurement data is a meaningful thing to do,

- sufficiently accurate that the updated parameter values will be useful for the application of interest, and
- sufficiently quick to run that decisions about the application can be made within the required timescale.

The availability of a model that satisfies these three criteria strongly affects which applications can most benefit from a digital twin, as was noted in the previous section. The criteria also affect the ways in which physics-based models for digital twins differ from physics-based models for other purposes such as safety verification or performance modelling, where high accuracy may be more important than a short run time because the models are safety-critical but are run less frequently. In many cases it is likely that models used at the design stage can be reused (after some adaptation) for twins, because they are often developed at a subsystem level and used for rapid design iteration. The main adaptation required would usually be the inclusion in the model of the parameter that is varying over time, unless the relevant phenomenon has been simulated at the design stage.

What still needs doing?

The most likely applications of digital twins discussed above are all either high value or safety-critical. It is therefore very important to be able to trust the predictions of the digital twin. This requirement means that there must also be trust in the data, trust in the model, and trust in the updating procedure.

Trust in the model requires verification and validation procedures. Whilst this is a mature field and best practice is available, there are still open questions. Many engineering applications for which digital twins are valuable require a linked chain of models, for instance an engine may require models of heat flow, fluid flow, and stress. Solving such models in a fully coupled way is typically computationally expensive, not least because different analyses require mesh refinement in different locations. Solving them sequentially typically requires using results of one model as input for another model, potentially requiring interpolation and processing of the results and the introduction of further errors and approximations that the verification will need to consider.

A further complication is that the existence of uncertainty means that validation (comparison with reality) needs to be treated as a statistical process. All measurements require associated uncertainties to be meaningful. This requirement means that model inputs, and hence model outputs, and validation data all have associated uncertainties, and so comparison of data with model results should generate an estimate of the probability that the values are consistent rather than using a “5% error is good enough” approach.

Uncertainty evaluation also gives a better understanding of how much trust can be placed in the model results. This trust is particularly important for models that include parameters that cannot be determined independently. These models are precisely the cases when the digital twin concept is so useful: it allows you to estimate what you cannot measure directly and thus improve your model.

Many publications offering introductions to, and guidance on, uncertainty evaluation exist. Examples include

- Publications from NAFEMS, particularly from the Stochastics working group [8, 9],
- the Guide to the Expression of Uncertainty in Measurement [GUM] and its supplements [10, 11], which are available for free download. In particular section 6 of supplement 1 offers advice on assigning input distributions to quantities for various sorts of information,
- a best practice guide [12] that focusses on the challenges associated with uncertainty evaluation for computationally expensive problems.

Trust in data gathered by sensors can be partially addressed by associating metadata with the data. Metadata (“data about the data”) captures aspects of the measurement process that may affect the reliability and future usability of the data, for instance:

- sensor type and capabilities (precision, standard uncertainty, known sensitivities to environment, etc.),
- date of last sensor calibration and any other traceability information,
- operator (if relevant),
- time of data collection,
- sensor location, etc.

For the metadata to be of the greatest possible use, standards for metadata need to be common across industries, because most types of sensor are used in multiple industries. Metadata also plays a key role in the use of curated and historical data. The structuring of metadata (“ontology”) has a strong effect on the ease of carrying out data searches, particularly for data sets with multiple levels of metadata. The ability to carry out this type of search efficiently underpins the ability to merge data sets that are gathered at different points in time and space but relate to a single object, a task crucial to the effective use of digital twins. This searchability is also key for data reuse and data traceability in the event of product failure.

The Data Science group at NPL is working to develop metadata protocols that ensure that data has appropriate, and appropriately structured, information associated with it. The work is using the “Findability, Accessibility, Interoperability, and Reusability” (FAIR) principles of data curation as a starting point, and is identifying the unique requirements of measurement data that will need special consideration.

Two current EU-funded metrology-focussed projects [13, 14] are addressing the need to define and transmit calibration and uncertainty information as metadata in a manufacturing environment, providing enhanced trust in the data. The project “Metrology for the Factory of the Future” [13] is developing calibration methods for digital-only industrial sensors and will establish the infrastructure and software needed to take account of measurement uncertainty and quality together with measurement data. The complementary project “Communication and validation of smart data in IoT-networks” [14] is defining a digital format for the secure transmission and unambiguous interpretation of measurement-related data and is developing secure digital calibration certificates for IoT-connected measurement devices so that calibration becomes simpler.

The details of the data and how they relate to the model should also be considered. The data set for a digital twin needs to go beyond what is required for definition of geometry, material properties, boundary conditions and loads, and beyond what is needed for model

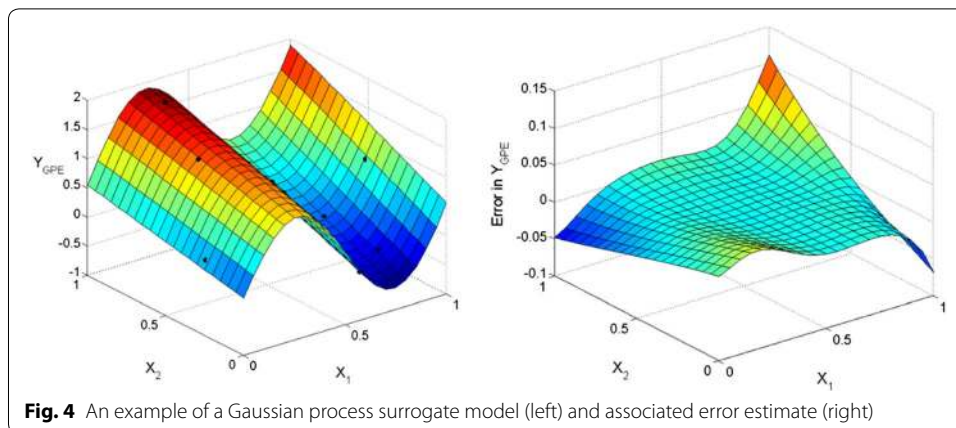
validation. The twin development process needs to identify a set of model parameters that are either poorly known or likely to change during manufacture or use, and the data needs to be sufficient to update these parameters. For instance a model of a power station may include efficiency curves for the various turbines, and these curves might be expected to change as the turbine ages. Using real-time power and angular velocity data to update estimates of the efficiency curves can support smart maintenance and reduce plant downtime.

The low cost of sensors and the easy access to cloud storage has led to widespread collection of extremely large datasets. These sets frequently consist of data from multiple sensors of varying types gathered at short time intervals. The challenge with using such data sets in digital twins is to identify which measurements at which locations or times have the most effect on the parameters to be updated within the twin. Data reduction techniques provide a way to address this challenge. The simplest method is singular value decomposition [15], which requires linearity of the model linking data and parameters, but efficient methods for data reduction are a lively area of current research and new techniques for handling nonlinear and transient models with various forms of data structures appear on a regular basis [16].

The best choice of process used to update the model depends on the size of the data being used to update the model, and the number of parameters that are available for updating. If only a small number of parameters are used than a simple optimisation approach may be appropriate. For larger numbers of parameters, data assimilation [17] may be more suitable. Data assimilation methods are widely used in meteorology where initial conditions of a system are only approximately known and better predictions can be obtained if the initial conditions are updated as data become available.

In many cases, the computational expense associated with these large multiphysics models of complex systems means that real-time updating, which may be required for effective process control, is not possible. Similarly, as was noted above, the computational expense can make a complete uncertainty evaluation too time-consuming. For these situations, replacing the computationally expensive model with an approximate model that is quicker to run can make these challenges easier to address. Such a model is known as a surrogate model or a metamodel.

One technique for surrogate model development that has been successful across a range of applications from fire safety to additive manufacturing [18, 19] is Gaussian process modelling [20], also known as kriging. This technique requires the user to have a set of results of the full model at a set of known values of the input quantities, called the training set. The technique constructs a model that interpolates the results at these points using a correlation function that effectively assumes that similar input quantity values will lead to similar model results, so that the closer together two points are in the input quantity space, the more strongly correlated they are. This approach is quite general, provided the model is broadly continuous, and a variety of correlation functions have been developed for different purposes. Another advantage is that when the surrogate model is evaluated at a set of input values where the full model result is unknown, it returns both an estimate of the model result for those values and an estimate of the error associated with that estimate. The error estimate means that it is easy to identify regions



of the input space where knowing the model result would add most benefit by reducing error the most, so the surrogate model can be developed iteratively.

A simple example is shown in Fig. 4. The left hand plot shows a training data set (black points) and the values predicted by the associated metamodel (coloured surface). The right hand plot shows the associated error estimate. The metamodel interpolates the training data and the error estimate is also zero at that point, but the error estimate is larger in magnitude further from the training data set, identifying regions where further training data could reduce error.

An alternative approach to improving computational efficiency for models based on partial differential equations is the application of model order reduction (MOR) methods [23]. The approach seeks to characterise the system being modelled in terms of a small number of functions or “modes”. The approach is analogous to the Fourier series decomposition of a time-dependent signal as a sum of sine and cosine terms. There is a slight complication as it is not always obvious what these modal functions should be, but there are methods to derive appropriate functions based on the results of finite element models. As with the kriging approach described above, the approach constructs a less computationally expensive model based on the results of a small number of runs of the full model. More advanced MOR techniques can include key parameters as an input to the reduced order model, making model updating simpler.

Conclusions

The concept of a digital twin pulls together several existing mature technologies. There is a danger that the over-use of the term will lead to under-use of the technology as potential users become cynical of the marketing buzz-words.

This paper has discussed what differentiates a “plain” model from a digital twin, highlighted some applications where digital twins can bring genuine benefit, and identified areas where research is still needed. Successful deployment of digital twins will require trust in the model, trust in the data, and trust in algorithms used to update the model based on the data. Once these elements are in place, we can have confidence in the decisions made using the technology.

Acknowledgements

Not applicable.

Authors' contributions

LW wrote the material related to data science and definition and applications of digital twins. SD wrote the material related to the Kibble balance application and provided additional suggestions for applications. Both authors read and approved the final manuscript.

Funding

The work reported here was Funded by the UK Government via the National Measurement System Programme, Data Science Theme.

Availability of data and materials

Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

Competing interests

The authors declare that they have no competing interests.

Received: 17 October 2019 Accepted: 12 February 2020

Published online: 11 March 2020

References

- West T, Blackburn M. Is digital thread/digital twin affordable? A systemic assessment of the cost of DoD's latest manhattan project. In: Complex adaptive systems, Chicago, USA, 2017.
- Digital Twin—looking behind the buzzwords, April 2018 edition of benchmark magazine. <https://www.nafems.org/publications/benchmark/archive/april-2018/>. Accessed 19 Feb 2020.
- [eBook] Forging the digital twin in discrete manufacturing: a vision for unity in the virtual and real worlds. <https://www.insresearch.com/research-library/research-articles/ebook-forging-the-digital-twin-in-discrete-manufacturing-a-vision-for-unity-in-the-virtual-and-real-worlds>. Accessed 19 Feb 2020.
- Richard P, Fang H, Davis R. Foundation for the redefinition of the kilogram. *Metrologia*. 2016;53(5):A6.
- Robinson IA, Schlamminger S. The watt or Kibble balance: a technique for implementing the new SI definition of the unit of mass. *Metrologia*. 2016;53(5):A46.
- Newcastle's 'digital twin' to help city plan for disasters. <https://www.theguardian.com/cities/2018/dec/30/newcastles-digital-twin-to-help-city-plan-for-disasters>. Accessed 19 Feb 2020.
- The Gemini Principles. Centre for Digital Built Britain. <https://www.cdbb.cam.ac.uk/Resources/ResoucePublications/TheGeminiPrinciples.pdf>. Accessed 19 Feb 2020.
- Stochastics Working Group. What is uncertainty quantification. NAFEMS publication. 2018. https://www.nafems.org/publications/resource_center/wt08/. Accessed 19 Feb 2020.
- Fortier M. Stochastics and its role in robust design. NAFEMS publication. https://www.nafems.org/publications/browse_buy/browse_by_topic/education/r0107/. Accessed 19 Feb 2020.
- BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP and OIML. Evaluation of measurement data—guide to the expression of uncertainty in measurement. Joint Committee for Guides in Metrology, JCGM. 2008;100.
- BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP and OIML. Evaluation of measurement data—supplement 1 to the guide to the expression of uncertainty in measurement—propagation of distributions using a Monte Carlo method. Joint Committee for Guides in Metrology, JCGM. 2008;101.
- Rasmussen K, et al. Best practice guide to uncertainty evaluation for computationally expensive models. 2015. <http://www.mathmet.org/publications/guides/index.php#expensive>. Accessed 19 Feb 2020.
- Metrology for the factory of the future. <https://www.ptb.de/empir2018/met4fof/home/>. Accessed 19 Feb 2020.
- Communication and validation of smart data in IoT-networks. <https://www.ptb.de/empir2018/smartcom/home/>. Accessed 19 Feb 2020.
- Golub GH, Kahan W. Calculating the singular values and pseudo-inverse of a matrix. *J Soc Ind Appl Math Ser B Numer Anal*. 1965;2(2):205–24. <https://doi.org/10.1137/0702016>.
- Chrétien S, Wei T. Sensing tensors with Gaussian filters. *IEEE Trans Inf Theory*. 2017;63(2):843–52.
- Rogers CD. Inverse methods for atmospheric sounding: theory and practice. Singapore: World Scientific Publishing Co.; 2000. ISBN 978-981-02-2740-1.
- Stroh R, et al. Assessing fire safety using complex numerical models with a Bayesian multi-fidelity approach. *Fire Saf J*. 2017;91:1016–25.
- Yang Z, et al. Investigating predictive metamodeling for additive manufacturing. In: ASME 2016 international design engineering technical conferences and computers and information in engineering conference.
- Rasmussen CE, Williams CKI. Gaussian processes for machine learning. Cambridge: MIT Press; 2006. <http://www.gaussianprocess.org/gpml/chapters/RW.pdf>. Accessed 19 Feb 2020.
- Wikipedia. Hype Cycle. https://commons.wikimedia.org/wiki/File:Gartner_Hype_Cycle.svg. Accessed 19 Feb 2020.
- Wikimedia Commons, GNU free documentation license, version 1. https://commons.wikimedia.org/wiki/Commons:GNU_Free_Documentation_License_version_1.2. Accessed 19 Feb 2020.
- Chinesta F, Cueto E, Abisset-Chavanne E, Duval JL, El Khaldi F. Virtual, digital and hybrid twins. *Arch Comput Methods Eng*. 2018. <https://doi.org/10.1007/s11831-018-9301-4>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.