

Digital Twin

ISSN: 2752-5783 (Online) Journal homepage: www.tandfonline.com/journals/tdtw20

Digital Twins through the lens of machine learning

Maximilian B. Toller, Franz M. Rohrhofer, David Cemernek, Sebastian Scher, Heimo Gursch, Lucas Iacono, Christian Laubichler, Mark Dokter, Dominik Kowald, Roman Kern & Bernhard C. Geiger

To cite this article: Maximilian B. Toller, Franz M. Rohrhofer, David Cemernek, Sebastian Scher, Heimo Gursch, Lucas Iacono, Christian Laubichler, Mark Dokter, Dominik Kowald, Roman Kern & Bernhard C. Geiger (10 Nov 2026): Digital Twins through the lens of machine learning, Digital Twin, DOI: 10.1080/27525783.2026.2756700

To link to this article: https://doi.org/10.1080/27525783.2026.2756700

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

Published online: 10 Nov 2026.

Submit your article to this journal

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=tdtw20

DIGITAL TWIN
https://doi.org/10.1080/27525783.2026.2756700

REVIEW ARTICLE

OPEN ACCESS Check for updates

Digital Twins through the lens of machine learning

Maximilian B. Toller[a,b], Franz M. Rohrhofer[a], David Cemernek[c], Sebastian Scher[d], Heimo Gursch[a], Lucas Iacono[e,f], Christian Laubichler[g], Mark Dokter[e], Dominik Kowald[h,i], Roman Kern[a,b] and Bernhard C. Geiger[a,j]

[a]Methods and Algorithms for AI, Know Center Research GmbH, Graz, Styria, Austria; [b]Institute of Machine Learning and Neural Computation, Graz University of Technology, Graz, Styria, Austria; [c]Student, Graz University of Technology, Graz, Styria, Austria; [d]Wegener Center for Climate and Global Change, University of Graz, Styria, Austria; [e]Data Management for AI, Know Center Research GmbH, Graz, Styria, Austria; [f]Institute of Human-Centred Computing, Graz University of Technology, Graz, Styria, Austria; [g]Area X: Simulation-based Development, LEC GmbH, Graz, Styria, Austria; [h]Fair AI, Know Center Research GmbH, Graz, Styria, Austria; [i]Department of Digital Humanities, University of Graz, Graz, Styria, Austria; [j]Signal Processing and Speech Communication Laboratory, Graz University of Technology, Graz, Styria, Austria

ABSTRACT

Research on digital twins is heterogeneous and covers numerous different real-world tasks in various domains. At the same time, machine learning has repeatedly demonstrated its potential as disruptive technology and has impacted many research fields. However, currently these two disciplines are somewhat disconnected, and key enabling machine learning technologies with high potential impact for digital twins remain niche topics. This review article presents digital twins from a machine learning perspective to encourage interdisciplinary research and to permit knowledge exchange between these two fields. We discuss the conceptual differences between digital twins in *Earth sciences*, *healthcare* and *industry*. Unlike previous reviews, we focus on machine learning-specific aspects such as data landscape, learning assumptions and objective functions for each domain. We systematically review 50 recent machine learning-based digital twins. We give a general, learning theory-based definition of digital twins to serve as starting point for theoretical research on machine learning for digital twins. Then, we describe all major aspects of the cyclical flow of information in digital twin systems on a domain-agnostic level, including data acquisition (measurement), pre-processing, data management, simulation, forecasting, control, intervention testing, and human-computer interaction.

ARTICLE HISTORY
Received 28 April 2026
Accepted 26 October 2026

KEYWORDS
Digital twin; machine learning; interdisciplinary research; survey; review

Introduction

In broad terms, a digital twin (DT) is an accurate virtual representation of a physical asset. DTs have become a relevant economic factor: In 2023 there was an estimated market size of USD 16.75 billion, which is projected to keep growing until 2030 [1]. Moreover, a large number of applied research articles have been published on the topic—in 2024 alone, there were more than 25 survey/review articles on DTs summarising thousands of original articles in various fields [2–27].

Currently, it is popular to use machine learning (ML) techniques for building or augmenting DTs and numerous works report a successful application. However, there is a notable disconnect between recent advances in ML research and DT applications: On the one hand, it was repeatedly demonstrated that novel ML-concepts can have a large impact on DTs—for example, Dang et al.'s [28] DT for structure health monitoring of a bridge based on neural networks. On the other hand, many ML researchers are only incidentally aware of relevant DT applications, and key enabling technologies such as physics-informed ML (PIML) [29] remain niche research topics.

CONTACT Maximilian Toller mtoller@know-center.at Methods and Algorithms for AI, Know Center Research GmbH, Sandgasse 34 8010, Graz, Styria, Austria.

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

An explanation for this disconnect is that DTs and ML require different expertise, and and developing expertise that effectively integrates both domains requires substantial effort and a broad interdisciplinary skill set:

- DTs require substantial expertise about the physical asset [30]. In many settings, the level of knowledge needed to build a DT is comparable to that needed to build/design the physical asset itself, including knowledge about the operational environment. Purely model-driven approaches are often too simplified, while purely data-driven approaches fail to detect causual relationships, since observational data alone only capture correlations [31,32].
- The precise usage of ML requires substantial expertise about statistical learning theory and experience with the practical behaviour of the involved algorithms. To control ML approaches, their hyperparameters need to be set manually [33], and for many ML methods, small changes in the learning setup lead to vastly different outcomes. According to the current state of learning theory, there is no reliable method for automatically calibrating an ML system to an application [34].

Another noteworthy segregating factor is the ML community's heavy focus on benchmark datasets, which is not realistic for DTs. It is conventional to evaluate a novel ML method using benchmark datasets, but the observed accuracy of results does not predict the method's suitability for practice. Applied research has shown that many novel, experimental methods fail to deliver in realistic use cases [35].

Since DTs can have high economic and societal impact, it is essential that ML research devises methods that perform well in various DT settings. However, this is challenging due to the aforementioned disconnect between the fields and the lack of interdisciplinary knowledge. The **aim of this review article** is to make DTs more accessible for ML researchers and to facilitate interdisciplinary research at the intersection of ML and DTs. To this end, we make the following contributions:

Review

- We describe DTs in three common domains through the lens of ML: *Earth science*, *healthcare*, and *industry*. For each domain, we highlight the subtleties related to data, assumptions, and learning objectives that affect the design of suitable ML systems and make recommendations for applied ML research in these domains.
- We review contemporary interdisciplinary research at the intersection of DTs and ML, with a focus on technology trends, research hotspots, opportunities and challenges.

Theory

- We formalise DTs across domains in a general way, including a rigorous definition that is suitable to serve as a starting point for theoretical research and that allows theoreticians to make more realistic initial assumptions. Our formalisation includes a definition of the involved physical, measurement, and control processes as learnable functions that are suitable for non-linear optimisation.
- We describe, on a domain-agnostic level, all phases of the cyclical flow of information associated with DTs from an ML viewpoint.

ML view of DT applications

There is a disconnect between the objectives of most contemporary ML research and the needs of DTs. To enable an increased practical impact of contemporary ML research, we describe three contrasting DT domains from an ML point of view. In each domain, we first present the predominant DT notion, data, learning assumptions and objectives. We then conclude the section with a summary and comparison of ML for DTs across the domains.

Earth sciences

DTs in Earth sciences are distinct in many ways—particularly in terms of what is considered a *high-fidelity* or *accurate* representation. A good example is Destination Earth [36–38], an international effort to construct a DT of the Earth itself using observations, satellite images, meteorological re-analysis data, big data mining, machine learning, physical models, and many other tools. A common goal of Earth Sciences DTs is to leverage high-resolution data to empower climate-related decision-making such as renewable energy planning, water management and disaster protection. One key idea is that DTs of the Earth will improve simulations and permit a detailed analysis of what-if scenarios since in-situ observations are usually scarce with limited coverage, and existing downscaling method accuracy relies on geographical features, grid resolution and reference datasets [39]. A visualisation of a DT in Earth Sciences is given in Figure 1.

From an ML point of view, there are several noteworthy aspects about DTs in Earth Sciences.

Data

In terms of volume, Earth sciences contain some of the largest spatio-termporal datasets in the DT context—for instance, Destination Earth [36] processes one petabyte per day. There is a multitude of heterogeneous observational data - including ground based weather stations, remote sensing data from satellites, precipitation radars, and radio sondes. All of these observations have different characteristics, particularly in terms of spatial and temporal sampling frequency, and their representation of the underlying physical variables - ranging from direct in-situ measurements of point-wise quantities, to indirect satellite based measurements, which with additional assumptions can be representative for mean values over larger areas. For many applications, however, it is desirable or necessary to know the state of the part of the Earth system of interest - e.g. the atmosphere - as a spatiotemporal grid of physical variables, and heterogeneous observations can not be used directly. In meteorological practice, this is solved by combining all available observations with numerical models to generate best guess estimates of spatiotemporal fields. This process is called data assimilation [40], and is both done operationally every couple of hours by weather forecasting centres, as well as retrospectively for several decades in reanalysis datasets [41]. From a DT perspective, this creates the unique position that there are two different potential starting points to develop an Earth science DT: 1) one can start from existing spatiotemporal gridded datasets - which are significantly easier to use than heterogeneous observations, and well suited for established deep learning architectures such as CNNs or transformer architectures. An example is provided in Figure 2. On the other hand, these datasets are also far from perfect, and machine learning might have the potential to improve on them. Therefore, 2), one can start directly from observational datasets - which are more challenging to use, but avoid the errors introduced by data assimilation systems. Additionally, both approaches could be combined - using both gridded datasets, as well as observational datasets as input to the DT.

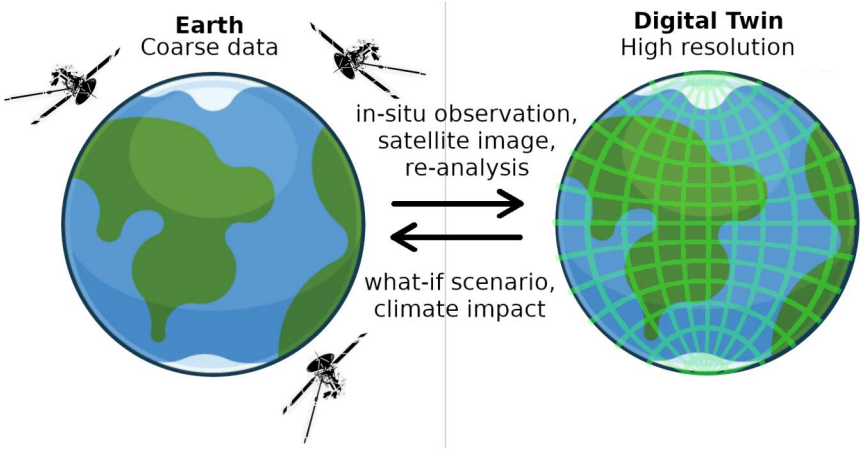


Figure 1. Common interpretation of DT in Earth sciences. The DT contains high-resolution data detailing a part of Earth, such as the atmosphere.

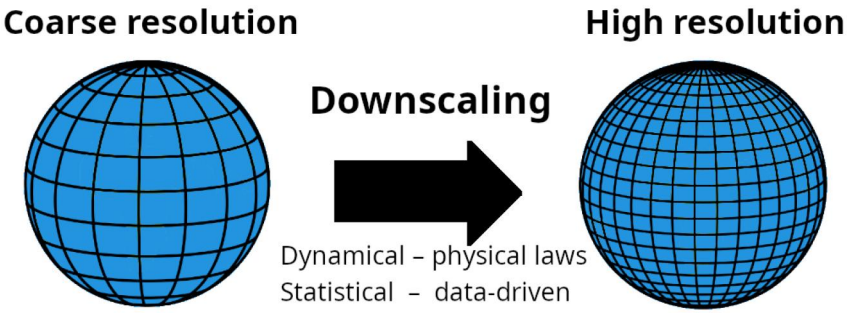


Figure 2. Simplified visualisation of downscaling, a typical DT task in the earth sciences. The objective is to generalise from a coarse resolution of some part of earth to a finer resolution, providing more accurate data. Dynamical downscaling uses physical laws encoded in numerical simulations, whereas statistical downscaling leverages spatio-temporal relationships via deep learning and similar methods.

Independent of which approach is used, there usually is sufficient data to overcome the sample complexity [42] of most learning algorithms. A downside of these large data quantities is that batch processing, distributed computation, and hardware acceleration are frequently necessary. This impedes exploratory data analysis and makes it more difficult to find sensible inductive biases, which are necessary for successful learning [43]. Moreover, despite the large amounts of available data, the relevant spatial region—often large parts of the Earth's surface—is almost always heavily under-sampled, with the grid cells of spatiotemporal datasets even in the best case being only representative for the mean over one square kilometre, and true observational data are only available for a small number of locations. One set of typical Earth science DT tasks often face the challenge to expand the available data to higher spatial and temporal resolution, or to new regions. In ML terms, this can be translated to a setting where a sparse tensor is given consisting of a grid/matrix for each timestamp, and the tensor needs to be populated, e.g. using interpolation, convolution or similar techniques. The tensor dimension increases accordingly if the Earth is modelled at multiple levels above the surface, for example to monitor upper air temperatures [44].

Another Earth science DT task is to model particular downstream problems, e.g the connection between weather events—which are likely directly modelled in the DT—and their impacts on a particular region, which might be available in additional datasets. From a ML point of view, this is a typical supervised learning problem. The Earth system is characterised by process with a wide range of spatial and temporal scales. In many Earth science DTs, the sampling frequency is usually between ten minutes and daily observations, yet there are exceptions such as seismology where high-frequency vibrations oscillate with split-second periodicity [45]. Observation data are typically heterogeneous and may require different data pre-processing for each station/satellite/observations system, yet many data sources have been harmonised in recent years [46]. Overall, the data landscape in the Earth Sciences is interesting for ML, especially due to the large amounts of available data that have similar stochastic properties, and the existing spatiotemporal datasets that ML can build on and potentially extend. This facilitates and encourages the use of deep learning methods, whose sample complexity is prohibitive in other domains [47].

Assumptions

From a statistical perspective, many Earth science variables such as temperature have 'well-behaved' properties which harmonise with typical ML assumptions. Many widely-used ML approaches assume smoothness of the underlying mapping to enable effective generalisation beyond training data—this is evident from the wide use of (smooth) kernel functions, as well as the differentiability condition tied to gradient descent [48]. Hence, one can expect that the learning-theoretical guarantees associated with this assumption, such as convergence to a local minimum, will usually hold for Earth Science DTs—at least in a partial manner. Similarly, the implicit assumption of light-tailed distributions with finite moments is commonly justified for modelling task such as air-quality prediction [49] and wind nowcasting [50], which further simplifies learning. However, the Earth Sciences also contain special tasks where these ML assumptions generally do not hold, which complicates the construction of DTs. For

instance, in hydrology, the distribution of events such as precipitation and flooding are skewed and heavy tailed, which requires statistically robust modelling approaches [51,52]. A further example is ice sheet monitoring, where the distribution of floe sizes follows a power law [53]. Moreover, the Earth Sciences frequently involve periodic, cyclical and seasonal processes that exhibit repeating behaviour such as tides, diurnal and seasonal cycles. Although these processes do not capture flawless repetitions, it is sensible and common to model them with cyclo-stationary ML models such as SARIMA (seasonal autoregressive integrated moving average) or with recurrent neural networks paired with feature engineering as shown in practice [54].

Learning objectives

Earth Sciences DTs have the interesting property that accuracy is not necessarily the primary learning objective, and other objectives such as physical plausibility and explainability are (at least) equally important [55,56]. This circumstance can be counterintuitive for ML researchers and engineers since the field mostly focuses on designing accurate predictors, classifiers and regressors. However, the ability to dissect the model and decode the logic behind the output is central in numerous Earth Sciences applications, and explains why traditional approaches such as Kriging [57] are still competitive with more complex ML approaches [58]. Additionally, particularly in meteorology, both estimates of the current state, and even more so forecasts of future states, are inherently uncertain [59]. Therefore, in addition to accuracy, reliable uncertainty estimates are equally important in many applications, which can be done with traditional ensemble approaches [60] and ML [61,62].

In general, Earth sciences commonly require that a DT operates on multiple different time scales simultaneously, each with a distinct expected behaviour. Accordingly, ML systems may also require different loss functions for different time scales or horizons.

Healthcare

In healthcare, the term DT is primarily associated with personalised medicine [63] for precise diagnosis and personalised treatment processes. This association entails the physical asset in healthcare DTs is often a human being, whose molecular status, physiology and lifestyle should be dynamically represented in the DT [64]. One of the key reasons for personalised treatment is the large variability in patient responses to drugs, where the desired effect is commonly only observed for a subset of patients [65]. It is projected that medical DTs may overcome this challenge and will enable the development of novel therapeutics via virtual clinical trials [66]. However, due to the complexity of the human body, many existing healthcare DTs focus on a specific organ or function that is relevant to the subfield, e.g. cardiac DTs aim to replicate the human heart [67], and in oncology simulations are run on DTs of cancer cells [68]. Besides virtually replicating the human body, secondary applications of DT technology in healthcare include DTs for hospital resource management and wearable technology which continuously transmits biomarkers such as blood pressure to the cloud [63].

From an ML perspective, there are several traits that affect the construction of DTs in healthcare:

Data

It is a characterising trait of the healthcare domain that available datasets from clinical studies are high-dimensional, with a small number of study participants but a large number of different measurements per patient, and each measurement is taken a small number of times per patient—often only once [69]. However, ML for medical DTs requires a different data landscape: To construct a personalised model, every measurement needs to be taken a large number of times for a single patient to give a reliable and dynamically updated estimate of the patient's condition over time. For most medical DT applications, personalised data is needed [70]—in other words, new data need to be dynamically collected for each new individual, as exemplified in Figure 3.

While this data requirement is feasible for some biomarkers that are easy to measure such as blood pressure or glucose, it can be prohibitive for invasive or less accessible biomarkers, e.g. Amyloid-beta measured for Alzheimer monitoring via positron emission topography [72]. This discrepancy manifests in a data landscape with large variation in data quantity—for some individuals, an accessible biomarker

Static data
Many patients
one measurement
per feature

Patient ID	Age	Anaemia	CPK	Diabetes	Ejection Fraction
1	75	0	582	0	20
2	55	0	7861	0	38
3	65	0	146	0	20
4	50	1	111	0	20
5	65	1	160	1	20
6	90	1	47	0	40

Dynamic data
One patient
many measurements
per feature

Patient ID	Age	Anaemia	CPK	Diabetes	Ejection Fraction
97	63	1	514	1	25
97	63	1	633	1	25
97	63	1	606	1	30
97	63	1	583	1	30
97	63	1	516	1	35
97	63	1	454	1	35

Figure 3. Comparison of classical data situation and DT data requirements in medicine. Top) The first six rows of the *Heart Failure* dataset [71], for each patient, there is a single measurement of multiple medical features. Such static data is common in medicine. Bottom) Data requirements for a healthcare DT. To enable dynamically adjusted, personalised diagnostics, the DT requires multiple measurements of multiple features for one—or several—patients.

is generously sampled via a wearable, whereas for a larger set of patients with diverse characteristics, only a few individual measurements are available. Hence, data fusion is a major challenge in the medical domain, and may impede the development of DT technologies. Moreover, this facet of the domain complicates the use of pre-trained ML approaches such as foundation models [73], which are trained over many patients, since insights cannot simply be transferred from one individual to another, and some biomarkers vary widely between individuals. In a wider sense, ML for healthcare DTs commonly needs to operate with either a) small *static* datasets with few instances (measurements) where learning is difficult for statistical reasons; or b) *dynamic* datasets with a large number of instances (measurements), but the population's variability is not well represented due a limited number of patients. In terms of data types, it is characteristic for the medical domain that categorical features and real-valued features are mixed—for instance, blood type (categorical) and blood pressure (numeric). However, image and text datasets are not uncommon, especially due to the wide use medical imaging technologies and since anamnesis is a world-wide standard procedure which generates a large amount of written protocols [74].

Assumptions

Besides the unbalanced data landscape, the healthcare domain is further characterised by challenging learning assumptions. The human body is a complicated 'system' that is subjected to many internal and external influences, which usually leads to an overspecified learning task, i.e. there are more constraints/equations than measurable features, which excludes a solution. Most prominently, there is an abundance of confounders for exceedingly many ML task such as heart disease prediction [75] or Alzheimer classification [76]. The construction of data-driven DTs is severely complicated by confounders since correlations between observed data and the target variable may be spurious, and it is recommended to assume there exists no accurate mapping between medical data and the learning target variable—even if ML training loss seems to suggest the opposite. For example, the enzyme creatine phosphokinase (CPK) is typically high in patients suffering from muscular disease, but an ML system cannot use CPK alone for predicting disease since CPK levels are also significantly affected by other factors such as regular intensive exercise [77]. It is thus essential to carefully validate any trained DT model, e.g. using held-out data from similar patients. However, despite the difficulties caused by confounders, successful ML is possible for constructing medicine DTs due to frequent (stochastic) redundancy in clinical datasets. Many important clinical states are reflected in multiple features—for instance, most types of lung cancer affect numerous biomarkers [78]—and the ML approach does not need to use all features for inference or prediction. Hence, the use of dimensionality reduction techniques and neural architectures with a sparse hidden layer is recommended and may possibly mitigate

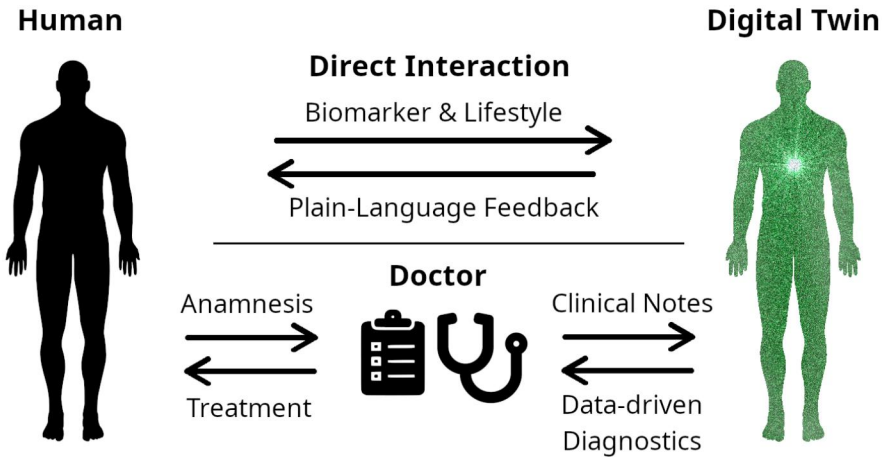


Figure 4. Characteristic schema of DT in healthcare. The DT receives biomarkers and lifestyle information from the patient, e.g. via a wearable, and clinical information from the treating professional. The DT provides the patient with high-level feedback (not medical advice), and the professional with data-driven diagnostics, treatment recommendations and similar.

the effect of confounders, given there is sufficient data for training—the higher the redundancy, the less likely it is that a confounder affects all predictive features.

Learning objectives

The main objective of most medicine DTs is directly related to an individual's health—see Figure 4. In most realistic applications, a DT may inform about an individual's health, but the actual decision on a potential treatment is made by a medical professional. Here, a personalised DT may be beneficial for discovering potential health risks, which a professional can then validate with a more accurate confirmatory test. While the DT should obviously be an accurate representation of an individual's health, the notion of accuracy is nuanced in this domain. Conventional measures commonly used in ML such as F_1-score or mean squared error may be suboptimal since the cost of errors in healthcare is not symmetrical. In classification tasks, due to the severe potential impact of health conditions, a false negative (patient has a condition but DT misses this) is usually more dangerous than a false positive (DT falsely reports a condition the patient does not have) [79], which implies that recall/sensitivity is more important than precision. Similarly, in regression settings an asymmetric loss function should be used which penalises underprediction errors of conditions more heavily than overprediction errors, for instance weighted asymmetric squared loss [80]. The asymmetric impact of errors is not wide-spread in ML, and it is crucial that the DT's learning objectives are selected appropriately to minimise health risks.

Industry

Industrial DT applications are diverse and have been extensively studied in previous research. As there is no single 'industry', there is much heterogeneity in industrial DTs, ranging from state-of-health estimation for battery management systems [81], over fault detection in manufacturing cells [82], to noise level monitoring for employee safety and health [83]. However, industrial DTs are usually not employed as an end in themselves, but are motivated by the desire to gain a competitive advantage such as increased reliability, detection of faults, and—most importantly—**reduction of cost** [84,85]. Especially in large industrial settings, it is common to observe systems of DTs interacting with each other rather than individual isolated DTs [86]. Such systems of DTs—or even systems of DT systems [87]—have high complexity and are a dominant characteristic of Industry 4.0 [88]. Here, interoperability and integration are important factors. The current lack of standards impedes the testing of different DT designs and makes the update of DT components difficult [89,90]. Currently, interfaces to legacy and third-party systems are required for the operation of DTs [11,91].

Despite their heterogeneity, there are several common traits shared across industries that affect the usage of ML of industrial DTs.

Data

It is evident that real data are needed to train and re-train ML models for industrial DTs, which is a surprisingly common limiting factor in practice—oftentimes, no suitable datasets are available, and data need to be manually collected and annotated. This can be a significant factor in cases where periodic re-training is needed [90,92]. Therefore, models that require less data are of particular interest [90,93,94]. Moreover, even if data are available, most machines, sensors, actuators, and logistics systems capable of collecting these data are not designed with ML in mind. Across most industries, data are not immediately available in the desirable numerical format that most ML approaches need, and multiple pre-processing steps are necessary—for instance missing value imputation and outlier removal [95], aggregation and data type conversions [96], as well as feature harmonisation and resolving naming inconsistencies [97]. Importantly, this implies that it is rarely a good idea to train the ML model directly on raw data, and commonly the pre-processing effort is larger than the effort needed to identify and train a suitable learner [98]. The three most common data types in the industries are a) time series, which are predominantly obtained from sensors and other embedded devices; b) tabular data, which is common in in operations management and logistics, e.g. via a DT of a facility or of a supply chain; c) text, which most frequently serves for documentation or is found in specification documents or quality management. Since most ML approaches cannot handle mixed data types, data fusion techniques such as encoding and embeddings [99] are frequently necessary.

Furthermore, relevant industrial data for DTs are often not collected, curated, and labelled due to high cost rather than technical impossibility. Therefore, one should carefully balance the ML component's expressive power with the cost of obtaining sufficient industrial data to meet the sample complexity requirements to avoid a situation where the DT requires more data than is cost-effective.

Assumptions

In many industrial applications the data-generating process is at least superficially understood and may also be backed by physical theories such as sensor profiles, physical measurement principles, or governing laws of the systems under consideration. This implies that an ML approach is advantageous if it can include this knowledge in the modelling phase—for instance via a physics loss term. From a data-driven perspective, it is noteworthy that typical industrial sensors and equipment pre-process the raw data they generate on a low level using the mean over a time window or similar statistics, e.g. a Kalman filter [100]. This step is often necessary to limit data quantity, e.g. if sensors measure in split-second intervals, but this implicitly imposes a strong assumption: Short-term trends and oscillations in the signal might be obfuscated. As such, industrial DTs are typically fed by data streams with already reduced information content, which thus makes some prediction/simulation tasks unsolvable. A noteworthy example for this is the (formerly) wide-spread computation of root mean squared intensity of vibration signals for monitoring turbines, which may hide or distort unexpected vibrations associated with a loss in gearbox efficiency [101]. There is no data-driven solution for overcoming this obstacle other than modifying the sensors or equipment. Hence, it is essential that the measurement resolution is considered when setting up the ML pipeline—any behaviour of the data-generating process that has characteristic time scales substantially shorter than the averaging window length cannot be learned accurately. Naturally, this limitation propagates in systems of DTs.

Beyond the limitations of sensors and measurement equipment, it is typical for industrial DTs to operate in dynamical settings where many common (implicit) ML assumptions are not satisfied—most notably that a) the training set has a similar distribution as the test set; and b) that the labels are correct. Distribution shifts and labelling noise induced by human error/disagreement need to be expected in settings where the DTs' environment can change and novel situations may arise. Reasonable approaches to overcome distribution shifts are robust learning techniques [102] and continual learning [103], whereas labelling noise can be addressed with robust losses [104] and confidence scores [105]. Figure 5 shows a visualisation of the challenging ML tasks associated with industrial DTs.

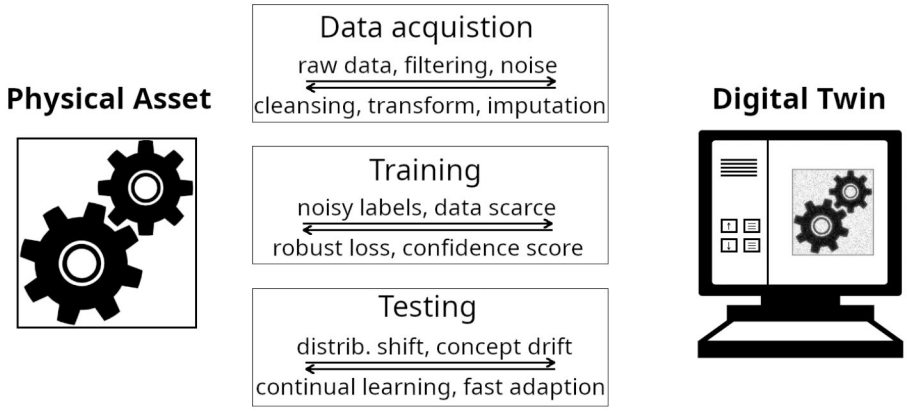


Figure 5. Challenges and solutions of ML for industrial DTs. In each ML phase, industrial DTs typically have specific obstacles that need to be overcome with appropriate ML-related mitigation.

Furthermore, real-time capabilities might be needed for the application. Many ML systems cannot process data arriving in real-time, whereas some DT systems would even require one to run entire simulations faster than real-time [106]. When real-time processing is needed, the DT should be designed according to principles of data stream processing and incremental learning [107], which focus on ML approaches that need to operate under time and memory constraints.

Learning objectives

Many novel ML systems behave poorly when deployed in industrial settings. A key reason for this deficit is underspecification [35], which implies that the loss function has a large number of identical (local) minima during training. Without further guidance, a trained ML system will converge to an arbitrary minimum due to random initialisation, and it cannot distinguish between the practical suitability of various minima in the loss landscape. To direct the ML system towards more industrial value, the objective function of an industrial DT should not only be data-driven, but also include further components as needed for the domain. For example, a vehicle DT should adhere to the physical laws relevant for driving, which can be modelled with a physics loss term in the learning objective. Similarly, an industrial need for stable DT performance under different heat conditions can be modelled via a loss term that tracks robustness against heat changes. Furthermore, it is essential that the training time and cost are at least considered when designing the learning objectives—operational necessity may pose hard limits on ML training resources.

Comparison

Although DTs in all three discussed domains have similarities, there are noteworthy differences from an ML perspective. The following superficial characterisation serves as guideline for the characteristic 'average' DT in the respective domain:

- Earth science DTs appear like an ideal 'playground' for ML due to the abundance of well-behaved data. Despite the complexity of the domain, many applications have minimal distribution drifts and the classical supervised ML setup is effective as long as spatio-temporal dependencies are considered. Uncertainty quantification is mandatory in many cases, limiting the utility of point-wise predictions or forecasts. Successful interdisciplinary research requires ML experts to comprehend the intricacies of various geo-data formats, the inherent uncertainty of non-observational ground truth, and what is physically plausible. Domain experts are encouraged to understand the statistical limitations of data-driven inter- and extrapolation, as well as the difficulty of ensuring consistent predictions. Compared to healthcare, Earth science DTs are built on much more data, and wider understood (geo)physical laws, which facilitates statistical inference but also permits extreme dynamical

complexities. Industrial DTs are typically also based on fewer data than Earth science DTs, but there are more similarities due to the underlying physical and chemical connections.

- Healthcare DTs are difficult to implement with ML due to confounders and limited data. ML should leverage the redundancy among features to mitigate spurious correlations with the target variable. The loss is asymmetrically biased towards positive detections to maximise screening efficiency and minimise chances of undetected health conditions. Interdisciplinary research requires knowledge of the inherent legal constraints, ethical dilemmas, and procedural difficulties associated with human subjects in experiments. ML researchers need to understand that most data used for DTs have confounders and available labels high uncertainty. Clinicians should comprehend how easily data quality can degrade, and that small studies do not produce sufficient data for ML. Compared to Earth sciences and industry DTs, medical DTs are rather different because the underlying system—the human body—is usually less understood, and there are no governing laws that could explain outcomes in an ideal setting.
- Industrial DTs are diverse and challenging. ML should not be trained on raw data due to frequent contamination. Distribution shifts and label noise are common. External constraints such as cost and physical requirements may outweigh detection accuracy, giving preference to approaches with weaker predictive power but higher resource efficiency and physical plausibility. To enable interdisciplinary research, ML experts must comprehend the operational limitations with respect to data collection, computing resources, response times and similar. Domain experts are recommended to tolerate certain enineering inaccuracies in the modelling phase, and to provide substantial assistance in data pre-processing. A core unique property of industrial DTs is that cost is a driving factor, which can be less restrictive in Earth sciences or healthcare DTs—although cost is always relevant, industrial DTs are often specifically designed to minimise cost.

A side-by-side comparison of the domains is depicted in Table 1, which also summarises the main points of our study of ML for DTs in these domains.

Interdisciplinary review

Since ML and DTs are both rapidly growing fields, there are constant changes and improvements to the used technology. In this section, we review recently published papers to capture contemporary developments at the intersection of DTs and ML. The insights obtained from this reveiw are later used as basis for the formal definitions we introduce later in the theoretical section.

Methodology

Database and query

Scopus and Google Scholar are selected as paper databases. We query for the terms *digital twin* and *machine learning*, which must occur anywhere within the publication. Note that, due to mixed terminology, alternative terms such as *digitisation*, *automation*, as well as *data mining*, *artificial intelligence* and *autonomous systems* are not uncommon in the literature. However, since the citations within most papers also use mixed terminology, our query is sufficient to identify the most relevant literature.

Selection criteria

We define several inclusion and exclusion criteria for papers. A publication must be recent, which we consider satisfied for all publications in the years 2025–2026. We only include primary research articles and exclude reviews, surveys, and position papers. Content-wise, a paper must describe any application of ML techniques in any DT-related task. Papers that predominantly focus on the development of novel ML techniques or ML-related innovation are excluded if they do not include a real DT application, study, or use case, or only speculate about future applied usage. Similarly, papers are also excluded if they focus on DTs and only mention ML techniques but do not apply any statistical procedure, learning algorithm, or other data-driven technique. We include publications in English and exclude publications in other languages. Finally, we exclude any publication whose full-text is not available in either

Table 1. High-level description of typical properties of DTs in the domains.

	Earth sciences	Healthcare	Industry
DT concept	High resolution replica of Earth or parts of Earth	Personalised replica of individual health and organs	Diverse; replica of product, machine or facility
Data	Large spatio-temporal datasets; batch-processing needed; clean data with low pre-processing effort but implicit missing values; numerical features	High-dimensional, undersampled tabular data; new personalised data needed;	Quantity and labels limited due to cost; mixed data types; high pre-processing effort.
Assumptions	Well-behaved processes; smooth, light-tailed distributions; cyclical behaviour	Many confounders which cause spurious correlations; feature redundancy facilitates modelling and encourages dimensionality reduction	Data generating process superficially understood; Pre-processing limits resolution; Distribution shifts and label noise are common
Learning objectives	(Geo-)physical plausibility central; Uncertainty quantification mandatory	Asymmetrical losses preferred to avoid health risks	Additional loss term as needed in domain; cost needs consideration
Models	Simple models such as Kriging popular; physics-informed ML recommended	Causal ML mitigates confounders; explainable 'white-box' models preferred by clinicians	Physics-informed ML or hybrid models recommended
Sources of diversity	Temporal diversity; different ecosystems and Earth system components; data originating from different organisations and collection methods	Equipment variety; biological diversity of patients; probabilistic nature of treatment outcomes	Equipment variety; operational constraints; sensor availability; safety requirements
Commonalities	*Healthcare*: Inherent uncertainty, temporal dynamics *Industry*: Data assimilation strongly related to online monitoring, physics information	*Earth sciences*: Inherent uncertainty, temporal dynamics *Industry*: Incomplete data, missing labels, data acquisition costly	*Earth sciences*: Data assimilation strongly related to online monitoring, physics information *Healthcare*: Incomplete data, missing labels, data acquisition costly
Interdiscipl. challenges	Large data volume; mixed ground truth (observation, assimilation, re-analysis)	Human subjects in experiments (legal, ethical, procedural); pure data-driven methods fail due to confounders	High specialisation; operational constraints limit feasible ML solutions; data not collected with ML in mind

open access or under a standard research institutional subscription for the publishers *Elsevier*, *Springer Nature*, *John Wiley & Sons*, *Taylor & Francis*, *Oxford/Cambridge University Press*, *IEEE*, and *ACM*.

Domain

For each article, we categorise its domain based on the described DT use case or concept. We include the categories *Earth Sciences* and *Healthcare* as defined in the previous section. Due to its diversity, *Industry* is split into the following more specific categories: *Manufacturing*, *Construction*, *Energy*, and *Aviation*. Corner cases, such as a paper describing battery technology for aircraft, are categorised based on the predominant categories of the journal (or conference) where they were published.

Deployment conditions

For each ML-driven DT use case, we observe the solution's designated speed and computing infrastructure conditions: We distinguish between a) (near) real-time systems, which permit a productive interaction with the DT under working conditions, and b) offline systems, where there is a notable delay between computation and an update of the DT. Similarly, we distinguish between ML-driven DT solutions that i) are suitable under restrictions on computation power such as *edge devices* or which permit an integration of *cloud computing*, or ii) are not suitable in such settings and only provide generic or unclear consideration of the available computing infrastructure.

Objective

We analyse the purpose of each ML-driven DT to observe common deployment objectives across domains, applications and technologies.

ML types

In each paper, we analyse the applied ML technologies, and denote the primary used ML type. The ML types found in the literature are diverse and range from *computer vision* over *reinforcement learning* to *physics-informed machine learning*. If a publication uses several ML types, we only report the dominant type. We track ML types across all domains, as well as for each domain separately.

ML maturity

We further determine the level of maturity of the observed ML technologies, based on the current state-of-the-art in ML. We distinguish the following five levels of maturity:

1. **Physical/chemical simulation** A domain knowledge-driven simulation or inference is conducted to inform the DT, e.g. using differential equations and the Finite Element Method (FEM).
2. **Statistical models** Data-driven methods based on classical inferential statistics such as *ridge regression*, *variational Bayes*, *ARIMA*, *Kalman filters* or similar.
3. **Established ML models** These include wide-spread algorithmic ML methods such as *random forest*, *XGBoost*, *support vector machines*, *k -NN*, as well as genetic algorithms. We also place *Multilayer perceptron* with a single hidden layer and *Gaussian process regression* in this level.
4. **Deep learning** Any NN-based method that either uses more that a single hidden layer, or uses advanced layer types such as convolution or recurrent. Examples are *Autoencoder*, *CNN*, *GNN*, *GAN*, *LSTM*, and *Transformer*.
5. **Autonomous AI system** ML methods in a system that autonomously sets goals and implements suitable actions. Prime examples are Agentic AI and similar systems whose control-flow is driven by large language models.

Note that we do not claim that a higher level of maturity implies technical superiority in any scenario. Rather, maturity levels are derived from the historical progression of ML research, which arguably began with descriptive statistics and now focuses on deep learning and agentic AI.

ML usage

For each publication, we determine whether existing ML techniques are applied 'off-the-shelf' or whether they are adapted to the DT's needs. The use of existing techniques with numerical adaptions of hyperparameters such as number of clusters, learning rates or similar is considered *applied usage*. Usage of custom loss functions, integration of domain-specific biases or similar is classified as *customised usage*.

ML failure

Not all applications of ML for DTs are successful. We explicitly include works that report negative results for ML-based DTs, and track the proportion of positive to negative results.

Technology transfer

Finally, we analyse whether the selected ML methods are applied in the setting they were designed for. Some ML methods are designed for specific data types, distributions or similar, and their behaviour beyond these settings is not well studied. For each identified publication, we observe whether or not ML techniques are used in an 'unintended' way, which may indicate creative solutions for positive results or may explain failure for negative results.

Results

The initial query yielded 8,572 documents on Scopus and about 591,000 results in Google Scholar. We selected all documents which appeared in both databases. Out of these documents exactly 50 papers were found to align with the selection criteria and selected for further study [108–157].

The distribution of the observed DT domains are depicted in Figure 6. Industrial DTs are most commonly linked with ML technology, with a total of 70% of the screened papers focusing on either

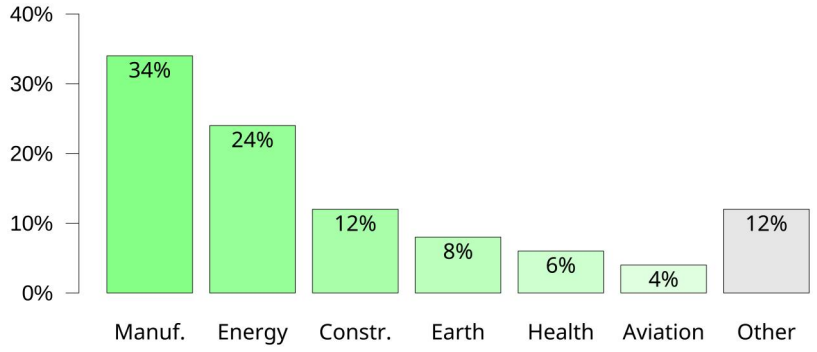


Figure 6. Distribution of ML-driven DTs by domain. ML-driven DTs were most commonly reported in manufacturing, followed by energy-related industries and construction. 12% of the reviewed papers did not fit into any category. For example, individual publications addressed ML-driven DTs in sports or agriculture.

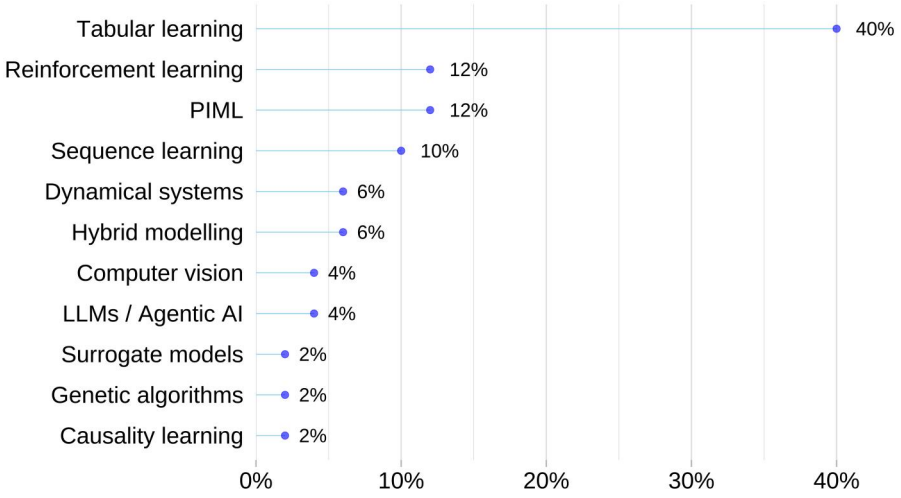


Figure 7. Distribution of ML techniques observed for DT applications. The most common global technique is tabular learning (40%), showing that tabular data is common in many DT applications.

manufacturing (34%), *energy* (24%), or *construction* (12%) industries. Earth sciences (8%) and healthcare (6%) DTs reach a combined total of 14%. Within our sample, 4% of ML-based DTs were observed in aviation, which is considerably lower than the partition reported in Tao et al.'s survey [158].

In terms of ML technologies, we observe that tabular learning techniques such as decision trees or gradient boosting [159] are most commonly used across domains (40%). In second place are reinforcement learning and physics-informed machine learning (PIML) with 12% each. Two papers reported incorporating large language models (LLMs) or agentic AI into DTs. Similarly, genetic algorithms, causal learning, surrogate models and computer vision were all observed in two or fewer DT applications, making these techniques rare. An overview of these numbers is depicted in Figure 7.

To assess which ML techniques are used in which DT domain, we depict a heatmap in Figure 8. Dissecting ML techniques by domain, we observe that tabular learning is not only the most commonly used ML-technique in total, but also within the majority of DT domains. Reinforcement learning is used predominantly in industrial DTs, whereas PIML is used in DT applications where adherence to physical laws is of particular interest such as Earth sciences, manufacturing, construction, and aviation. Sequence learning was most common in energy-related DT applications, which is unsurprising since the energy industry produces mostly time series data.

Concerning the DT's main objective/purpose, we observed four primary directions which can be associated with domain needs and ML tasks. The most common DT objective was related to either prediction or forecasting (44%), in the sense that future states of the DT or predictive inference tasks were considered, e.g. predict scrap rates from given production data. Optimisation and control (22%) were

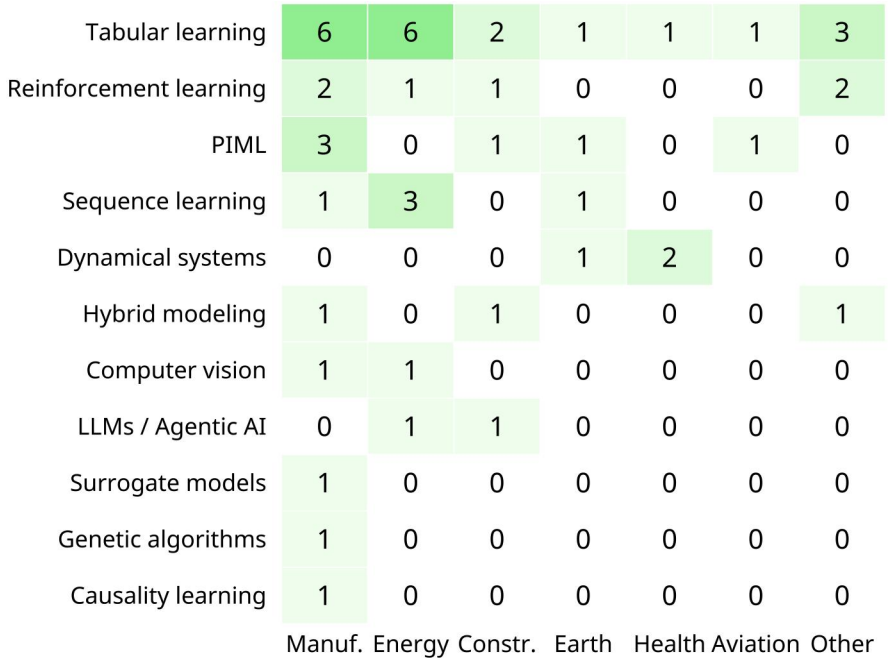


Figure 8. Frequency of ML-techniques per domain. Tabular learning is most common in total, and in most domains it is also the most commonly used ML type. One possible explanation is the wide use and availability of tabular data.

also common objectives, and used the ML-based DT to either ensure that that the physical asset remains in a desired state, or to improve its state based on the DTs state. The third common DT objective was anomaly detection and monitoring (22%), where the DT's purpose is to identify deviations, faults, or other undesirable developments in the physical asset. Four papers reported scenario studies and what-if analyses as the DT's main objective, whereas two papers reported unique objectives that fit in no category. An overview of the DT objectives is depicted in Figure 9.

A careful survey of the ML-based DTs permits an assessment of their technical maturity, from an ML perspective (see Figure 10). We observed two DTs using LLMs or similar techniques to integrate an autonomous (or agentic) AI system (4%). The most common ML maturity level was deep learning, with exactly half of the reviewed publications using NNs to inform, control, or improve the DT (50%). Established ML models such as random forest were used in 17 DT applications (34%), whereas statistical models where reported in three papers (6%). Three further DT papers considered ML techniques, e.g. for comparison, but the proposed DT solution was based on a physical model or simulation.

Figure 11 shows the observed distributions of ML usage, reported ML failures, and technology transfer. The majority of examined publications (62%) applies existing ML techniques either off-the-shelf or with minimal modifications (e.g. hyperparameter tuning). The remainder (38%) either adapted existing ML techniques in non-trivial ways, or designed a unique ML method which was integrated in the DT. Almost all publications reported a successful integration of ML into the DT (90%). Only 10% reported that the examined ML approaches were insufficient in some way. It is noteworthy that all of these paper instead proposed a DT solution based on a physical or statistical model. Concerning technology transfer, 24% of the examined DT publications applied ML in a non-designated way. The most common transfer was to use static, memory-less models such as random forest for sequence learning. However, the majority of papers (76%) applied ML in the designated way.

Discussion

After presenting the results of the review in a neutral way, we now discuss trends, potential opportunities and challenges we perceive.

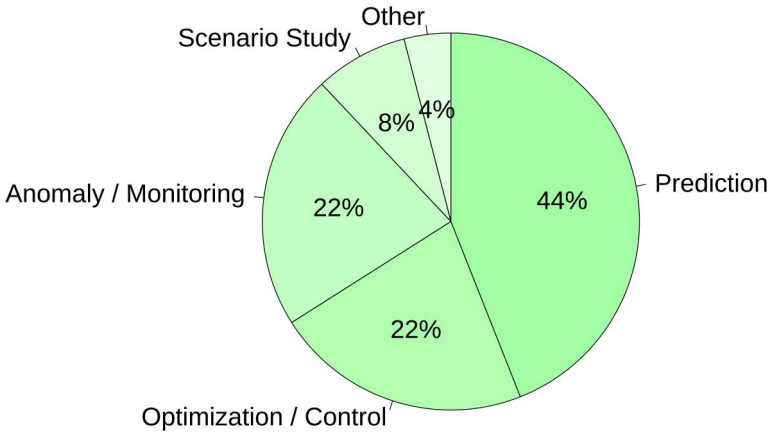


Figure 9. Distribution of observed ML-based DT objectives. Most DTs are used for predicting a future state of the physical asset, or for predicting a quantity of interest conditional on the DTs data.

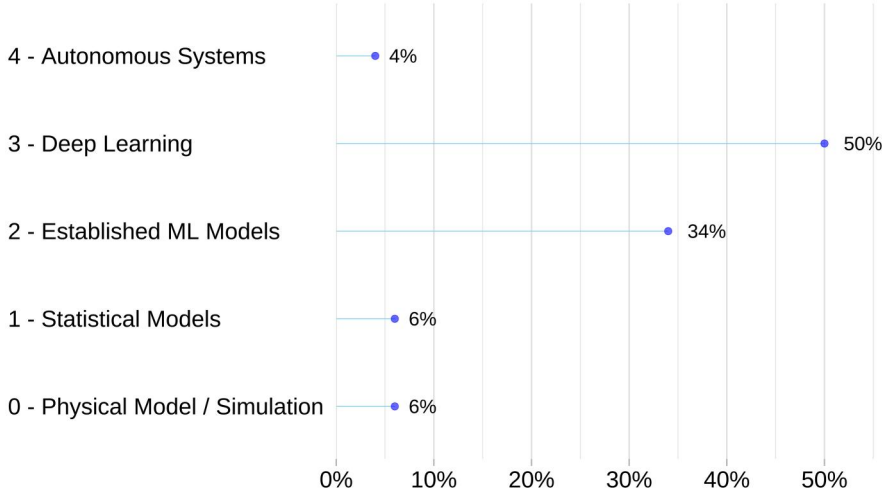


Figure 10. Observed ML maturity levels in ML-based DTs. Deep learning is wide-spread in DT applications, indicating that several ML advancements made in recent years have spread to the DT literature. First works have started experimenting with autonomous AI systems and AI agents in the context of DTs.

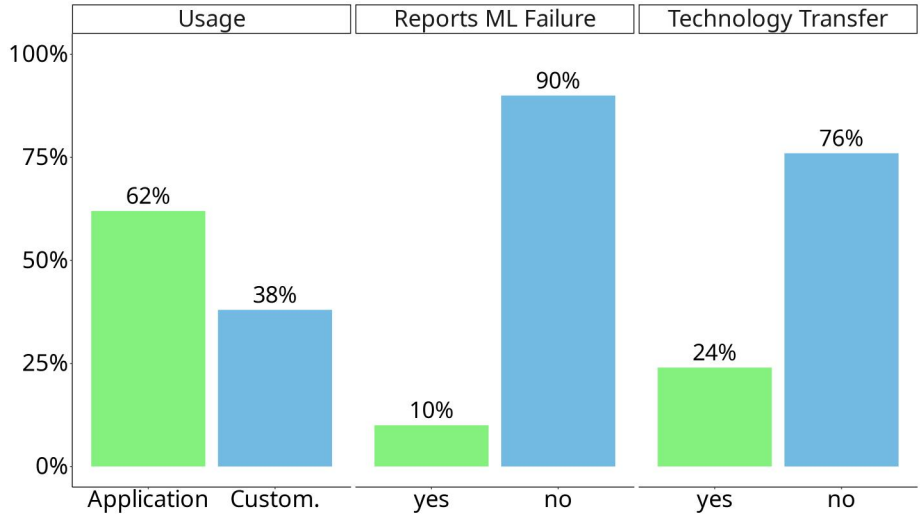


Figure 11. Distribution of ML usage types, reported ML failure, and technology transfer. An application of existing ML techniques was more common than a customisation for the DT. Most works reported a successful integration of ML into the DT. The majority used ML techniques in the designated way.

Trends

Deep learning for tabular data was a major focus in ML research a few years ago [160]. This past ML trend seems to have proliferated to contemporary DT research, given the large number of matching publications. It seems reasonable to speculate that applied fields such as DT research naturally experience a delay to the most recent developments in basic ML research—the transition from theory to practice requires additional effort, and it may be safer to prefer established techniques over experimental novelties. It is further noteworthy that we observed no tabular NNs applied in tabular regression/classification tasks, indicating that these novel technologies have not yet arrived in practice.

The distribution of DT domains shows that still most DT applications are industrial. However, we believe there are rising trends towards DTs in Earth Sciences and Health care, with programs such as Destination Earth [36] and the European Virtual Human Initiative [161] providing substantial funding to reinforce these trends. Advancements of DTs in these areas may proliferate to dependent fields such as agriculture, where contemporary works [145] indicate multiple opportunities for data-driven solutions and ML. Similarly, as human life expectancy is globally rising [162], increasing health-related data quantities and heightened interest in ML-driven healthcare DTs can be predicted with reasonable certainty.

Technology-wise, it is noteworthy that first DTs are adopting PIML. Given that PIML is a natural, data-driven yet physically consistent alternative to classical approaches for dynamical simulation or parameter inference, we expect an increasing number of successful DTs based on physics-informed ML. Despite the current challenges for PIML research [163], these techniques address a wide spectrum of engineering challenges, and are promising candidate for dominating industrial and Earth science DTs in the near future.

Opportunities

Generative ML for DTs. It is noteworthy that almost none of the examined DTs were strongly linked with established generative ML methods. Only two works reported using transformer-based LLMs [164] as generators. In broad terms, NN architectures such as autoencoder, generative adversarial network (GAN), or diffusion models, are ML models trained to generate new data such as images, text, or numerical data based on a prompt, which may be numeric or verbal. In a DT setting, there are numerous promising applications for generative ML:

- *Robustness testing*—Adversarial generative ML is designed to produce inputs that bring a system in an undesirable state. In a DT setting, this enables an indirect study of possible failure modes of both the DT and the physical asset, which can, e.g. reduce the costs of testing. testing [165,166].
- *Rapid documentation*—DTs typically are complex and require accurate documentation for reproducibility and maintenance. Generative ML can rapidly create specification documents and application programming interfaces (APIs), which reduces documentation effort [167].
- *Visualisation*—Productive interaction between humans and DTs often requires effective visualisations [168]. Stable diffusion and similar methods can rapidly produce visualisations of DT components or what-if scenarios, which facilitates interaction with and communication about DTs.

Besides these main applications, generative ML may also take over specific supporting roles such as rapid translation of the DT interface within multi-national corporations, or to identify bugs or security weaknesses in existing DT software.

Dedicated sequence ML. Although sequence learning was among the most commonly used ML techniques in the examined literature, it is currently still outnumbered by tabular learning for prediction tasks even when the underlying DT data streams are sequential, e.g. which is typical in Earth sciences and energy-related industries. A possible explanation is the strong performance of tabular methods such as XGBoost in recent—mostly finance-oriented—forecasting competitions, where sequence learning techniques such as LSTMs or transformers underperformed [169]. However, in other domains sequence ML has shown encouraging performance [170], which motivates the idea that DTs may likely benefit from

employing such techniques—especially in Earth sciences and energy-industry. DTs naturally generate continuous time series, where sequence models can explicitly learn the temporal dependencies which are difficult to capture with tabular models. Moreover, as DT research and operation continues, more sequential data is collected, which mitigates the high data quantity requirements of many sequence ML methods.

Autonomous AI systems. Currently, there is a surge of rapid development in autonomous AI systems, where an AI agent triggers actions on its own, across multiple research fields. It is natural to expect that this trend will proliferate to DT research, with two works in our survey already reporting a partial integration of autonomous AI agents in the DT solution [121,149]. Autonomous AI systems open several opportunities for DTs, out of which we consider the following two major opportunities:

- *Decision Instead of Prediction*—Currently, prediction is the most common DT objective. However, an autonomous AI empowered DT could take an additional step and decide upon necessary actions based on the prediction [171], and even act on it (subject to human oversight).
- *Interconnection*—Many autonomous AI systems are capable of launching multiple agents simultaneously. A multi-agent AI system, where each agent controls one DT, has the potential to overcome the current difficulties in interconnecting multiple DTs to a system of DTs [172].

Other opportunities for improving DTs with autonomous AI systems include explanations via chatbots, integration of foundational models for multi-objective DTs, self-learning DTs that continuously adapt to changing conditions, and trustworthy [173] DTs by enabling autonomous risk assessment and certification.

Open challenges

Despite the large proportion of a successful ML-based DT, we observe four challenges that current interdisciplinary research faces.

Challenge 1: single-use solutions. The effort needed to create an ML-based DT is high and specific. The reviewed DT solutions do not seem to generalise well, and the idea that promising ML techniques can be re-used for several similar DTs is an illusion. Although contemporary DT literature brings forth many interesting successful applications, we observe a lack of reflection on alternative ML setups, comparison with similar DTs, and evaluative reproduction of existing DTs. To overcome this challenge, the field is encouraged to more strongly focus on testing and evaluating existing ML-based DTs instead of building DTs from scratch when existing work looks promising. Moreover, when proposing a novel ML-based DT, it is encouraged to examine its ability to generalise—for example by testing on datasets from different domains, adapting DT tasks, as well as conducting parameter sensitivity analyses and ablation studies.

Challenge 2: ML irreproducibility. The majority of the DT literature we reviewed omits certain details about the selected ML approaches. Many publications mostly focused on the application and its background, while the ML approaches were only briefly described. Critical details such as learning rate, weight initialisation strategy, NN layer sizes, selected optimiser, and further hyperparameters are commonly not mentioned, which makes it difficult to reproduce the ML part of the DT. Pure ML literature typically provides such details in an appendix, but this information is mostly absent in the DT literature. A possible solution to overcome this challenge is to publish the necessary details, and make the used software components such as source code accessible.

Challenge 3: positivity bias. The low proportion of reported ML failures indicates the field has a preference towards publishing successful ML-based DTs. Such a skew is not unusual and can be found in multiple fields including basic ML research. However, due to the high degree of specialisation in DTs, a lack of negative results aggravates the low rate of generalisable or re-usable DT solutions. For

example, it would be highly interesting to know if a DT of a machine is restricted to a specific production series, or if a healthcare DT does not work for certain ethnicities. To overcome this challenge, we encourage both researchers and journals to actively publish negative or inconclusive results and document specific boundary conditions and limitations of their ML-based DTs.

Challenge 4: lack of specialization. Our review indicates that 62% of ML-based DTs apply existing ML methods with only minimal adaption. Due to the highly specific nature of most DTs, we believe there is currently much unused potential for improvement. A highly customised ML-based DT may provide better characterisations of the physical asset and may extract more useful patterns from the available data, e.g. with a specialised physically plausible ML objective describing the DTs control error. A push for synergetic DT solutions may yield higher accuracy, but also makes generalisation more difficult. As such, this challenge is particularly difficult to overcome, yet pioneer solutions exist, e.g. [139], which in a next step could be studied for generalisation.

Theoretical definition of DTs

There have been substantial efforts to define DTs and DT systems in various domains. Here, we give a theoretical, ML-centric definition of DTs to serve as foundation for general theoretical ML research on this topic.

Terminology and notation

For the most basic terms, we rely on the international standard [174], which provides the following definitions: an *entity* is a 'thing (physical or non-physical) having a distinct existence'; a *physical entity* is an 'entity in the physical world that can be the subject of sensing and/or actuating'; a *digital entity* is a 'computational element comprising data elements and procedural elements'. In the context of DTs, there are two relevant entities: A physical entity which we refer to as *physical twin* (PT), and a digital entity which we call *digital twin* (DT). The two different twins—physical and digital—are usually part of a so-called DT system which may have additional components, depending on the setting. The standard [174] summarises such a *DT system* as a 'system providing functionalities for the digital twin composed of inter-operating target entities, digital entities, data connections, and models, data and interfaces involved in the data connection process'. In a DT system, the exchange of information is bidirectional. Since the DT operates with data, we refer to the flow of information from PT to DT and the related technologies as *data layer*. The opposite flow of information from DT to PT is called *information backflow*.

Without loss of generality, we restrict ourselves to vectors (denoted by lower case letters such as a) and use matrices (denoted by upper case letters, e.g. A) only if notationally convenient. If any of these symbols is succeeded by round brackets, this denotes a function of the respective dimension, e.g. $A(\cdot)$ denotes a matrix-valued function. Calligraphic letters, e.g. $\mathcal{A}$, denote sets. We index the elements of a vector with integers, e.g. $a = [a_1, ..., a_m]$. The symbol d, as in da represents the total derivative. We use β as a placeholder symbol for additional parameters of a function that are not further specified, e.g. $a(t, \beta)$ denotes a vector-valued function of t (time) and possible other parameters, depending on the context. The symbol t denotes a specific point in time; in the context of digital information processing, we indicate the discrete nature of time by an integer subscript k, as in t_k. Functions of time are not always explicitly denoted as such when it is clear from the context, e.g. we may abbreviate $a(t)$ with a to avoid excessive notation.

Mathematical foundation: states, sensors, and control signals

We construct the mathematical formulation of DTs and DT systems based on several aspects derived from a physical point of view. A widely valid and practically relevant fundamental assumption is that the PT is in a state denoted by $z \in \mathbb{R}^d$, which evolves over time and can therefore be modelled as a continuous-time dynamical system[1], i.e. $z(t)$ with $t \in \mathbb{R}$.

The rate of change of this state, expressed mathematically as dz/dt or just $\dot{z}$, is influenced by various internal and external driving forces. On the one hand, these forces arise from fundamental physical principles, such as the laws of motion or conservation laws, and govern the intrinsic physical behaviour of the PT. Mathematically, these principles are expressed in the form of differential and algebraic equations. On the other hand, the state may also be influenced by external forces $u(t)$, which represent control signals applied by an external observer and which may appear as source terms in the differential equations. The interaction between the state and its driving forces—both intrinsic and extrinsic—can be mathematically described by a differential-algebraic equation (DAE) system of the form:

$$f(z(t), \dot{z}(t), u(t), \eta(t); \beta) = 0, \tag{1}$$

where f is a vector-valued function that encodes the relationship between the system's state and its driving forces.

Since it is generally not feasible—or even desirable—to capture all environmental noise effects within the state z, we explicitly include these noise effects as an additional driving force $\eta(t)$ in equation (1) to account for such (potentially time-varying) uncertainties and stochastic influences in the process. Additionally, β denotes physical system parameters—quantities that remain constant with respect to the system's state but characterise inherent properties of the PT, such as material constants, diffusion coefficients, and similar physical attributes. We further note that the DAE in (1) will usually be non-linear and not sparse, impeding the application of many analysis techniques which primarily focus on linear or sparse dynamical systems [48], and calling for the development of new mathematical theory [175].

We proceed by arguing that the flow of information from the PT to the DT, i.e. the data generation process, can be modelled as a raw signal—broadly speaking, obtained through observations, such as those provided by sensors:

$$\tilde{x}(t_k) = g^{(1)}(z(t_k), u(t_k), \eta(t_k)) \tag{2}$$

at discrete time points t_k, with $k \in \mathbb{N}$, reflecting the fact that in realistic scenarios the raw signal is available only at specific, discrete measurement times $t_k \in \mathcal{T} = \{t_1, ..., t_n\}$ due to finite sampling rates. Here, $g^{(1)}$ is a vector-valued function with m elements that maps the PT's state—again incorporating environmental noise η—to the raw observed signal $\tilde{x}$. Note that η is present in (1) and (2), highlighting that the raw signal $\tilde{x}$ is doubly influenced by the environment—once via the effect on the PT, and once via measurement noise. For all times t not covered by sampling, i.e. for $t \notin \mathcal{T}$, we define $g^{(1)}(z(t), u(t), \eta(t)) = \aleph$, where $\aleph$ denotes missing data. In practice, it may also occur that only a subset of signal channels $\tilde{x}_i$ is observed at time t_k, while other channels $\tilde{x}_j$ are missing—e.g. due to sensors attached to the PT operating at different sampling frequencies. In this case, using the notation introduced so far, we have $\tilde{x}_j = g_j^{(1)}(z(t_k), u(t_k), \eta(t_k)) = \aleph$

Denoting with $\tilde{x}$ the vector function collecting all observations, we assume that the raw signal is processed to a mixed-type dataset

$$X = g^{(2)}(\tilde{x}) \quad X \in \mathcal{X}^{m \times n} \tag{3}$$

where the instance space $\mathcal{X}$ is the set of all possible measurements, i.e. the codomain of $g^{(1)}$; in general, elements of $\mathcal{X}$ can be real-valued, categorical, binary, textual, etc. Hence, the dataset X consists of m different measurement channels which have one data point for each time $t_k \in \mathcal{T}$. Since we assume that $\aleph \notin \mathcal{X}$, this implies that $g^{(2)}$ is, or at least contains, an imputation function which replaces missing data with reasonable valid data. For simplicity, and abusing notation, we summarise the data generation process via the operator g and write

$$X = g(z, u, \eta) \tag{4}$$

as a shorthand for

$$X = g^{(2)}((g^{(1)}(z(t_i), u(t_i), \eta(t_i)); t_i \in \mathcal{T})). \tag{5}$$

Definitions

The literature attributes numerous different functionalities to DTs, e.g. simulation, control, optimisation, intervention capability. Here, we argue that the union of these tasks defines a DT in a general sense. Moreover, this work's focus on ML motivates us to assume that these tasks are *learnable* [176]. Hence, we define a DT in the following way.

Definition 1. *A digital twin is a collection of $r \in \mathbb{N}$ learned[2] functions*

$$\begin{aligned} &\mathcal{H} = \{h^{(1)}, ..., h^{(r)}\} \\ &\text{such that} \qquad\qquad h^{(j)} : \mathcal{X}^{m\times n} \times \mathcal{S}_j \to \mathcal{Y}_j \quad \forall j \in \{1, ..., r\}, \end{aligned} \tag{6}$$

where $\mathcal{Y}_j$ is the target space of and where $\mathcal{S}_j$ provides side information to the j^{th} learning task. In particular, the function $h^{(1)}$ was learned to accurately infer the state $z(t)$ of the PT, i.e.

$$\begin{aligned} ||z(t_k) - \hat{z}(t_k)|| < \varepsilon \qquad &\forall t_k \in \mathcal{T}, \forall X \in \mathcal{X}^{m\times n}, \\ \text{where} \qquad\qquad &\hat{z}(t_k) = h^{(1)}(X, t_k), \end{aligned} \tag{7}$$

for some small $\varepsilon > 0$.

This states that each function $h^{(j)}$ maps from (a subset of) our instance space to a specific target, e.g. a binary label or a real-valued number. The function $h^{(1)}$ in Definition 1 is central in the sense that it ensures that the DT produces a representation of the PT. Note that this representation *is not* a low-dimensional transformation of X as commonly used in the ML-field of representation learning. Instead, a DT has a virtual representation, which 'realistically approximates the geometries, physical properties, behaviors, and rules of the physical world' [158]. Since, by the notation we introduced, z summarises these properties and behaviours of the PT, (7) is a key defining equation of at DT.

To complete the bi-directional information flow, (some learnable function of) the estimated state $\hat{z}(t_k)$ needs to be compared against the desired (future) state $z^\star(t_k + \delta)$,with $\delta \geq 0$ and where $z^\star$ is either specified by a controller, e.g. to maintain system stability, or by a human, e.g. to perform an intervention. A comparison to a desired future state requires a forecast of how the PT will evolve. This *forecast* can be framed as yet another learning task, which we assume is solved by the learned function $h^{(2)}$, i.e.

$$\hat{z}(t_k + \delta) = h^{(2)}(X, u|_\tau, t_k, \delta) \tag{8}$$

for some hypothetical (future) control signal $u|_\tau$, $\tau = [t_k; t_k + \delta]$.[3] A successful DT will solve this learning problem accurately, at least for some finite horizon $\delta \leq \Delta$, where Δ may depend on the set of observations X.

Finally, there is the task of designing appropriate control signals $u|_\tau$ which either maintain $\hat{z}$ or implement a desired change. Note that there are usually many different suitable $u|_\tau$, and in concrete optimisation tasks there commonly is an additional objective or constraint which guarantees a unique solution. For instance, $u|_\tau$ may need to be bounded, or its total variation $\int_{t_k}^{t_k+\delta} |\dot{u}(t)|dt$ or energy $\int_{t_k}^{t_k+\delta} |u^2(t)|dt$ shall be minimised. Finding such an optimal $u(t)$ can be framed as yet another DT learning task, resulting in an (explicit or implicit) function h_k.

Combining all terms discussed in this subsection, we define a DT system as follows:

Definition 2. *A digital twin system is the tuple $(f, g, \mathcal{H}, z^\star)$ consisting of the PT, observation functions, a DT, and desired states. It satisfies the following two properties:*

$$\|z(t_n + \delta) - h^{(2)}\Big(g(z, u, \eta), u|_{\tau=[t_n;t_n+\delta]}, t_n, \delta\Big)\| \leq \varepsilon \tag{9a}$$

for all U and $\mathcal{T}$ and all $0 \leq \delta \leq \Delta(X)$, and where $z(t_n + \delta)$ is obtained by applying the control signal $u(t)$, $t \in [t_n; t_n + \delta]$ to f. Second, for every $z^\star(t_n + \delta)$, the DT determines an optimal control signal $u(t)$ such that

$$\|z^\star(t_n + \delta) - z(t_n + \delta)\| \leq \varepsilon. \tag{9b}$$

In words, a DT system consists of a PT and a DT such that

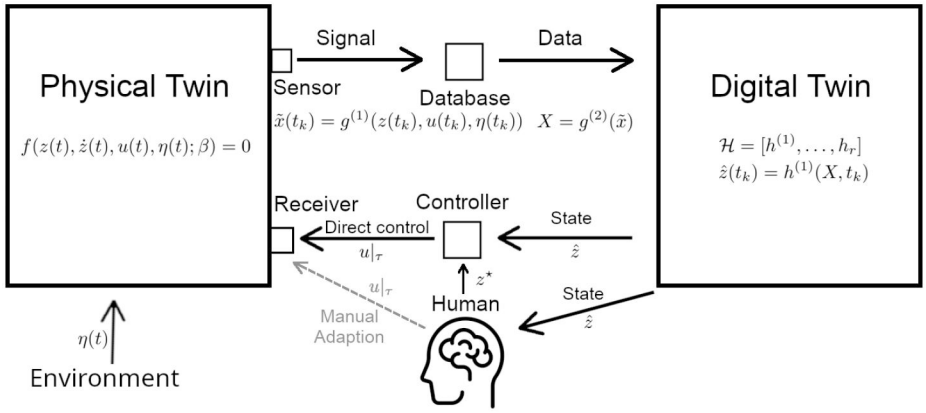


Figure 12. Formalised description of a DT system and its components. There is a cyclical flow of information, with data flowing from PT to DT, and control information flowing from DT to PT. The PT's state z is indirectly observed in the raw signal obtained from sensing technology and stored in a database. The raw signal is then pre-processed and passed to the DT to estimate $\hat{z}$ (since z is not directly measurable) and potentially to estimate other quantities of interest. Depending on the setting, the estimated state $\hat{z}$ is passed to a human or a controller. If intended, a human may update the desired state $z^\star$ and the controller then infers control signals to bring (or keep) the PT in this state. Alternatively, a human may also manually adapt the PT in some settings.

- the DT estimates the PT's current state using the observed data and can estimate its future state for arbitrary control signals (Eq. (9a));
- for a desired (future) state of the PT, there is a functionality for inferring control signals that achieve this desired state change (Eq. (9b)).

Intuitively, a DT system can be understood as a physical asset equipped with technology to digitise its state, and a mechanism for (automatically) adapting the physical asset based on the digital information. Figure 12 gives an overview of this definition and the surrounding notation.

Domain-agnostic description of DT components

This section describes all phases of the cyclical flow of information associated with DTs from an ML viewpoint—data layer, intelligence layer, information backflow, and physical and external requirements. No particular application is assumed, and concepts are described in a general domain-agnostic way. An overview is depicted in Figure 13.

Data layer

In this section, we describe the flow of information from PT to DT. The data collected from the PT must be transmitted, stored, and pre-processed before most ML systems can handle them. However, specific ML tasks also need to be solved at this stage, and the limits of computing must be considered.

Data acquisition and synchronization

To enable a data flow from PT to DT, it is necessary to measure and digitise information about the PT's current state $z(t)$. In many DT systems, this is achieved by equipping the PT with sensing technology [177] such as sensors, transmitters, and instruments, which results in a data stream. Formally, this data stream $g^{(1)}$ maps $z(t)$ and $u(t)$ to a raw data signal $\tilde{x}(t)$—note that in many cases, $u(t)$ is simply mapped through an identity function to preserve it for later analysis (cf. [178]). However, in realistic settings we cannot expect that z is fully measurable. For example, in wind turbine gears, it is not feasible to directly measure the gear's torque. Instead, it is widespread to use strain gauges to indirectly measure a signal which can be mapped to the torque via virtual sensing [179,180]. In many cases, these virtual sensors are learned functions, hence part of the DT—in cases where virtual sensors are obtained

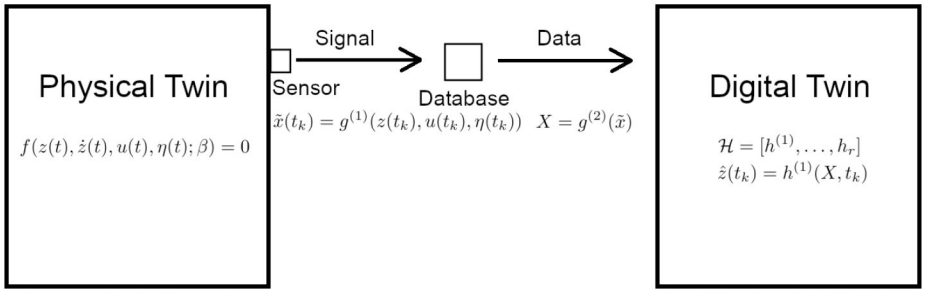


Figure 13. Data-related components of DT system.

from first-principles considerations, within the context of our work they can be included in the formulation of $g^{(1)}$.

In data streams, data may also arrive in batches, which can be obtained from a range of information management systems, such as customer relationship management platforms, electronic health records, or manually maintained registries. While batched data typically reflects historical or aggregated information about $z(t-\delta)$ for $\delta \geq 0$, it often also lacks precise synchronisation, i.e. it is not clear if we should associate $\tilde{x}(t)$ with $z(t)$, or $z(t \pm \delta)$. Data synchronisation for digital twins is an often overlooked problem that is insufficiently discussed in many DT papers:

Example 1 (Building Management Systems). *In [181], the authors collected heterogeneous data (temperature, humidity, motion, CO levels, equipment vibration, and others) using a wireless sensor network deployed inside a building. In the paper, data were used to synchronise the building and its digital representation, enabling real-time monitoring and the deployment of ML-based services such as predictive maintenance, anomaly detection, and asset usage optimisation. While the system effectively demonstrates the integration of streaming sensor data into a digital twin framework, it provides limited detail on how these data streams are temporally aligned with other heterogeneous sources.*

Example 2 (Human-robot Interaction). *Lv* et al. *developed in [182] a DT supporting collaborative human–robot assembly using data streams from RGB-D cameras, point clouds, and motion tracking systems. These data streams enable the synchronisation between the physical and virtual workspaces, providing real-time adaptation of the behaviour of the robot through ML-based robot-control policies based on deep reinforcement learning. Although the system effectively demonstrates effective PT-DT coordination, the description of how sensor inputs are temporally integrated is limited.*

Example 3 (Drilling Operations). *In [183], Brackel* et al. *presented a DT for drilling operations, where real-time sensor data from both surface and downhole equipment were collected using industrial protocols such as OPC UA and ModbusTCP. The data was used to (i) align the DT with the physical rig, and (ii) to enable ML-based functions such as performance optimisation and fault detection. While the paper addresses integration with heterogeneous devices, it does not detail how streaming and static data sources are coordinated over time.*

One strategy to successfully address the synchronisation problem in data acquisition is the Flooding Time Synchronisation Protocol (FTSP), or extensions thereof [184]. In simple terms, a key aspect of FTSP is a single designated node which broadcasts synchronisation messages that are flooded through the sensing technology. However, the synchronisation problem of coordinating the data streams of multiple digital twins is still an active research area [185].

Data cleansing

Realistic $\tilde{x}$ are often considered 'dirty' and are not suitable for most ML approaches [95,186]. This lowered data quality often arises because the data acquisition has functional limitations such as power and memory constraints, communication failures, measurement noise, and errors [187]. Additional causes are (harsh) environmental conditions of the PT [188] which may cause sensors to degrade, or limited spatial or temporal coverage by the measurements [189]. The interval at which individual sensors, actuators, etc. perform can differ, so the act of sampling the continuous-time state process $z(t)$ may

lose information via time-discretization. Furthermore, sampling may be event-based, hence non-equidistant, and different measurement channels may be temporally misaligned. Moreover, even if all channels are synchronised at some t_k, there might be a small drift or delay in one of the sensors which accumulates over time, requiring new synchronisation in the future.

This can lead to an unfortunate situation such as the following:

$$\left[\tilde{x}(t_1), \tilde{x}(t_2), \tilde{x}(t_3), \tilde{x}(t_4), \ldots\right] = \begin{bmatrix} x_1(t_1) & \aleph & \aleph & x_1(t_4) & \ldots \\ x_2(t_1) & x_2(t_2) & \aleph & x_2(t_4) & \ldots \\ \aleph & x_3(t_2) & x_3(t_3) & \aleph & \ldots \\ x_4(t_1) & \aleph & x_4(t_3) & x_4(t_4) & \ldots \\ \aleph & x_5(t_2) & \aleph & x_5(t_4) & \ldots \end{bmatrix},$$

where one only rarely observes all signal channels simultaneously. Missing data $\aleph$ or incorrect data [190] can bias analysis and cause failure in many ML systems [191]. Although there exist robust ML techniques which can produce reliable results despite the presence of missing or incorrect data, e.g. [192], it is preferable to work with a clean dataset since this permits the use of methods which have higher statistical efficiency (cf. [193]). Hence, in DT systems, it is necessary to pre-process the raw signals with data cleansing methods to obtain a high-quality dataset on which ML techniques can be performed.

Missing data. In the literature, missing values are commonly described as an obstacle to developing DTs [194,195]. One of the main problems is that in several applications, (implicitly) missing data are naturally part of the modelling task and cannot be avoided, for example in Earth sciences. This necessitates the use of imputation, interpolation, or simulation to obtain a high-quality dataset X. There are many ML techniques for imputing missing data, and it was recently shown that several of these techniques offer only minimal gain over conventional heuristic imputation methods for simple imputation tasks or small datasets [196]. However, in the context of DT systems, there are two areas where ML provides substantial improvements. The first area is super-resolution, where neural networks offer an elegant imputation approach. The underlying learning task for super-resolution is the following: Given partially observed (or sparse) input $\tilde{x}$, learn the maps $g^{(2)}$ and $h^{(1)}$ s.t. $h^{(1)}(g^{(2)}(\tilde{x})) \approx \hat{z}$, where $\hat{z}$ is the output of a physical model. Each of these maps can be learned by a suitable neural network, which was demonstrated in [197], where a convolutional neural network with an attention module and a localisation guided augmentation model were used. The second main area where ML can be effectively leveraged is to handle irregular or non-equidistant sampling in $\tilde{x}$, i.e. the difference $t_k - t_{k-1}$ is not constant for all k, or at least for some channels. Traditionally, this issue was approached with compressed sensing or (physical) differential equations, yet recently, it was demonstrated that neural differential equations [198] can be extended to irregularly sampled data [199], thus providing an efficient data-driven approach. Here, the task is to learn a map from $\tilde{x}$ to a so-called path, which is a natural cubic spline with knots at the observed times t_k, and then to map from the path to state z. Essentially, this is a refined interpolation approach which is directly integrated into a neural approximation of a differential equation.

Data contamination and robustness. Besides missing data, it is also not uncommon in DT systems that at specific times t_k, we observe $\tilde{x}(t_k)$ but some $\tilde{x}_j(t_k)$ are filled with incorrect values, e.g. due to noise or faults in the measurement devices [95]. Such data contamination can be more problematic than missing data, because it is difficult to detect [190] and can bias analysis/simulation results [200]. Directly removing data contamination is especially challenging in many DT systems due to the temporal dependencies between $\tilde{x}(t_k)$ and $\tilde{x}(t_{k+a}), a \in \mathbb{N}$. These temporal dependencies can cause data contamination to appear in a sequence $S = [\tilde{x}(t_k), \ldots, \tilde{x}(t_{k+a})]$ where each $s \in S$ might seem like a normal measurement but when viewed collectively, S is deviating [201]. Although many ML methods have been proposed for detecting such S, many are not reliable [202,203]. The development of robust DT techniques seems to be a promising research direction to finding and removing data contamination [204]. From the ML perspective, the idea is to use (potentially) contaminated data $\tilde{x}$ rather than clean data X as input to map $\mathcal{H}$ or as training data to learn $\mathcal{H}$. Clearly, it is not sufficient to use conventional

techniques to learn $\mathcal{H}$ from $\tilde{x}$ because contamination breaks typical (stochastic) assumptions. Robust neural networks [205], which abstain from making a prediction in difficult regions, are a suitable alternative for some settings, but research is ongoing. Recently, a robust DT system for power grids [206] was constructed using federated learning.

Data management

Data management can be essential for ML to permit efficiency in training and operation. In the context of DTs, an ML system uses observed data X to predict, monitor, optimise, and simulate physical systems f—via functions $\mathcal{H}$ learned from data. Data are acquired, stored, curated, prepared, and then analysed. These phases of the data science pipeline can directly impact the quality of the data used by DTs, and thus their overall performance [207]. Each stage of the pipeline requires the application of specialised data management techniques to preserve the integrity, consistency, and utility of the processed data, thereby mitigating negative effects on the performance of DTs.

Data storage. There are different challenges related to data storage infrastructures that support ML for DTs. In cases where DTs use real-time data to provide predictions, real-time data ingestion is a relevant topic. Other scenarios like the ones on which DTs are supplied by large volumes of batched data require scalability and long-term data preservation. In both scenarios (data streams and batched data), fault tolerance is also a significant challenge to address. Different studies have proposed different solutions to tackle the above mentioned challenges. In [208], a communication architecture for DTs is presented in the context of transportation based on Apache Kafka. The author demonstrates that Kafka enables the development of architectures that provide high availability, high throughput and low-latency real-time data ingestion. Jedermann et al. demonstrated in [209], that Apache Kafka can be combined with IoT protocols such as MQTT to supply live data efficiently to a pipeline that includes a DT to monitor fruits during sea transportation, and an InfluxDB for long-term retention and historical data analysis. Long-term data preservation is a requirement in domains such as environmental monitoring and smart grids, where historical data is required for model training, historical analysis, and regulatory compliance. Eshaghi et al. [210] discuss how cloud storage services, SQL and NoSQL on-premise databases can fulfil the storage requirements of DTs applied to environmental monitoring. These technologies can be combined in hybrid infrastructures to provide real-time analytics and historical forecasts. In [211], Mchirgui et al. review the role and implementation challenges of DTs in the smart grid domain. Among other things, the authors emphasise the need for scalable cloud/edge storage architectures to provide data to DTs that handle large volumes of data.

Extraction, transformation and loading (ETL). DTs operate with a wide variety of heterogeneous data. The effectiveness of a DT depends on its ability to integrate and process this dissimilar information. In this context, ETL workflows emerge as a fundamental component for preparing and unifying these data, making them suitable for consistent use within the DT ecosystem. In [212], Yan et al. presented a common data environment to facilitate data management in DTs for smart cities. Their approach stresses the need for a robust 'data pipeline' that enables ETL processes of heterogeneous data types (building information modelling, geographic information systems, IoT sensors, drone data, etc.). Also, the paper discusses some of the main challenges of ETL processes in DTs: 1) lack of interoperability between data from different sources and actors, and 2) the need to resolve syntactic and semantic heterogeneity to ensure data consistency and quality. De Conto et al. [213] introduced 'Function + Data Flow' (FDF), a domain-specific language devised to simplify the development of artificial intelligence pipelines in the domain of DTs. FDF can handle hybrid data by combining information from physics-based models (e.g. data from finite element simulations) with sensor measurements. Although FDF does not focus on conventional ETL processes, its relevance lies in the effective orchestration and manipulation of ML functions and models in a manner analogous to traditional ETL processes. This characteristic of FDF facilitates the data preparation for use in highly complex DTs. Merino et al. addressed in [214] the challenges inherent to data integration for DTs in the specific domain of construction. The article highlights that the application of ETL in the architecture, engineering and

construction domain, presents particular difficulties due to the natural segregation and isolation of data in silos composed by domain-specific systems such as Building Automation Systems, and document databases. To mitigate this fragmentation, the authors propose a method based on the implementation of federated data models and ontologies. This architecture, based on the use of data lakes, facilitates the ingestion of information from multiple sources through independent ETL processes, thus standardising the integration of complex and heterogeneous data for a unified representation in the DTs. Overall, we observe that suitable protocols, tools, and services are readily available for enabling data management for ML in the context of DTs. This availability has spread over the last years, which partially explains the surge in research on ML-based DTs.

$\mathcal{H}$—intelligence/simulation layer

In this section, we describe the DT in more detail. An overview is given in Figure 14. Several aspects are relevant to describing the DT's functionalities—essentially, per Definition 1 it is a list of functions $\mathcal{H}$ where at least one function provides an estimate $\hat{z}(t)$ of the current state of the PT. Stemming from the PT, it is further relevant to equip the DT with the capability of modelling physical phenomena not just in the sense that errors are minimised, but that the DT also adheres to laws of nature, and permits causal modelling.

Forecasting and simulation

Forecasting a future state $z(t+\delta)$ of the PT is relevant in many DT applications. Indeed, if one knows that some forecast $\hat{z}(t+\delta)$ is close to the true $z(t+\delta)$, one can leverage this information to enhance decision-making [10,215] and implement predictive maintenance for the PT [216]. From a technical perspective, we can obtain a forecast that improves upon a random guess (in expectation) if and only if data X contain measurable temporal dependencies, i.e. for some $t_i, t_j \in \mathcal{T}$ we have

$$P(X(t_i)|X(t_j)) \neq P(X(t_i))P(X(t_j)).$$

This requirement is typically fulfilled, e.g. due to a predictable trend in X, but also many stationary processes can be considered—for example Markov chains. Furthermore, this requirement is fulfilled for many control problems, as the control signal $u(t)$ may be explicitly included in X as per eq. (2). In a viable setting with such dependencies, we can obtain forecasts by learning a dedicated function $h^{(2)}$,

$$\hat{z}(t_k+\delta) = h^{(2)}(X, u|_\tau, t_k, \delta), \tag{10}$$

which shall be close to the true $z(t+\delta)$ for all control signals U and for δ within a maximal forecasting horizon Δ, cf. Definition 2

The ML task of forecasting allows us to draw a logical link to simulation: Let us consider the physical perspective where the development of z is described by a DAE system (1). If f, $z(t)$, t, $\eta(t)$ and β are known, we can use knowledge about the control signal $u|_\tau$, $\tau = [t; t+\delta]$ and solve (1) to obtain exact states $z(t+\delta)$ in the future of t. In the more realistic setting where only the control signal $u|_\tau$, an

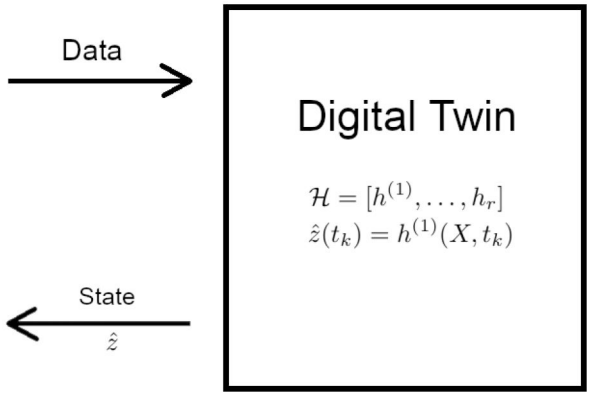


Figure 14. Intelligence and simulation components of DT system.

approximation $\hat{\beta}$ of the DAE parameters, and an estimate $\hat{z}(t)$ of the current state are known, the solution of (1) using, e.g. Euler's method [217], yields an approximation $\hat{z}(t+\delta)$ of the future state. Essentially, this approach is the numeric counterpart of data-driven forecasting $\hat{z}(t+\delta)$ using via the learned function $h^{(2)}$.

If measurement and data processing are well understood, i.e. $g^{(1)}$ and $g^{(2)}$ are available, then forecasting techniques can also be used for simulation of *What-If* scenarios. Suppose we are interested in the consequences of bringing the PT in a specific state. For example, it could be of interest if a constructed bridge can bear several heavy trucks while exposed to strong lateral wind. If the hypothetical initial state of this scenario $\tilde{z}(t)$ is known, we can leverage the DT's forecasting capability to simulate this scenario. More explicitly, rather than using first-principles simulations of f starting from the hypothetical initial state $\tilde{z}(t)$ and the scenario-defining input $u|_\tau$—which requires knowledge of the system's DAEs f—one can rather use knowledge of g to map the hypothetical initial state $\underline{\tilde{Z}}(t)$ to some hypothetical observations $\tilde{x}$. Then, the forecasting function (10) can be used on $\tilde{x}$ and $u|_\tau$ to simulate the scenario in the sense of obtaining an estimate of $z(t+\delta)$.

There are numerous techniques that could—in principle—be used as components of $h^{(2)}$ since X is a batch from a data stream and hence closely related to time series. These techniques include classical statistical models such as (variants of) ARIMA and exponential smoothing [218], as well as ML-based techniques such as gradient boosting machines [219] and deepAR [220]. However, the concrete choice for $h^{(2)}$ is not arbitrary and should be harmonised with the DT system's properties. Indeed, unlike typical forecasting tasks, DT systems are often governed by complex physical equations that make purely data-driven forecasting non-trivial. Hence, while black-box ML techniques such as deepAR can achieve strong performance on benchmark datasets, they may be insufficient for forecasting the behaviour of DT systems and produce forecasts that violate governing equations [221]. To address this issue, one must either couple these black-box techniques with numerical models following the DT system's governing equations, or incorporate these equations during the design and training of ML models such as $h^{(2)}$.

Methods for learning/including governing laws

The literature offers a wide array of approaches that, in contrast to purely data-driven models, aim to incorporate prior knowledge into the design and training of ML models. Collectively referred to as knowledge-guided [222], theory-inspired [223], or physics-informed ML (PIML), these methods can be broadly categorised into three complementary strategies for embedding prior knowledge into learning algorithms [29]: observational bias, inductive bias, and learning bias. These categories are not mutually exclusive and can often be employed in combination within a single application to enhance model fidelity, generalisability, and consistency with governing equations.

First, the observational bias is addressed through the use of specific data and targeted data transformations that bring the data into a representation more closely aligned with the underlying laws. Often serving as the first guiding principle in the design of a DT, a suitable data representation influences every subsequent step—including the structure of the ML model and the way first principles are incorporated via the training data. These transformations are effectively realised through the data preprocessing function $g^{(2)}$ in (3). They can be employed, for instance, to account for symmetries or invariances reflected in the PT ([224], Chapter 3), [225]. As a concrete example, raw sensor data obtained in a radially symmetric setting may benefit from a coordinate transformation—e.g. from Cartesian to cylindrical or spherical coordinates—that enables a mathematical simplification of the system. This can be advantageous for using reduced-order models [226,227] or for simplifying the learning task itself [228]. A relevant example of physically-informed data preprocessing is nondimensionalization [229], which typically simplifies problems by expressing physical quantities in terms of characteristic scales. This not only reduces the number of free parameters but also highlights the dominant physical effects. Furthermore, nondimensionalization helps recover characteristic time and length scales or resonance frequencies, and yields a system representation that is more naturally aligned with the mathematical structure of governing differential equations.

Second, the inductive bias is introduced through the deliberate choice of ML models and architectures that inherently encode the laws governing the PT. This bias can manifest at various levels—from high-level decisions such as the selection of a particular ML model for the task at hand, to more granular choices like the design of the network architecture or the use of specific activation functions. An important example involves the conservation of physical quantities, expressed in simplified form as $\sum_i z_i = C$. A straightforward way to enforce this conservation law is through the use of a softmax layer. When the outputs of a model are passed through a softmax function and scaled by the conserved quantity C, the result is: $C \cdot \sum_i \sigma(z)_i = C$, where σ_i denotes the softmax activation of the i-th output neuron, and z is the pre-activation input vector. This construction ensures that the sum of the outputs remains conserved, regardless of the specific values predicted [230]. Importantly, inductive bias provides a robust mechanism for embedding governing laws directly into the ML model. Once integrated, its effect is independent of the training data distribution or learning procedure—it is inherently encoded in the model's structure and therefore consistently enforced. Further examples of inductive bias include Lagrangian [231] and Hamiltonian neural networks [232], which embed conservation principles such as energy or momentum into the model dynamics, as well as equivariant networks, e.g. [233] and other physics-informed architectures [234–236].

Third, the learning bias—the encoding of governing laws directly into the training objective—has gained significant attention in deep learning, largely due to the availability of efficient training algorithms that allow neural networks to optimise complex loss functions. Arguably the most prominent example of learning bias in PIML is the class of physics-informed neural networks, or short PINNs [237]. PINNs have demonstrated remarkable flexibility and effectiveness in directly incorporating governing differential equations into the loss function of a neural network. Leveraging automatic differentiation, PINNs include a term in the loss function that penalises candidate solutions violating the PT's governing function f as defined in (1). This term, when minimised along with data-driven losses and boundary or initial conditions, encourages the network to learn solutions that conform to the behaviour prescribed by f. Furthermore, the PINN framework also allows certain parameters of f—such as the system parameters β—to be treated as trainable variables. These parameters can thus be inferred directly from the data during the learning process, making PINNs particularly well-suited for solving inverse problems, where unknown physical quantities must be estimated from observed system behaviour [238,239]. Despite the fact that PINNs have been shown to suffer from severe training instabilities, cf. [240–243], they have been successfully applied across a wide range of scientific and engineering disciplines [244–246].

Causal considerations

The initial insight stemming from the causal perspective concerns (2). This formula reflects the process of observing the PT via $g^{(1)}$, an act that is inherently destructive. Among other things, the cause-effect relation loses the direction and can only be observed as correlation. The task of causal discovery aims to reconstruct the causal structure from observation, with a wide array of methods being proposed [247]. Many of these methods rely on strong assumptions, for example, additive noise models assume that $g^{(1)}(z(t_k), u(t_k), \eta(t_k)) = g^{(1)}(z(t_k), u(t_k)) + \eta(t_k)$.

Causal modelling is the task to collect existing causal knowledge, for example the laws f governing the PT. Depending on the available information, two levels of abstraction are suitable for causal modelling. 1) Structural causal models [248] provide structural equations for all endogenous variables that also model the influence of exogenous noise variables. These equations represent the causal mechanism, but are elusive if the mechanisms are not available or fully understood. 2) If only the causal variables and their causal relations, i.e. the causal structure, is known, causal diagrams [249] are a suitable modelling choice. Conceptually, such causal diagrams are directed acyclic graphs (DAGs). This representation allows for understanding bias and has implications on ML as implementation of the DT.

A mechanistic/physical model of the DT arguably requires much knowledge about the PT. At the other end of the spectrum a purely data-driven, statistical model does not require domain knowledge about PTs to provide predictions, but commonly assumes that data are independent and identically distributed (IID) and that no shifts in distribution are expected. One can argue that this assumption is

not realistic in a DT system. On the one hand, the IID assumption is often violated due to the fact that the observations X are almost always (temporally dependent) data streams. Further, in (9a), $u(t_n)$ represents a control signal that is often based on the output (e.g. prediction) of the DT. Causally speaking, this constitutes an intervention, which may—in addition to introducing a dependency between $z(t_n)$ and $z(t_{n-1})$—change the distribution of X. This informs us that for modelling the DT, traditional ML models that rely on the IID assumption are not suitable to conduct 'What-If' analyses, or introduce an unwanted bias. Instead, ML models should be informed and constrained by the available causal knowledge to avoid bias and establish the desired generalisation abilities.

One noteworthy approach of causal models inspired by additive noise models refers to the measurement noise during observation. In half-sibling regression [250], the aim is to denoise the observed signal X and remove the influence of an unobserved confounder, e.g. as represented by $\eta(t)$. This requires one of the m measurement channels to be directly influenced by the state of the PT of interest, which is only available under the influence of noise stemming from an unobserved confounder. This confounder also affects other measurement channels. A regression of the other measurement channels on the target channel helps to reduce the noise and provide an improved estimate of the state of PT.

Apart from informing the modelling of ML for DTs, causality also helps in addressing communication bottlenecks by tailoring the representation to align with semantics of the domain. In [251] they propose a causally informed communication based on imitation learning, aiming for semantic communication.

DT systems also pose unique challenges for causal modelling. Modelling the interaction between the PT and the DT is not possible with traditional methods, such as structural causal models or DAGs, as the closed loop between PT and DT constitutes a cycle. Hence, methods that relax causal models for cyclic relationships are required, with promising works already being proposed [252].

Information backflow

Even with all its predictive capabilities, a DT has little purpose of its own. To benefit from a DT, the information about the PT stored in the DT must be used to assess or adapt the PT. In the following, we discuss control mechanisms to reach desired states, interventions to discover latent system variables, and how to present the information in the DT to enable a human to specify a reasonable desired state $z^\star$. An overview is visible in Figure 15.

Human control vs. direct control

According to the now classical interpretation of Michael Grieves and John Vicker [253], a DT system is primarily used to make information about the PT accessible without needing a human in the PT's immediate proximity, thus enabling remote decision making. Another interpretation is that of a coupled DT system [86], where DT and PT mirror each other's actions, and where the DT can directly control the PT when needed or desirable [178]. Although these two interpretations seem conceptually different, we now demonstrate that—through the lens of ML—these two interpretations of human-

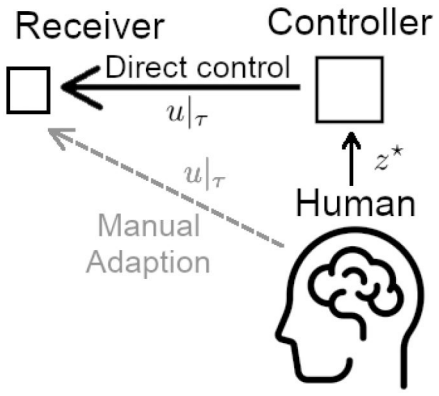


Figure 15. Human and control components of DT system.

centered information presentation and automatic control are intrinsically linked, representing complementary perspectives of the same underlying concept.

Suppose we have a DT whose functions $h^{(1)}, h^{(2)}$ give reasonable approximations of the PT's current and future states $z(t)$ and $z(t+\delta)$ for $\delta > 0$. Regardless of whether a human or the DT should control the PT, either has to provide a suitable control signal $u(t)$ to guide the PT into the desired state $z^\star(t+\delta)$. Since the current state $z(t)$ is not fully observable or directly measurable, both will have to instead base their control signal on an estimate $\hat{z}(t)$. In any case, the relationship how $u(t)$ influences the state $z(t)$—which is explicit in the DAE (1)—must be understood to find a control signal that causes $z(t)$ to change towards $z^\star(t+\delta)$.

One can summarise this need in a function ν that maps the current state estimate $\hat{z}(t)$ and the desired future state $z^\star(t+\delta)$ to a function $u|_\tau$, $\tau = [t, ; t+\delta]$ constituting the corresponding control signal. Such a function ν—be it explicit or implicit, e.g. via the solution of an optimisation problem—enables the DT to compute a suitable control input for the PT, but interestingly, it enables a human to do the same. Although a human will undoubtedly process the available information differently, and will tap different information sources, if we fix the desired state $z^\star(t+\delta)$, then both human and DT are necessarily required to find a suitable ν_{Human} and ν_{DT}, respectively. While it is unlikely that V_{Human} matches ν_{DT} (the desired change in $z(t)$ may be achieved in several ways, hence ν is usually not unique), both functions must achieve $z(t+\delta) = z^\star(t+\delta)$. Hence, through the lens of ML, the only difference between human control and direct control is that a human exclusively defines $z^\star(t+\delta)$—or formal criteria which define what constitutes $z^\star(t+\delta)$—whereas the underlying task of learning a suitable control function as per Eq. (9a) is the same for human control and direct control.

Direct control and optimization

We shall now focus on direct control, i.e. we assume that the task of determining the optimal control signal $u(t)$ is solved by the DT. If a desired state $z^\star(t+\delta)$ is given, then it is a major learning task to identify a function ν_{DT} as described above. Indeed, a function that maps the desired state and the matrix X of observations into the space of control signals can be written implicitly as the solution of the following optimisation problem:

$$\begin{aligned} \min_{u|_\tau} \quad & \gamma \|\hat{z}(t+\delta) - z^\star(t+\delta)\|_2 + (1-\gamma) R(u|_\tau) \\ \text{where} \quad & \hat{z}(t+\delta) = h^{(2)}(X, u|_\tau, t_k, \delta). \end{aligned} \tag{11}$$

In (11), the term $R(u|_\tau)$ is a regularisation term on the control signal, e.g. its energy $\int_t^{t+\delta} \|u(t)\|_2^2 dt$. In essence, this optimisation problem involves moving the DT into the desired estimated state in a way that ensures a smooth transition and adheres to the laws of nature governing the PT. The optimisation problem (11) is parametrised by γ, which defines the trade-off between an accurate match[4] of $z^\star(t+\delta)$ and a smooth transition from $\hat{z}(t)$.

Optimisation problems similar to (11) are rather new in the ML literature and have only been tackled recently. A first step is taken in [254], where f is assumed to be an ordinary differential equation system. The proposed solution is a dual network architecture, where one part focuses on optimising the control signal, and the other on solving the differential equations and estimating the state. Another recent work [255] focuses on so-called neural DAEs, which are an extension of neural ordinary differential equations, which—in turn—are neural methods for approximating ordinary differential equations. The authors propose to use simultaneous collocation, which is a novel method for training PINNs [256]. Although all of these approaches seem promising candidates for optimising DT control, further research and rigorous review is needed to validate their effectiveness.

For example, it is intuitive to think of $z^\star$ as a desired position that the PT should take, e.g. an autonomous vehicle parking at a specific spot. However, $z^\star$ could also be linked with internal properties of the PT, e.g. energy consumption. In such a setting, it is common that $u(t)$ controls the power provided to the PT [10,178], and hence $u(t)$ may itself affect the optimisation objective, e.g. if one seeks to minimise energy consumption. This is possible by selecting a small γ close to zero, which encourages the use of control signals which small norm and allows the DT system to operate efficiently.

Interventions

While control focuses on achieving or maintaining a desired state $z^\star$, intervention testing focuses on studying the impact of changes to the DT system. As such, interventions are a counterpart to system control, since here we do not study alternative control signals $u(t)$ that are necessarily designed to achieve or maintain $z^\star(t)$. Instead, we are interested in determining how a specific $u(t)$ affects the estimated state $\hat{z}(t)$ and measurements X to uncover their hidden relationship. This is a powerful setting that goes beyond the possibilities of classical data science, since controlled interventions permit the discovery of cause-effect relationships [257] (see Section on causal considerations) and latent variables such as β [258].

More concretely, using interventions one can—in theory—tackle the problem of discovering DAE f, as well as determining environmental influences $\eta(t)$ which affect both the physics and the observational measurement function $g^{(1)}$.

This possibility is a key benefit of DT systems that is not achievable in other systems

Many approaches based on PIML, e.g. [255,256] assume that the physical equations f are fixed, yet in a DT system we can observe, test, and validate f. This opens a potent avenue for future research on DT-driven physics. To proceed in this direction, it is required to overcome the limitations of existing causal ML techniques due to the cyclical nature of DT systems. Moreover, from an ML perspective, intervention testing is often out-of-distribution, which challenges conventional learning theory [259]. First steps in this direction were taken in the medical domain, where intervention testing is firmly established. In [260], partial DTs of humans based on ML techniques were tested to propose adaptions to the human's lifestyle such as dietary changes. However, further research is needed to explore the high potential of DT-based intervention testing.

Human-centered information presentation

Since ultimately a human must specify the desired state $z^\star$ of the PT, it is essential to prepare and present the information generated by a DT in a manner that allows humans to make informed decisions and interact effectively with the PT to fully capitalise on the benefits of a DT. Given the inherent complexity of the ML models underpinning DTs, merely outputting predictions or raw data is often insufficient. Instead, user-centered information presentation and explainable AI (XAI) are important for boosting trust, gaining new insights, and encouraging people to use DT systems in real life.

ML models, especially deep learning approaches used in DTs for forecasting and control (cf. Definition 2), typically operate as black boxes. The opacity of the DT can significantly impede the human capacity to comprehend, validate, or troubleshoot its functionality. Accountability and transparency are fundamental ethical considerations, particularly for DTs of humans, social systems, or critical infrastructure. These challenges are addressed by XAI methods, which offer methods to render the decisions and predictions of ML models comprehensible [261]. For a DT, this means making understandable not only what is predicted but also why a specific prediction or control recommendation $u(t)$ was generated. This is particularly relevant when the human operator ultimately assumes control over the PT. Furthermore, the reliability of these ML models is inherently related to the quality of the input data, processed through the data layer. .

The quality of the processed data is crucial for the effective presentation of DT information, which combines interpretable ML techniques, embedding, and sophisticated visualisation. It is imperative that these data are effectively managed throughout the pipeline to ensure the overall performance of DT systems [207].

- *Interactive Data Visualisation:* Beyond static charts, interactive dashboards offer the capability to dynamically explore complex datasets and model outputs. This includes visualising the data stream $\tilde{x}(t_k)$, the inferred state $\hat{z}(t_k)$ and the forecasted state $\hat{z}(t+\delta)$. 3D visualisations of $\hat{z}$ can provide intuitive insights into the geometric and physical properties of the PT, highlighting deviations. A clear indication of the prediction intervals or uncertainty surrounding predictions should also be included in these visualisations; this is particularly important in industrial applications where dependability is crucial.

- *Explainable AI Interfaces:* To guarantee interpretability, DT systems should directly incorporate XAI outputs into the user interface. This can encompass three methods: First, feature importance can indicate which input features in X contribute the most to a particular prediction or control recommendation. This makes it easier for users to comprehend the key features of the data pipeline. Second, local explanations can provide case-specific rationales, such as counterfactual explanations that exhibit which minimal change in inputs would have resulted in a different outcome. This is essential for debugging and comprehending particular operational scenarios. Third, scenario analysis enables users to simulate 'What-If' scenarios to understand the impact of various control strategies $u(t)$ or external disturbances $\eta(t)$ presenting the results in an easily digestible format. Finally, the outcomes of these analyses must be in accordance with the laws of nature.
- *Contextualisation and Actionability:* Integrating governing laws can facilitate visualising the integration of boundary conditions or conservation laws and the impact of their adherence or violation on the results. This helps users reconcile ML predictions with the known physical reality. Levels and warnings can transfer model predictions into clear, actionable alerts when important levels are crossed. This is in direct correlation with the necessity of robust data cleansing and anomaly detection. Concerning recommendations for $z^\star$, when the DT performs control tasks or generates recommendations for human interventions, the interface should clearly state the recommended action and its expected impact and the underlying rationale. This aligns with the discussion above regarding direct versus human control. Providing metrics on the quality of the incoming data, such as information about misspelled values, imputed data, and detected anomalies is further helpful for contextualisation. This allows users to assess the trustworthiness of the DT's outputs, especially when the DT relies heavily on $g^{(2)}$ for imputation [95].

There are numerous obstacles associated with the implementation of human-centered, effective information presentation in DT systems. The presentation must strike a balance between the complexity of the underlying models and the simplicity of the presentation. Oversimplification can result in the omission of critical information, while an excessive amount of detail can be overwhelming for the user. Additionally, real-time interpretability is required for DTs to function in real-time, which can be challenging in dynamical systems that are developing quickly. The standardisation of metrics for explainability in the DT context, the development of adaptive interfaces that adapt to user expertise, and the integration of multimodal feedback (e.g. audio or haptic signals) to enhance decision-making are all potential future research directions. To transform DTs from powerful analytical tools to trustworthy and collaborative partners for human users, the synergy between XAI techniques and advanced user interface design will be essential. This will ensure that the insights derived from the data layer are clearly communicated and understood.

Physical and external requirements

In the following, we discuss and highlight several requirements that restrict ML-based modelling the PT. These requirements may stem from the nature of the PT (which is often, but not always, a physical object), from the intended application of the DT system (e.g. in industrial settings), and from the fact that DT systems are interacting—directly or indirectly—with humans, leading to ethical and legal challenges. An overview is given in Figure 16.

Physical requirements

We now focus on the PT f and its (potentially time-varying) environment $\eta(t)$. From an analytical perspective, eq. (1) is central, as it—in theory—fully describes the PT. However, in practice f and η are rarely known, which impedes exact modelling and motivates approximations. The physical reality of the PT commonly imposes constraints which must be satisfied.

From a physical perspective, several factors can critically influence the efficacy of learning tasks associated with the bidirectional flow of information between a PT and its corresponding DT. Central to this discussion is the ability of the DT to accurately represent the PT with sufficient temporal (and spatial) resolution [262]. It is important to recall that the PT's state $z(t)$ evolves in continuous time, i.e.

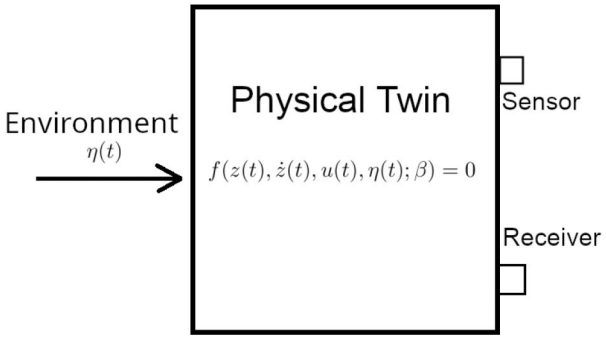


Figure 16. Physical and external parts of DT system.

$t \in \mathbb{R}$, as governed by f. In contrast, the DT can only observe and represent this state at discrete time instances $\hat{z}(t_k)$, where $t_k \in \mathcal{T}$ corresponds to a finite sampling set. This disparity introduces fundamental limitations in capturing the full dynamics of the PT, particularly in systems with fast-evolving or high-frequency behaviour [263–265]. In general terms, if the PT's inherent process evolves over time scales shorter than the sampling rate of the data acquisition system, meaningful learning becomes infeasible [266]. More precisely, the underlying system dynamics f are characterised by specific temporal (and spatial) scales that must be adequately resolved for the DT to provide an accurate representation through $\tilde{x}(t_k)$ and allow successful adaption through control in (11) [30,267].

Many physical systems exist for which high temporal resolution is not critical. Diffusion-driven processes—such as drug absorption [268], air pollution dispersion [269], or catalytic reactions [270]—typically evolve over long time scales. In such cases, there is sufficient time to measure and process the state of the PT, assuming reasonably capable hardware. In contrast, other systems are governed by much shorter characteristic time scales, which impose stringent requirements on the hardware infrastructure. For example, processes such as combustion [271], neural activity [272], or high-frequency mechanical vibrations [273] exhibit stiff dynamics and can unfold within millisecond or even sub-millisecond intervals. These cases demand high-frequency sensing and low-latency data processing to support timely and effective control interventions [274,275]. Additionally, when DTs are deployed on edge devices, further constraints arise due to limited computational resources and power budgets. Ensuring real-time data acquisition and processing under such conditions requires careful selection and optimisation of the hardware platform [276].

Moreover, additional challenges arise when modelling chaotic systems—a class that encompasses a wide range of natural and engineered processes, from turbulent fluid flows in mechanical systems to population dynamics in biological and ecological contexts [277,278]. Even when the DT provides an accurate short-term approximation of the PT—that is, when $\hat{z}$ closely matches z—the intrinsic sensitivity of chaotic systems to initial conditions imposes fundamental limits on the predictive time horizon as governed by f. Small perturbations—whether introduced directly through the control input u or indirectly through unmodeled effects represented by η—can quickly cause significant divergence between the DT and PT trajectories. As a result, learning tasks in such contexts must be carefully constrained to reduced time windows over which predictive accuracy and control effectiveness can be reasonably maintained [279,280].

Continuing the discussion on physical requirements, additional considerations are necessary to ensure a stable, long-term connection and effective bidirectional flow of information between the PT and the DT. The PT f is often subject to fundamental physical principles—such as conservation laws, symmetries, and invariances under certain transformations (e.g. translations and rotations). Ideally, these physical principles should also be reflected in the DT by means of physics-informed approaches as discussed above. Failing to do so may introduce discrepancies into the bidirectional data flow, ultimately leading to diverging states and an increase in the error term ε in Equation (9b) over time. Nonetheless, it is often unclear how and where to enforce such conservation laws within the DT framework. Moreover, measurement noise and environmental disturbances—represented by $\eta(t)$ in (1) and (2)—may prevent these physical constraints from being directly inferred from the raw sensor data $\tilde{x}$. In

such cases, it becomes necessary to incorporate these laws either through the data preprocessing function $g^{(2)}$ in Equation (3) [48,281], via the design of the ML model itself [230,231], or directly within the learning process defined in Equation (9a)—for example, by employing methods from PIML [29,237].

Ethical and legal requirements

The ambiguity and lack of transparency in the application of ML to DT systems give rise to significant ethical challenges in the light of trustworthy AI [173]. Moreover, there is a growing body of legislation concerned with ML and AI [282], such as the General Data Protection Regulation (GDPR) and the EU AI act [283] as well as many smaller AI bills [284]. Therefore, this subsection addresses the ethical dimensions by first identifying the types of DT systems where these considerations are most pressing and then proposing how ML techniques can be adapted to meet ethical and the associated legal requirements.

For which types of DT systems are ethical and legal considerations relevant?. We identify three main types of DT system, in which ethical considerations are especially relevant: Digital Twins of Human Beings: The most ethically sensitive category is the 'digital patient' or any DT that models an individual's physiology, health status, or behaviour. Such DTs, often proposed for personalised medicine or lifestyle coaching, handle extremely sensitive personal data. The core ethical challenges include privacy and consent, effects on user autonomy, and the potential for bias if models are not representative of the population [285]. With respect to *Privacy and Consent*, questions like to following arise: How is data collected, stored, and used? Is the consent informed and ongoing, especially as the DT evolves with new data? For *Autonomy and Agency*, we refer to the central question of whether a DT's ML model predicts a future health risk, and how it affect an individual's choices and autonomy. In terms of *Bias and Equity*, we argue that if the ML models are trained on data from a specific demographic, the resulting DT may be less accurate or even harmful for individuals from underrepresented groups, exacerbating existing health disparities. Digital Twins of Social Systems and Human Environments: This category includes DTs of smart homes, office buildings, or even entire cities. These systems model the interaction between humans and their environment, making them ethically complex. There are several key ethical issues to consider here [286]. A DT of a building or city inherently involves monitoring the behaviour of its occupants, hence it is connected to *Surveillance and Control*. This raises concerns about mass surveillance, a loss of privacy, and the potential for behavioural manipulation. In terms of *Fairness and Justice*, ML models used for resource allocation (e.g. public transport) could mimic or even amplify existing societal biases, directing resources towards wealthier neighbourhoods and neglecting others (cf. smart city DT). Finally, questions regarding *Accountability* arise if things go wrong (e.g. resource depletion due to inadequate allocation). The lack of clarity makes assigning accountability nearly impossible, which is a critical gap in algorithmic governance [287]. Digital Twins of Critical Infrastructure: This category includes DTs for assets like bridges, power plants, and autonomous vehicle fleets. While not directly modelling individuals, their failure can have severe consequences for human life and societal stability [288]. One ethical consideration is around *Safety and Reliability:* The central ethical mandate is to ensure that the system is safe. Given the opacity of many ML models, the main question is how can we guarantee their reliability? Another ethical consideration concerns *Accountability and Transparency:* In the event of a failure, it is ethically imperative to understand why it occurred. A 'black box' ML model that cannot explain its reasoning, which makes post-mortem analysis and future prevention extremely difficult.

How can we adapt ML techniques to meet ethical and legal requirements?. Addressing these ethical challenges requires moving beyond a narrow focus on predictive accuracy and incorporating ethical principles directly into the ML development lifecycle. This involves adapting the employed techniques in the following ways:

- *Transparency and Explainability*: The lack of clarity in current DT literature must be countered with a commitment to XAI. Instead of using opaque models, we should prioritise interpretable alternatives or employ post-hoc explanation techniques to understand model behaviour [261]. As an example, for a DT monitoring a bridge, an XAI system could highlight which specific sensor readings and physical parameters led to a damage prediction, allowing engineers to verify the reasoning.
- *Privacy Preservation:* For DTs of humans or social systems, the default mode of data handling must shift towards privacy preservation. This can be achieved through techniques such as federated learning and differential privacy [289,290].
- *Bias and Fairness:* To address bias in DTs that affect humans, ML models must be explicitly optimised for fairness alongside accuracy. This could involve bias audits (i.e. systematically auditing datasets for demographic imbalances and stereotypical correlations before any model is trained) and fairness-aware models (i.e. employing ML algorithms designed to mitigate bias by enforcing fairness constraints) [291].
- *Compliance:* Relevant existing legislation must be considered in advance of developing an ML-based DT system. ML systems must be designed in a way to allow easy adaption to foreseeable regulatory requirements [292]. Since ensuring full compliance of DT systems is likely not achievable [293], legal stakeholders should be involved and consulted in appropriate ways to prevent compliance risk during operation.

Summed up, to leverage ML in DT, we must treat transparency, privacy, fairness and compliance as primary design goals. This requires a fundamental shift in how ML techniques are selected, implemented, and validated within the DT framework, moving from the current state of ambiguity to one of clear, accountable, and human-centric design. These ethical considerations in the light of trustworthy AI are also important when considering the certification of ML-backed DTs [294].

Summary and recommendations

In this paper, we have reviewed DTs through the lens of ML. We have described DTs in the domains *Earth Sciences*, *Healthcare*, and *Industry*, with a focus on data, learning assumptions and learning objectives. We have reviewed 50 recent ML-based DTs, highlighted trends, key technologies, opportunities, and challenges. We have given a formal definition of DTs intended as starting point for dedicated interdisciplinary research on machine learning for DT. We have described all components of a DT system in a domain-agnostic setting with a focus on ML, including discussions of promising niche topics such as physics-informed ML. The main motivation for this endeavour was to facilitate interdisciplinary research and to encourage ML researchers to investigate DT settings. As such, our work serves as guideline for machine learning researchers and practitioners who research, design or contribute to DT systems.

Despite the inherently applied and highly specific nature of DTs, we observe that digital twin systems are amenable to statistical learning theory, and that ML has the potential to enhance all their components—and is already commonly used for this purpose.

Based on the insights gained in this review, we make the following recommendations:

1. DTs differ considerably between and within domains, yet there are also surprisingly many shared concepts and similarities. To adequately address the needs of the application, ML objectives needs align with the DT's specific objective. This may require the definition of custom losses or architectures, depending on the domain's requirements and restrictions. We explicitly discourage using ML methods 'off-the-shelf' for an operational DT—the default learning objectives, hyperparameters and architectures are often suboptimal, and in our review we frequently observed customised setups performing better.
2. Although arguably most DT applications are industrial, other domains such as Earth sciences and healthcare have different DT concepts and different operational constraints. Traditional key targets such as cost and response times are less relevant when creating a long-term representation of a part of Earth, where interoperability and plausibility are more relevant—this needs to be considered

when selecting learning approaches, since an accurate yet opaque deep learning method may offer a worse trade-off than an explainable lightweight regressor with uncertainty quantification. Moreover, when human subjects are involved, DTs require a different optimisation approach that favours safety over accuracy, which may require asymmetrical loss functions.

3. We observe a trend in contemporary DTs towards deep learning, yet novel technologies involving agentic AI and LLMs remain rare. This is a promising direction for future DTs, which may leverage autonomous agents to enhance interoperability, generate comprehensive explanations, ensure safety and compliance, and to gain a competitive advantage.
4. 'Single-use' DT solutions present an open challenge, which is aggravated by the positivity bias and irreproducibility commonly found in ML-research. To overcome these difficulties, we encourage extensive, open evaluations of ML-driven DTs, which requires secondary experimentation such as sensitivity or ablation studies. Moreover, we encourage the publication of negative or inconclusive results to permit a more objective perspective on the field.
5. Many ML approaches assume that a good training dataset is available. This is not realistic for most DT applications we examined. Data should be modelled as (dependent) data streams from which the ML system should learn and complete inference tasks—ideally in real-time. Labels noise must be considered, as well as inherent uncertainties in supposedly accurate ground-truth data.
6. One of the most promising avenues for ML is to include the PT's properties as a learning objective. Domain knowledge-based approaches, such as physics-informed ML, can help to enrich the validity of the DT and avoid predictions of impossible states. Therefore, such approaches are prime candidates for many DT applications.
7. One of the main benefits of a complete DT system, consisting of a coupled DT and PT, is its capability to make interventions, which can be used to generate realistic training datasets and to permit causal discovery and modelling.
8. The evaluation of ML for DTs is typically multi-dimensional—besides accuracy, other aspects such as robustness, plausibility, explainability, efficiency, ethical/legal requirements, and cost can be relevant in realistic digital twin settings.

To conclude, we note that both DTs and ML are growing and quickly evolving fields, and that many of the current challenges are likely to be overcome in the near future. We envision that future work at the intersection of these fields produces significant interdisciplinary skills and knowledge, giving rise to a new generation of ML-based DT systems which are more realistic, more cost-efficient, amenable to theoretical analyses, explainable to humans, and require less resources to reduce energy consumption.

Notes

1. We omit explicit spatial dependencies $z(t, x)$, as these can be implicitly represented within the state z through spatial discretization.
2. Aiming for maximal generality of our definition, we do not specify the learning-theoretic background. E.g., we do not explicitly assume that the functions can be analyzed within the framework of probably approximately correct (PAC) learning, cf. [176] ; although in many cases of practical interest that may very well be the case.
3. Here and in the remainder of this work, $u|_\tau$ refers to the restriction of the function U to the set τ.
4. Note that, as per (9a), the forecast $\hat{z}(t+\delta)$ is an accurate estimate of the actual PT state $z(t+\delta)$.

Author contributions

CRediT: **Maximilian B. Toller**: Conceptualization, Formal analysis, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing; **Franz M. Rohrhofer**: Formal analysis, Investigation, Methodology, Validation, Writing – original draft, Writing – review & editing; **David Cemernek**: Writing – original draft, Writing – review & editing; **Sebastian Scher**: Visualization, Writing – original draft, Writing – review & editing; **Heimo Gursch**: Writing – original draft, Writing – review & editing; **Lucas Iacono**: Writing – original draft, Writing – review & editing; **Christian Laubichler**: Writing – original draft, Writing – review & editing; **Mark Dokter**: Writing – original draft, Writing – review & editing; **Dominik Kowald**: Funding acquisition, Writing – original draft, Writing – review & editing;

Roman Kern: Funding acquisition, Writing – original draft, Writing – review & editing; **Bernhard C. Geiger**: Conceptualization, Formal analysis, Funding acquisition, Methodology, Project administration, Supervision, Writing – original draft, Writing – review & editing.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

Know Centre Research GmbH and LEC GmbH are COMET Centres within the COMET—Competence Centres for Excellent Technologies Program and are funded by BMK, BMAW, as well as the co-financing provinces Styria, Vienna and Tyrol. COMET is managed by FFG.

ORCID

Maximilian B. Toller http://orcid.org/0000-0002-7206-1991
Franz M. Rohrhofer http://orcid.org/0000-0002-8228-7759
David Cemernek http://orcid.org/0000-0001-5828-6628
Sebastian Scher http://orcid.org/0000-0002-6314-8833
Heimo Gursch http://orcid.org/0000-0001-7465-6553
Lucas Iacono http://orcid.org/0009-0003-0528-3070
Christian Laubichler http://orcid.org/0000-0001-7062-6939
Mark Dokter http://orcid.org/0000-0003-4960-5740
Dominik Kowald http://orcid.org/0000-0003-3230-6234
Roman Kern http://orcid.org/0000-0003-0202-6100
Bernhard C. Geiger http://orcid.org/0000-0003-3257-743X

References

[1] Grand View Research. ; 2023 Digital twin market size, share & trends analysis report by solution (component, process, system), by deployment (cloud, on-premise), by enterprise size, by application, by end-use, by region, and segment forecasts, 2024 - 2030. Technical Report GVR-2-68038-494-9, Grand View Research.
[2] Nagaraj A. Digital twin technology: a review. AI-driven digital twin and industry 4.0; 2024. p. 49–67.
[3] Igbokwe EC, Ezeh FC, Oliha AO. The relevance of digital twin technology in surveying and geoinformatics. Int J Multidiscip Res Growth Evaluation. 2024;5(3):520–526.
[4] Sicard B, McCafferty-Leroux A, Appuhamy R, et al. Aerospace digital twins: examining applications of digital twin technology to unmanned aerial vehicles and satellites. Autonomous Systems: sensors, Processing, and Security for Ground, Air, Sea, and Space Vehicles and Infrastructure 2024. Vol. 13052. SPIE; 2024. p. 57–66.
[5] Balasubhramanyam A, Ramesh R, Sudheer R, et al. Revolutionizing healthcare: a review unveiling the transformative power of digital twins. IEEE Access. 2024:12.
[6] Irfan MS, Dasgupta S, Rahman M. Towards transportation digital twin systems for traffic safety and mobility: a review. IEEE Internet Things J. 2024;11(14):24581–24603. doi: 10.1109/JIOT.2024.3395186
[7] Wang D, Song Y, Zhang Y, et al. Digital twin of optical networks: a review of recent advances and future trends. J Lightwave Technol. 2024a;42(12):4233–4259. Cham, Switzerland. doi: 10.1109/JLT.2024.3401419
[8] El-Agamy RF, Sayed HA, Akhatatneh AMAL, et al. Comprehensive analysis of digital twins in smart cities: a 4200-paper bibliometric study. Artif Intell Rev. 2024;57(6):154. doi: 10.1007/s10462-024-10781-8
[9] da Silva Puppim B, de Souza DGB, Silva GV, et al. Enhancing efficiency in the oil and gas industry through digital twins: a survey on operational and environmental monitoring variables. In: International Joint Conference on Industrial Engineering and Operations Management. Cham, Switzerland: Springer; 2024. p. 447–461. doi: 10.1007/978-3-031-80785-5_34
[10] Kandemir E, Hasan A, Kvamsdal T, et al. Predictive digital twin for wind energy systems: a literature review. Energy Inform. 2024;7(1):68. doi: 10.1186/s42162-024-00373-9
[11] Zhu H, Hwang B-G, Tan YZ, et al. Building on digital twin: overcoming barriers and unlocking success in the construction industry. J Constr Eng Manage. 2024;150(10):04024142. doi: 10.1061/JCEMD4.COENG-13991
[12] Kumari S, Kumar Tyagi A. Internet of things for smart home: a survey. In: Digital twin and blockchain for smart cities; 2024. p. 371–405.

[13] Verma S, Sharma A, Tran B, et al. A systematic review of digital twins for electric vehicles. J Traffic Transp Eng (Engl Ed). 2024;11(5):815–834. doi: 10.1016/j.jtte.2024.04.004
[14] Ruoying LIU, Weiyu ZOU, Shaoqian HU, et al. Survey of digital twin modeling and applications in power system. Comput Mod. 2024;9:61.
[15] Liu J, Yeoh W, Gao S, et al. Blockchain-based digital twin for supply chain management: a survey. In: Digital twins for smart cities and villages; 2025a. p. 581–600.
[16] Shashidhar S, Sunil CK, Vineetha B. 2024 IEEE Conference on Engineering Informatics (ICEI). A Survey of Digital Twins in Smart Farming Systems. IEEE; 2024. p. 1–7. doi: 10.1109/ICEI64305.2024.10912241
[17] Kabir MR, Halder D, Ray S. Digital twins for iot-driven energy systems: a survey. IEEE Access. 2024;12: 177123–177143. doi: 10.1109/ACCESS.2024.3506660
[18] Zami MBA, Shaon S, Quy VK, et al. Digital twin in industries: a comprehensive survey. IEEE Access. 2025;13:47291–47336. doi: 10.1109/ACCESS.2025.3551532
[19] Tobing V, Hamzah M, Chew J, et al. 2024 IEEE International Conference on Computing (ICOCO). Towards Digital Twin Integrated Border Control. IEEE; 2024. p. 114–121. doi: 10.1109/ICOCO62848.2024.10928191
[20] Teofilo A, Sun Q. Towards sustainable urban rooftop solar energy planning through spatial digital twins paradigm: a systematic literature review. In: Digital twin computing for urban intelligence; 2024. p. 15–51. doi: 10.1007/978-981-97-8483-7_2
[21] Galagedarage Don M, Wanasinghe TR, Gosine RG, et al. Digital twins and enabling technology applications in mining: research trends, opportunities, and challenges. IEEE Access. 2025;13:6945–6963. doi: 10.1109/ACCESS.2025.3526881
[22] Nadeem M, Kostic S, Dornhöfer M, et al. A comprehensive review of digital twin in healthcare in the scope of simulative health-monitoring. Digit Health. 2025;11:20552076241304078. doi: 10.1177/20552076241304078
[23] Hasan MA, Mustofa R, Hossain NUI, et al. Smart health practices: strategies to improve healthcare efficiency through digital twin technology. Smart Health. 2025;36:100541.
[24] Wu D, Zheng A, Yu W, et al. Digital twin technology in transportation infrastructure: a comprehensive survey of current applications, challenges, and future directions. Appl Sci. 2025;15(4):1911. doi: 10.3390/app15041911
[25] Alhumam N, Rahman MMH, Aljughaiman A. A comprehensive review on cybersecurity of digital twins issues, challenges, and future research directions. IEEE Access. 2025;13:45106–45124.
[26] Roopa MS, Venugopal KR. Digital twins for cyber-physical healthcare systems: architecture, requirements, systematic analysis and future prospects. IEEE Access. 2025;13:44963–44996. doi: 10.1109/ACCESS.2025.3547991
[27] Ghorbani Z, Messner JI. An ontology for digital twin operations and maintenance use cases in the aeco industry. ITcon. 2025;30:335–351. doi: 10.36680/j.itcon.2025.014
[28] Dang HV, Tatipamula M, Nguyen HX. Cloud-based digital twinning for structural health monitoring using deep learning. IEEE Trans Ind Inf. 2022;18(6):3820–3830. doi: 10.1109/TII.2021.3115119
[29] Karniadakis GE, Kevrekidis IG, Lu L, et al. Physics-informed machine learning. Nat Rev Phys. 2021;3(6): 422–440. doi: 10.1038/s42254-021-00314-5
[30] Sharma A, Kosasih E, Zhang J, et al. Digital twins: state of the art theory and practice, challenges, and open research questions. J Ind Inf Integr. 2022;30:100383. doi: 10.1016/j.jii.2022.100383
[31] Lippi M, Martinelli M, Picone M, et al. Enabling causality learning in smart factories with hierarchical digital twins. Comput Ind. 2023;148:103892. doi: 10.1016/j.compind.2023.103892
[32] Schölkopf B. Causality for machine learning. In: Probabilistic and causal inference: the works of Judea Pearl; 2022. p. 765–804.
[33] Bischl B, Binder M, Lang M, et al. Hyperparameter optimization: foundations, algorithms, best practices, and open challenges. Wiley Interdiscip Rev: data Min Knowl Discov. 2023;13(2):e1484.
[34] Xiao L, Pennington J, Schoenholz S. Disentangling trainability and generalization in deep neural networks. In: International Conference on Machine Learning PMLR; 2020. p. 10462–10472.
[35] D'Amour A, Heller K, Moldovan D, et al. Underspecification presents challenges for credibility in modern machine learning. J Mach Learn Res. 2022;23(226):1–61.
[36] Hoffmann J, Bauer P, Sandu I, et al. Destination earth–a digital twin in support of climate services. Clim Serv. 2023;30:100394. doi: 10.1016/j.cliser.2023.100394
[37] Hazeleger W, Aerts JPM, Bauer P, et al. Digital twins of the earth with and for humans. Commun Earth Environ. 2024;5(1):463. doi: 10.1038/s43247-024-01626-x
[38] Li X, Feng M, Ran Y, et al. Big data in earth system science and progress towards a digital twin. Nat Rev Earth Environ. 2023;4(5):319–332. doi: 10.1038/s43017-023-00409-w
[39] Devadarshini E, Geethalakshmi V, McDermid SP, et al. A systematic review of climate downscaling and extremes in coupled model intercomparison project 6 (cmip6). Environ Dev. 2025;56:101280. doi: 10.1016/j.envdev.2025.101280

[40] Navon IM. Data assimilation for numerical weather prediction: a review. In: Data assimilation for atmospheric, oceanic and hydrologic applications; 2009. p. 21–65.
[41] Hersbach H, Bell B, Berrisford P, et al. The era5 global reanalysis. Quart J Royal Meteoro Soc. 2020; 146(730):1999–2049. doi: 10.1002/qj.3803
[42] Golowich N, Rakhlin A, Shamir O. Size-independent sample complexity of neural networks. In: Conference on Learning Theory. PMLR; 2018. p. 297–299.
[43] Wolpert DH. The lack of a priori distinctions between learning algorithms. Neural Comput. 1996;8(7): 1341–1390. doi: 10.1162/neco.1996.8.7.1341
[44] Steiner AK, Ladstädter F, Randel WJ, et al. Observed temperature changes in the troposphere and stratosphere from 1979 to 2018. J Clim. 2020;33(19):8165–8194. doi: 10.1175/JCLI-D-19-0998.1
[45] Boore DM. Stochastic simulation of high-frequency ground motions based on seismological models of the radiated spectra. Bull Seismol Soc Am. 1983;73(6A):1865–1894.
[46] Wedi N, Sandu I, Bauer P, et al. Implementing digital twin technology of the earth system in destination earth. J Eur Meteorol Soc. 2025;3:100015. doi: 10.1016/j.jemets.2025.100015
[47] Pallikkavaliyaveetil N, Chandrasekaran S. Small data, big challenges: machine-and deep-learning strategies for data-limited drug discovery. Adv Drug Deliv Rev. 2026;229:115762. page doi: 10.1016/j.addr.2025.115762
[48] Brunton SL, Kutz JN. Data-driven science and engineering: machine learning, dynamical systems, and control. Cambridge, United Kingdom: Cambridge University Press; 2022.
[49] Madan T, Sagar S, Virmani D. Air quality prediction using machine learning algorithms–a review. 2020 2nd International Conference on Advances in Computing, Communication Control and Networking (ICACCCN); 2020. IEEE, p. 140–145. doi: 10.1109/ICACCCN51052.2020.9362912
[50] Zucatelli PJ, Nascimento EGS, Santos AÁB, et al. Nowcasting prediction of wind speed using computational intelligence and wavelet in Brazil. Int J Comput Methods Eng Sci Mech. 2020;21(6):343–369. doi: 10.1080/15502287.2020.1841335
[51] Mandelbrot BB, Wallis JR. Noah, Joseph, and operational hydrology. Water Resour Res. 1968;4(5):909–918. doi: 10.1029/WR004i005p00909
[52] Adlouni SE, Bobée B, Ouarda TB. On the tails of extreme event distributions in hydrology. J Hydrol. 2008;355(1-4):16–33. doi: 10.1016/j.jhydrol.2008.02.011
[53] Zhao C, Mei B, Zhang Y, et al. Observations on size distribution characteristics of ice floes in icebreaking channel. In: International Joint Conference on Civil and Marine Engineering. Singapore: Springer; 2024. p. 377–383.
[54] Azad AS, Islam N, Nabi MN, et al. Developments and trends in water level forecasting using machine learning models–a review. IEEE Access. 2025;13:63048–63065. doi: 10.1109/ACCESS.2025.3557910
[55] Mamalakis A, Ebert-Uphoff I, Barnes EA. Explainable artificial intelligence in meteorology and climate science: model fine-tuning, calibrating trust and learning new science. In: International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers; 2020. Cham, Switzerland: Springer. p. 315–339.
[56] Gevaert CM. Explainable ai for earth observation: a review including societal and regulatory perspectives. Int J Appl Earth Obs Geoinf. 2022;112:102869. doi: 10.1016/j.jag.2022.102869
[57] Chilès J-P, Desassis N. Fifty years of kriging. In: Handbook of mathematical geosciences: fifty years of IAMG. Cham, Switzerland: Springer; 2018. p. 589–612.
[58] Christianson RB, Pollyea RM, Gramacy RB. Traditional kriging versus modern gaussian processes for large-scale mining data. Stat Anal. 2023;16(5):488–506. doi: 10.1002/sam.11635
[59] Slingo J, Palmer T. Uncertainty in weather and climate prediction. Philos Trans A Math Phys Eng Sci. 2011;369(1956):4751–4767. doi: 10.1098/rsta.2011.0161.
[60] Palmer T. The ECMWF ensemble prediction system: looking back (more than) 25 years and projecting forward 25 years. Quart J Royal Meteoro Soc. 2019;145(S1):12–24. doi: 10.1002/qj.3383
[61] Scher S, Messori G. Predicting weather forecast uncertainty with machine learning. Q J R Meteorol Soc. 2018;144(717):2830–2841. doi: 10.1002/qj.3410
[62] Brecht R, Bihlo A. Computing the ensemble spread from deterministic weather predictions using conditional generative adversarial networks. Geophys Res Lett. 2023;50(2):e2022GL101452. doi: 10.1029/2022GL101452
[63] Erol T, Mendi AF, Doğan D. The digital twin revolution in healthcare. In: 2020 4th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT); 2020. IEEE. p. 1–7. doi: 10.1109/ISMSIT50672.2020.9255249
[64] Bruynseels K, Santoni de Sio F, Van den Hoven J. Digital twins in health care: ethical implications of an emerging engineering paradigm. Front Genet. 2018;9:31. doi: 10.3389/fgene.2018.00031
[65] Senn S, Rolfe K, Julious SA. Investigating variability in patient response to treatment–a case study from a replicate cross-over study. Stat Methods Med Res. 2011;20(6):657–666. doi: 10.1177/0962280210379174
[66] Laubenbacher R, Mehrad B, Shmulevich I, et al. Digital twins in medicine. Nat Comput Sci. 2024;4(3): 184–191. doi: 10.1038/s43588-024-00607-6

[67] Loewe A, Díaz PM, Nagel C, et al. Cardiac digital twin modeling. In: Innovative treatment strategies for clinical electrophysiology. Singapore: Springer; 2022. p. 111–134. doi: 10.1007/978-981-19-6649-1_7
[68] Mollica L, Leli C, Sottotetti F, et al. Digital twins: a new paradigm in oncology in the era of big data. ESMO Real World Data Digit Oncol. 2024;5:100056. doi: 10.1016/j.esmorw.2024.100056
[69] Konietschke F, Schwab K, Pauly M. Small sample sizes: a big data problem in high-dimensional data analysis. Stat Methods Med Res. 2021;30(3):687–701. doi: 10.1177/0962280220970228
[70] Vallée A. Digital twin for healthcare systems. Front Digit Health. 2023;5:1253050. doi: 10.3389/fdgth.2023.1253050
[71] Chicco D, Jurman G. Machine learning can predict survival of patients with heart failure from serum creatinine and ejection fraction alone. BMC Med Inform Decis Mak. 2020;20(1):16. doi: 10.1186/s12911-020-1023-5
[72] Mattsson N, Palmqvist S, Stomrud E, et al. Staging β-amyloid pathology with amyloid positron emission tomography. JAMA Neurol. 2019;76(11):1319–1329. doi: 10.1001/jamaneurol.2019.2214
[73] Subramanian S, Harrington P, Keutzer K, et al. Towards foundation models for scientific machine learning: characterizing scaling and transfer behavior. Adv Neural Inf Process Syst. 2023;36:71242–71262.
[74] Bickley L, Szilagyi PG. Bates' guide to physical examination and history-taking. Philadelphia, PA, USA: Lippincott Williams & Wilkins; 2012.
[75] Jindal H, Agrawal S, Khera R, et al. Heart disease prediction using machine learning algorithms. IOP Conf Ser: Mater Sci Eng. 2021;1022: 012072. doi: 10.1088/1757-899X/1022/1/012072
[76] Long X, Chen L, Jiang C, and Alzheimer's Disease Neuroimaging Initiative, et al. Prediction and classification of Alzheimer disease based on quantification of mri deformation. PLoS One. 2017;12(3):e0173372. doi: 10.1371/journal.pone.0173372
[77] Johnsen SH, Lilleng H, Wilsgaard T, et al. Creatine kinase activity and blood pressure in a normal population: the Tromsø study. J Hypertens. 2011;29(1):36–42. doi: 10.1097/HJH.0b013e32834068e0
[78] Villalobos P, Wistuba II. Lung cancer biomarkers. Hematol Oncol Clin North Am. 2017;31(1):13–29. doi: 10.1016/j.hoc.2016.08.006
[79] Iragorri N, Spackman E. Assessing the value of screening tools: reviewing the challenges and opportunities of cost-effectiveness analysis. Public Health Rev. 2018;39(1):17. doi: 10.1186/s40985-018-0093-8
[80] Zhou X, Liu X, Jiang J, et al. Asymmetric loss functions for learning with noisy labels. In: International Conference on Machine Learning. PMLR; 2021. p. 12846–12856.
[81] Li W, Rentemeister M, Badeda J, et al. Digital twin for battery systems: cloud battery management system with online state-of-charge and state-of-health estimation. J Storage Mater. 2020;30:101557. doi: 10.1016/j.est.2020.101557
[82] Tsakalos D, Krimpas GA, Mourtzis D. Digital twin for early-stage fault detection in a hybrid manufacturing cell. Digital Twin; 2026. p .2622705.
[83] Yi L, Ruediger-Flore P, Karnoub A, et al. Is it possible to develop a digital twin for noise monitoring in manufacturing? Digital Twin. 2025;2(2):2520118. doi: 10.12688/digitaltwin.17931.2
[84] Khajavi SH, Motlagh NH, Jaribion A, et al. Digital twin: vision, benefits, boundaries, and creation for buildings. IEEE Access. 2019;7:147406–147419. doi: 10.1109/ACCESS.2019.2946515
[85] Javaid M, Haleem A, Suman R. Digital twin applications toward industry 4.0: a review. Cogn Robot. 2023;3:71–92. doi: 10.1016/j.cogr.2023.04.003
[86] Jiang Y, Yin S, Li K, et al. Industrial applications of digital twins. Philos Trans A Math Phys Eng Sci. 2021;379(2207):20200360. doi: 10.1098/rsta.2020.0360
[87] Song Z, Hackl CM, Anand A, et al. Digital twins for the future power system: an overview and a future perspective. Sustainability. 2023;15(6):5259. doi: 10.3390/su15065259
[88] Lasi H, Fettke P, Kemper H-G, et al. Industry 4.0. Bus Inf Syst Eng. 2014;6(4):239–242. doi: 10.1007/s12599-014-0334-4
[89] Attaran M, Celik BG. Digital twin: benefits, use cases, challenges, and opportunities. Decis Anal J. 2023;6:100165. doi: 10.1016/j.dajour.2023.100165
[90] Saporiti N, Cannas VG, Pozzi R, et al. Challenges and countermeasures for digital twin implementation in manufacturing plants: a delphi study. Int J Prod Econ. 2023;261:108888. doi: 10.1016/j.ijpe.2023.108888
[91] Deepu TS, Ravi V. Exploring critical success factors influencing adoption of digital twin and physical internet in electronics industry using grey-dematel approach. Digital Business. 2021;1(2):100009. doi: 10.1016/j.digbus.2021.100009
[92] Kober C, Medina FG, Benfer M, et al. Digital twin stakeholder communication: characteristics, challenges, and best practices. Comput Ind. 2024;161:104135. doi: 10.1016/j.compind.2024.104135
[93] VanDerHorn E, Mahadevan S. Digital twin: generalization, characterization and implementation. Decis Support Syst. 2021;145:113524. doi: 10.1016/j.dss.2021.113524
[94] Kober C, Polzer J, Kulasingham E, et al. Success criteria for value-oriented digital twin development and implementation: insights from an industry case study. Int J Prod Res. 2025:1–25. doi: 10.1080/00207543.2025.2514728

[95] Wieser R, Fan Y, Yu X, et al. Large scale, data driven, digital twin models: outlier detection and imputation. In: 2024 IEEE 52nd Photovoltaic Specialist Conference (PVSC); 2024. IEEE. p. 0902–0905. doi: 10.1109/PVSC57443.2024.10748985
[96] Spaney P, Becker S, Ströbel R, et al. A model-driven digital twin for manufacturing process adaptation. In: 2023 ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C). IEEE; 2023. p. 465–469. doi: 10.1109/MODELS-C59198.2023.00081
[97] Mattila J, Ala-Laurinaho R, Autiosalo J, et al. Interoperability of digital twins for automation with digital twin schema. IEEE Access. 2025;13:200595–200608. doi: 10.1109/ACCESS.2025.3633736
[98] Maharana K, Mondal S, Nemade B. A review: data pre-processing and data augmentation techniques. Glob Transit Proc. 2022;3(1):91–99. doi: 10.1016/j.gltp.2022.04.020
[99] Rodríguez P, Bautista MA, Gonzàlez J, et al. Beyond one-hot encoding: lower dimensional target embedding. Image Vis Comput. 2018;75:21–31. doi: 10.1016/j.imavis.2018.04.004
[100] Welch G, Bishop G, et al. An introduction to the kalman filter; 1995.
[101] Igba J, Alemzadeh K, Durugbo C, et al. Analysing rms and peak values of vibration signals for condition monitoring of wind turbine gearboxes. Renewable Energy. 2016;91:90–106. doi: 10.1016/j.renene.2016.01.006
[102] Guerraoui R, Gupta N, Pinot R. Robust machine learning. Vol. 1. Singapore: Springer; 2024.
[103] Wang L, Zhang X, Su H, et al. A comprehensive survey of continual learning: theory, method and application. IEEE Trans Pattern Anal Mach Intell. 2024b;46(8):5362–5383. doi: 10.1109/TPAMI.2024.3367329
[104] Ghosh A, Kumar H, Sastry PS. Robust loss functions under label noise for deep neural networks. AAAI. 2017;31(1):1919–1925. doi: 10.1609/aaai.v31i1.10894
[105] Berthon A, Han B, Niu G, et al. Confidence scores make instance-dependent label-noise learning possible. In: International Conference on Machine Learning. PMLR; 2021. p. 825–836.
[106] Nishihara R, Moritz P, Wang S, et al. Real-time machine learning: the missing pieces. In: Proceedings of the 16th Workshop on Hot Topics in Operating Systems; 2017. p. 106–110.
[107] Gomes HM, Read J, Bifet A, et al. Machine learning for streaming data: state of the art, challenges, and opportunities. SIGKDD Explor Newsl. 2019;21(2):6–22. doi: 10.1145/3373464.3373470
[108] Yin C, Hu Z, Hua L, et al. Stress state perception and prediction of isothermal forging hydraulic press driven by digital twin and machine learning. Expert Syst Appl. 2026;331:133209.
[109] Han F, Liu L, Sun J. Logistics equipment condition monitoring and prediction based on digital twin and machine learning: f. han et al. Sci Rep. 2026;16(1):1–42. doi: 10.1038/s41598-026-43380-4
[110] Qin R, Zi X, Ge X. ; 2026. Digital twin-based intelligent risk assessment and decision support system for university student entrepreneurial projects. Sci. Rep.
[111] Briggs JK, Stroh JN, Park S, et al. Towards a physics informed digital twin to predict cerebral blood flow and cerebral vascular regulation. npj Digit Med. 2026.
[112] Wu Z. Personalized skill transfer optimization in swimming training through multi-agent reinforcement learning driven digital twin environments. Sci .Rep. 2026.
[113] Xin P, Isleem HF, Khishe M. A digital twin approach for sustainable construction: predictive optimization of concrete strength using industry 4.0 principles. Sci Rep. 2025;16(1):2443. doi: 10.1038/s41598-025-32276-4
[114] Gupta H, Kundu P. Optimising digital twin parameters to synchronise virtual-physical system response for condition monitoring solutions development. Digit Eng. 2026:100125.
[115] Kanti PK, Kumar HGP, Elsaeedy HI, et al. Digital twin analysis of graphene–silicon solar schottky-heterojunction cell efficiency for terrestrial photovoltaics. Sci Rep. 2026.
[116] Ahanger TA, Abdibayev Z, Sagnayeva S, et al. Smart wastewater management in hydro-technical systems using digital twin technology. Sci Rep. 2026.
[117] Vaidya S, Jethava G, Jethava S. Enhancing smart factory performance via hybrid scheduling and intelligent resource management. Sci Rep. 2026. doi: 10.1038/s41598-026-49107-9
[118] Abdel-Mooty MN, Coulibaly P, El-Dakhakhni W. A scalable data driven geospatial framework for climate risk assessment. Sci Rep. 2025.
[119] Lu H, Liu L, Gao Z, et al. Spatiotemporal collaborative digital twin structural health monitoring based on dynamic multi-fidelity models. Comput Ind. 2026;180:104502. doi: 10.1016/j.compind.2026.104502
[120] Dominguez-Gomez P, Zingaro A, Balzotti C, et al. Ai-enhanced cardiac digital twins extend drug proarrhythmic risk assessment through experimental data uncertainty propagation and overdose exploration: a loperamide case study. Regul Toxicol Pharmacol. 2026;171:106138. doi: 10.1016/j.yrtph.2026.106138
[121] Boyu WANG, Soto BGd, Qian WANG. Human-inspired zero-shot semantic segmentation of construction point clouds based on adaptive observation and mllm reasoning. Adv Eng Inf. 2026;76:104953. doi: 10.1016/j.aei.2026.104953
[122] Cuzzocrea A, Benlaredj I. Multidimensional risk analysis and prediction over digital twins via big data paradigms. Inf Sci. 2026;753:123582. doi: 10.1016/j.ins.2026.123582
[123] Zhang L, Zhang X, Chen K, et al. Digital twin-enabled synchronous jacking control of an aerial building machine using a scale model. Autom Constr. 2026a;190:107076. doi: 10.1016/j.autcon.2026.107076

[124] Bianchi G, Freddi F, Giuliani F. Ai–digital twin framework for automated predictive maintenance of bonded insulated rail joints. Autom Constr. 2026;190:107102. doi: 10.1016/j.autcon.2026.107102
[125] Cumbajin E, Rodrigues N, Frazão L, et al. Digital twin architecture for ceramics 5.0. Future Gener Comput Syst. 2026;183:108529. doi: 10.1016/j.future.2026.108529
[126] Huang M, Huang S, Gao L, et al. Digital twin-driven multi-agent collaborative online optimization of production regulation for smart reconfigurable manufacturing systems with human-robot collaboration. Rob Comput Integr Manuf. 2026;101:103322. doi: 10.1016/j.rcim.2026.103322
[127] Oussmou B, Saaidi A, Abderafi S. Data-driven model to control the alkaline electrolyzer efficiency, for hydrogen production. Comput Chem Eng. 2026;213:109745. doi: 10.1016/j.compchemeng.2026.109745
[128] Romano G, Bilancia P, Locatelli A, et al. A machine learning–based tool for enhancing position accuracy in industrial robots with a reduced dataset. Rob Comput Integr Manuf. 2026;101:103289. doi: 10.1016/j.rcim.2026.103289
[129] Meng X, Liu X, Zhang H, et al. A physics-informed neural network framework for wave-induced dynamic response in a poroelastic seabed. Comput Geotech. 2026;198:108250. doi: 10.1016/j.compgeo.2026.108250
[130] Wan A, He J, Li P, et al. Digital twin-oriented decision support for central air conditioning systems: integrated load prediction and fault diagnosis via hybrid machine learning. Energy. 2026;360:141697. doi: 10.1016/j.energy.2026.141697
[131] Keçeci M, Oktal H. Physics-informed hybrid digital twin for electric aircraft battery state of function: a probabilistic mission success approach. J Storage Mater. 2026;172:123020. doi: 10.1016/j.est.2026.123020
[132] Paldino GM, Caelen O, Oueslati M, et al. On the integration of domain adaptation and causal discovery in digital twins: a case study for plastic injection molding. Future Gener Comput Syst. 2026;182:108427.
[133] Mohamed MC, Ambrosone R, Masood MU, et al. Reinforcement learning to enhance digital twin based edfa fault prediction in optical networks. IEEE Photonics Technology Letters; 2026.
[134] Jiang Z, Moghaddas SA, Bao Y. Engineering mechanics-aware digital twin for pipeline condition assessment. Autom Constr. 2026;189:107070. doi: 10.1016/j.autcon.2026.107070
[135] Tabassum S, Vijay Babu AR, Dheer DK, et al. A digital twin framework driven energy management strategy for hybrid dc microgrids with integrated xev charging applications. Sustain Comput: Inf Syst. 2026;51:101364. doi: 10.1016/j.suscom.2026.101364
[136] Chen J, Liang Y, Lu K, et al. Digital twin-based energy-efficient framework for smart grid management using RF algorithm. Sustain Comput: Inf Syst. 2026;51:101387.
[137] Zhu J, Peng C, Lai H, et al. A novel aggregation modeling method for lithium-ion batteries based on artificial intelligence and digital twin. Energy. 2026;358:141345. doi: 10.1016/j.energy.2026.141345
[138] Yan J, Ding H. A hybrid rigid-compliant computational framework for large-scale assembly: integrating torsor kinematics with physics-informed graph learning. Eng Struct. 2026;362:122965. doi: 10.1016/j.engstruct.2026.122965
[139] Zhang Y, Tong Z-X, Ni S, et al. A two-stage physics-informed graph neural network approach for rapid and accurate turbine blade temperature field prediction and data fusion. Aerosp Sci Technol. 2026b;176:112713. doi: 10.1016/j.ast.2026.112713
[140] Al Muttaki MR, Afrin S, Anil AIA, et al. Demand-based anomaly detection algorithm in electric vehicle charging systems using machine learning and digital twin simulation. J Storage Mater. 2026;170:122795. doi: 10.1016/j.est.2026.122795
[141] Abadi BNR, Hattel JH, Sarhadi A. Machine learning-driven digital twin for real-time temperature prediction and adaptive process control in large-scale additive manufacturing. J Manuf Syst. 2026;87:62–77. doi: 10.1016/j.jmsy.2026.05.003
[142] Wang Z, Zhou G, Zhang C, et al. A tool wear prediction model integrating mechanism-data fusion with transfer learning for machine tool digital twins. Rob Comput Integr Manuf. 2026a;100:103235. doi: 10.1016/j.rcim.2026.103235
[143] Cavus M, Jiang J, Allahham A, et al. A physics-informed machine learning framework for climate-aware digital twins in decentralised energy systems. Appl Energy. 2026;416:128013. doi: 10.1016/j.apenergy.2026.128013
[144] Gupta V, Mangal D. A deep learning approach based small object detection for dynamic digital twins in smart manufacturing. Comput Electr Eng. 2026;136:111246. doi: 10.1016/j.compeleceng.2026.111246
[145] Cheikhrouhou S, Turki M, Dammak B. X-fedagritwin: privacy-preserving crop recommendation via digital twins and explainable federated learning. In: Information and Software Technology; 2026. p. 108173.
[146] Mavridis C, Barbosa FS, Farhadi H, et al. Learning a network digital twin as a hybrid system. Nonlinear Anal Hybrid Syst. 2026;61:101746. doi: 10.1016/j.nahs.2026.101746
[147] Guzmán E, Andrés B, Torres-Polo M. A cooperative digital twin–multi-agent reinforcement learning for circular supply chains: balanced control across production, logistics, and sustainability. Comput Ind Eng. 2026;218:112044. doi: 10.1016/j.cie.2026.112044
[148] Zeng X, Yang C, Xue J. Energy consumption optimization control model for highway electromechanical systems integrating digital twins and reinforcement learning. Eng Appl Artif Intell. 2026;177:114882. doi: 10.1016/j.engappai.2026.114882

[149] Shao C, Qin X, Li X, et al. Lightweight graph retrieval-augmented generation agent for safety monitoring digital twins in small-and medium-scale hydropower projects. Autom Constr. 2026;188:107041. doi: 10.1016/j.autcon.2026.107041
[150] Cai W, Ma C, He J, et al. Physics-guided thermoelastic state-space learning for spindle thermal error in machine tools. Expert Syst Appl. 2026;322:132444. doi: 10.1016/j.eswa.2026.132444
[151] Qiu Y, Liu C, Zeng Y, et al. Physics-guided, uncertainty-aware machine learning for non-invasive deep thermal interface characterization in hybrid-bonded 3d ics without prior geometric knowledge. Int J Heat Mass Transf. 2026;263:128717. doi: 10.1016/j.ijheatmasstransfer.2026.128717
[152] Olca KD, Yücel Ö. Deep learning applications for biomass gasifier monitoring: a dual approach with direct and recursive long short-term memory forecasting. Renewable Energy. 2026;268:125779.
[153] Wan Y, Xu Q. Machine learning-aided digital twin framework with noninvasive parameter estimation for voltage source converters. IEEE Trans Ind Electron. 2026;73(7):10918–10930. doi: 10.1109/TIE.2026.3657019
[154] Wang Z, Man Z, Meng L, et al. Ai-driven digital twin and delay-aware surrogate mpc framework for biogas production. Comput Chem Eng. 2026b;210:109637. doi: 10.1016/j.compchemeng.2026.109637
[155] Thosar D, Bhakte A, Li Z, et al. Adaptive physics-informed neural network-based digital twins integrated with ensemble kalman filter. Comput Chem Eng. 2026;210:109625.
[156] Feng X, Jiang Y, Huang H, et al. Fast prediction of oceanic nearshore circulation (FPOCH) system with multi-model fusion strategies. Appl Ocean Res. 2026;170:105059. doi: 10.1016/j.apor.2026.105059
[157] Lee D, Kang BK, Gyeom An T, et al. A field scale case study of digital twin framework for precision tomato cultivation: integrating xgboost-based prediction with model predictive control. Sci Hortic. 2026;362:114877. doi: 10.1016/j.scienta.2026.114877
[158] Tao F, Qi Q, Wang L, et al. Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: correlation and comparison. Engineering. 2019;5(4):653–661. doi: 10.1016/j.eng.2019.01.014
[159] Alpaydin E. Introduction to machine learning. Cambridge, MA, USA: MIT Press; 2020.
[160] Borisov V, Leemann T, Sebler K, et al. Deep neural networks and tabular data: a survey. IEEE Trans Neural Netw Learn Syst. 2024;35(6):7499–7519. doi: 10.1109/TNNLS.2022.3229161
[161] European Commission. European Virtual Human Twins (VHT) initiative. 2025 [accessed 2026 Jul 2]. Available at: https://digital-strategy.ec.europa.eu/en/policies/virtual-human-twins.
[162] United Nations. Department of Economic and Social Affairs, Population Division. World population prospects 2024: Summary of results. 2024. URL UN DESA Policy Brief UN DESA/POP/2024/TR/NO. 1, United Nationshttps://un.org.
[163] Meng C, Griesemer S, Cao D, et al. When physics meets machine learning: a survey of physics-informed machine learning. Mach Learn Comput Sci Eng. 2025;1(1):20. doi: 10.1007/s44379-025-00016-0
[164] Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. Adv Neural Inf Process Syst. 2017;30:5998–6008.
[165] Zhang G, Huang H, Zhou R, et al. High-risk trajectories generation for safety testing of autonomous vehicles based on in-depth crash data. IEEE Trans Intell Transport Syst. 2025;26(8):11619–11630. doi: 10.1109/TITS.2025.3576236
[166] Liu X, Huang H, Bian J, et al. Generating intersection pre-crash trajectories for autonomous driving safety testing using transformer time-series generative adversarial networks. Eng Appl Artif Intell. 2025b;160:111995. doi: 10.1016/j.engappai.2025.111995
[167] Bracken A, Reilly C, Feeley A, et al. Artificial intelligence (AI)–powered documentation systems in healthcare: a systematic review. J Med Syst. 2025;49(1):28. doi: 10.1007/s10916-025-02157-4
[168] Palazhchenko Y, Shendryk V, Shendryk S. Digital twins data visualization methods. Problems of human interaction: a review. In: International Conference 'New Technologies, Development and Applications. Springer; 2023. p. 478–485.
[169] Makridakis S, Spiliotis E, Hollyman R, et al. The m6 forecasting competition: bridging the gap between forecasting and investment decisions. Int J Forecast. 2025;41(4):1315–1354. doi: 10.1016/j.ijforecast.2024.11.002
[170] Su L, Zuo X, Li R, et al. A systematic review for transformer-based long-term series forecasting. Artif Intell Rev. 2025;58(3):80. doi: 10.1007/s10462-024-11044-2
[171] Hribernik K, Cabri G, Mandreoli F, et al. Autonomous, context-aware, adaptive digital twins—state of the art and roadmap. Comput Ind. 2021;133:103508. doi: 10.1016/j.compind.2021.103508
[172] Xu L, Mak S, Schoepf S, et al. Multi-agent digital twinning for collaborative logistics: framework and implementation. J Ind Inf Integr. 2025;45:100799. doi: 10.1016/j.jii.2025.100799
[173] Kowald D, Scher S, Pammer-Schindler V, et al. Establishing and evaluating trustworthy ai: overview and research challenges. Front Big Data. 2024;7:1467222. doi: 10.3389/fdata.2024.1467222
[174] ISO. Iso/iec 30173:2023 digital twin—concepts and terminology. November. Available at: 2023https://www.iso.org/standard/81442.html.
[175] Antil H. Mathematical opportunities in digital twins. In: Proceedings of the Workshop on Mathematical Opportunities in Digital Twins; 2023. p. 1–51.

[176] Shalev-Shwartz S, Ben-David S. Understanding machine learning: from theory to algorithms. New York, NY, USA: Cambridge University Press; 2014.
[177] Zhang R, Wang F, Cai J, et al. Digital twin and its applications: a survey. Int J Adv Manuf Technol. 2022; 123(11-12):4123–4136. doi: 10.1007/s00170-022-10445-3
[178] Liu G-P. Control strategies for digital twin systems. IEEE/CAA J Autom Sinica. 2024;11(1):170–180. doi: 10.1109/JAS.2023.123834
[179] Cappelle C, Cattebeke M, Bosmans J, et al. Sensor selection for cost-effective virtual torque measurements on a wind turbine gearbox. Forsch Ingenieurwes. 2021;85(2):325–334. doi: 10.1007/s10010-021-00464-z
[180] Liu L, Kuo SM, Zhou M. Virtual sensing techniques and their applications. In: 2009 International Conference on Networking, Sensing and Control. IEEE; 2009. p. 31–36.
[181] Lu Q, Parlikad AK, Woodall P, et al. Developing a digital twin at building and city levels: Case study of west cambridge campus. J Manage Eng. 2020;36(3):05020004. doi: 10.1061/(ASCE)ME.1943-5479.0000763
[182] Lv Q, Zhang R, Sun X, et al. A digital twin-driven human-robot collaborative assembly approach in the wake of covid-19. J Manuf Syst. 2021;60:837–851. doi: 10.1016/j.jmsy.2021.02.011
[183] Brackel H-U, Macpherson J, Mieting R, et al. An open approach to drilling systems automation. In: SPE Asia Pacific Oil and Gas Conference and Exhibition. SPE; 2018. p. D021S015R001. doi: 10.2118/191939-MS
[184] Coviello G, Avitabile G, Florio A. The importance of data synchronization in multiboard acquisition systems. In: 2020 IEEE 20th Mediterranean Electrotechnical Conference (MELECON). IEEE; 2020. p. 293–297.
[185] Tan B, Matta A. The digital twin synchronization problem: framework, formulations, and analysis. IISE Trans. 2024;56(6):652–665. doi: 10.1080/24725854.2023.2253869
[186] Picado J, Davis J, Termehchy A, et al. Learning over dirty data without cleaning. In: Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data; 2020. p. 1301–1316.
[187] Alsheikh MA, Lin S, Niyato D, et al. Machine learning in wireless sensor networks: algorithms, strategies, and applications. IEEE Commun Surv Tutorials. 2014;16(4):1996–2018. doi: 10.1109/COMST.2014.2320099
[188] Kong T, Hu T, Zhou T, et al. Data construction method for the applications of workshop digital twin system. J Manuf Syst. 2021;58:323–328. doi: 10.1016/j.jmsy.2020.02.003
[189] Saroj AJ, Roy S, Guin A, et al. Impact of connected corridor volume data imputations on digital twin performance measures. Int J Transp Sci Technol. 2023;12(2):476–491. doi: 10.1016/j.ijtst.2022.05.001
[190] Hua M, Pei J. Cleaning disguised missing data: a heuristic approach. In: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining; 2007. p. 950–958.
[191] Emmanuel T, Maupong T, Mpoeleng D, et al. A survey on missing data in machine learning. J Big Data. 2021;8(1):140. doi: 10.1186/s40537-021-00516-9
[192] Yu Q, Miche Y, Eirola E, et al. Regularized extreme learning machine for regression with missing data. Neurocomputing (Amst). 2013;102:45–51. doi: 10.1016/j.neucom.2012.02.040
[193] Soen A, Sun K. On the variance of the fisher information for deep learning. Adv Neural Inf Process Syst. 2021;34:5708–5719.
[194] Liu J, Xu Z, Xiong H, et al. Digital twin-driven robotic disassembly sequence dynamic planning under uncertain missing condition. IEEE Trans Ind Inf. 2023;19(12):11846–11855. doi: 10.1109/TII.2023.3253187
[195] Zhou B, Cai H, Jiang H, et al. Processing method for missing data in digital twin system. In: International Workshop of Advanced Manufacturing and Automation. Springer; 2023. p. 135–143.
[196] Sun Y, Li J, Xu Y, et al. Deep learning versus conventional methods for missing data imputation: a review and comparative study. Expert Syst Appl. 2023;227:120201. doi: 10.1016/j.eswa.2023.120201
[197] Park S, Singh K, Nellikkattil A, et al. Downscaling earth system models with deep learning. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining; 2022. p. 3733–3742.
[198] Chen RTQ, Rubanova Y, Bettencourt J, et al. Neural ordinary differential equations. Adv Neural Inf Process Syst. 2018:31.
[199] Kidger P, Morrill J, Foster J, et al. Neural controlled differential equations for irregular time series. Adv Neural Inf Process Syst. 2020;33:6696–6707.
[200] Huber PJ, Ronchetti EM. Robust statistics. Hoboken, NJ, USA: John Wiley & Sons; 2011.
[201] Chandola V, Banerjee A, Kumar V. Anomaly detection: a survey. ACM Comput Surv. 2009;41(3):1–58. doi: 10.1145/1541880.1541882
[202] Liu Q, Paparrizos J. The elephant in the room: towards a reliable time-series anomaly detection benchmark. Adv Neural Inf Process Syst. 2024;37:108231–108261.
[203] Saquib Sarfraz M, Chen M-Y, Layer L, et al. Position: Quo Vadis, unsupervised time series anomaly detection? In: International Conference on Machine Learning. PMLR; 2024. p. 43461–43476.
[204] Zhang X-Y, Liu C-L, Suen CY. Towards robust pattern recognition: a review. Proc IEEE. 2020;108(6):894–922. doi: 10.1109/JPROC.2020.2989782

[205] Leino K, Wang Z, Fredrikson M. Globally-robust neural networks. In: International Conference on Machine Learning PMLR; 2021. p. 6212–6222.
[206] Zhou Y, Ge Y, Jia L. Double robust federated digital twin modeling in smart grid. IEEE Internet Things J. 2024;11(24):39913–39931. doi: 10.1109/JIOT.2024.3450433
[207] Correia JB, Abel M, Becker K. Data management in digital twins: a systematic literature review. Knowl Inf Syst. 2023;65(8):3165–3196. doi: 10.1007/s10115-023-01870-1
[208] Fennell C. A communication architecture for transportation digital twins using Apache Kafka. Ireland: Trinity College Dublin; 2022.
[209] Jedermann R, Singh K, Lang W, et al. Digital twin concepts for linking live sensor data with real-time models. J Sens Sens Syst. 2023;12(1):111–121. doi: 10.5194/jsss-12-111-2023
[210] Sadegh Es-Haghi M, Anitescu C, Rabczuk T. Methods for enabling real-time analysis in digital twins: a literature review. Comput Struct. 2024;297:107342. doi: 10.1016/j.compstruc.2024.107342
[211] Mchirgui N, Quadar N, Kraiem H, et al. The applications and challenges of digital twin technology in smart grids: a comprehensive review. Appl Sci. 2024;14(23):10933. doi: 10.3390/app142310933
[212] Yan J, Lu Q, Li N, et al. Common data environment for digital twins from building to city levels. Autom Constr. 2025;174:106131. doi: 10.1016/j.autcon.2025.106131
[213] de Conto E, Genest B, Easwaran A. Function + data flow: a framework to specify machine learning pipelines for digital twinning. In: Proceedings of the 1st ACM International Conference on AI-Powered Software; 2024. p. 19–27.
[214] Merino J, Xie X, Moretti N, et al. Data integration for digital twins in the built environment based on federated data models. Proc Inst Civ Eng -Smart Infrastruct Constr. 2023;176(4):194–211. doi: 10.1680/jsmic.23.00002
[215] dos Santos CH, Lima RDC, Leal F, et al. A decision support tool for operational planning: a digital twin using simulation and forecasting methods. Production. 2020;30:e20200018.
[216] D'Urso D, Chiacchio F, Cavalieri S, et al. Predictive maintenance of standalone steel industrial components powered by a dynamic reliability digital twin model with artificial intelligence. Reliab Eng Syst Saf. 2024;243:109859.
[217] Biswas BN, Chatterjee S, Mukherjee SP, et al. A discussion on Euler method: a review. Electron J Math Anal Appl. 2013;1(2):2090–2792.
[218] Hyndman RJ, Athanasopoulos G. Forecasting: principles and practice. Melbourne, Australia: OTexts; 2018.
[219] Natekin A, Knoll A. Gradient boosting machines, a tutorial. Front Neurorobot. 2013;7:21. doi: 10.3389/fnbot.2013.00021
[220] Salinas D, Flunkert V, Gasthaus J, et al. Deepar: probabilistic forecasting with autoregressive recurrent networks. Int J Forecast. 2020;36(3):1181–1191. doi: 10.1016/j.ijforecast.2019.07.001
[221] Rohrhofer FM, Posch S, Gößnitzer C, et al. Data vs. physics: the apparent Pareto front of physics-informed neural networks. IEEE Access. 2023a;11:86252–86261. doi: 10.1109/ACCESS.2023.3302892
[222] Karpatne A, Kannan R, Kumar V, editors. Knowledge guided machine learning. accelerating discovery using scientific knowledge and data. Boca Raton, FL, USA: CRC Press; 2023.
[223] Hoffer JG, Ofner AB, Rohrhofer FM, et al. Theory-inspired machine learning – towards a synergy between knowledge and data. Weld World. 2022;66:1291–1304. doi: 10.1007/s40194-022-01270-z
[224] Arfken GB, Weber HJ, Harris FE. Mathematical methods for physicists: a comprehensive guide. Waltham, MA, USA: Academic Press; 2011.
[225] Khotanzad A, Lu J-H. Classification of invariant image representations using a neural network. IEEE Trans Acoust, Speech, Signal Process. 1990;38(6):1028–1038. doi: 10.1109/29.56063
[226] Lucia DJ, Beran PS, Silva WA. Reduced-order modeling: new approaches for computational physics. Prog Aerosp Sci. 2004;40(1-2):51–117. doi: 10.1016/j.paerosci.2003.12.001
[227] Li X, Deng Z, Feng R, et al. Deep learning-based reduced order model for three-dimensional unsteady flow using mesh transformation and stitching. Comput Fluids. 2024;285:106441. doi: 10.1016/j.compfluid.2024.106441
[228] Levine S, Finn C, Darrell T, et al. End-to-end training of deep visuomotor policies. J Mach Learn Res. 2016;17(39):1–40.
[229] Buckingham E. On physically similar systems; illustrations of the use of dimensional equations. Phys Rev. 1914;4(4):345–376. doi: 10.1103/PhysRev.4.345
[230] Rohrhofer FM, Posch S, Gößnitzer C, et al. Utilizing neural networks to supplant chemical kinetics tabulation through mass conservation and weighting of species depletion. Energy AI. 2024;16:100341. doi: 10.1016/j.egyai.2024.100341
[231] Cranmer M, Greydanus S, Hoyer S, et al. Lagrangian neural networks. arXiv Preprint arXiv:2003.04630. 2020.
[232] Greydanus S, Dzamba M, Yosinski J. Hamiltonian neural networks. Adv Neural Inf Process Syst. 2019;32.
[233] Satorras VG, Hoogeboom E, Welling M. E(n) equivariant graph neural networks. In Meila M, Zhang T, editors Proc. Int. Conf. on Machine Learning (ICML), volume 139 of Proceedings of Machine Learning Research, 18–24 Jul. PMLR; 2021. p. 9323–9332. Cambridge, MA, USA.

[234] Cohen T, Welling M. Group equivariant convolutional networks. In: International Conference on Machine Learning. PMLR; 2016. p. 2990–2999.
[235] Schütt KT, Sauceda HE, Kindermans P-J, et al. Schnet–a deep learning architecture for molecules and materials. J Chem Phys. 2018;148(24):241722. doi: 10.1063/1.5019779
[236] Thomas N, Smidt T, Kearnes S, et al. Tensor field networks: rotation-and translation-equivariant neural networks for 3d point clouds. arXiv preprint arXiv:1802.08219. 2018.
[237] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Comput Phys. 2019;378:686–707. doi: 10.1016/j.jcp.2018.10.045
[238] Jagtap AD, Mao Z, Adams N, et al. Physics-informed neural networks for inverse problems in supersonic flows. Comput Phys. 2022;466:111402. doi: 10.1016/j.jcp.2022.111402
[239] Chen Y, Lu L, Karniadakis GE, et al. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. Opt Express. 2020;28(8):11618–11633. doi: 10.1364/OE.384875
[240] Rohrhofer FM, Posch S, Gößnitzer C, et al. On the role of fixed points of dynamical systems in training physics-informed neural networks. 2023 TransMachLearnRes, 2023b. https://openreview.net/forum?id=56cTmVrg5w.
[241] Daw A, Bu J, Wang S, et al. Mitigating propagation failures in physics-informed neural networks using retain-resample-release (R3) sampling. In: Proc. Int. Conf. on Machine Learning (ICML), Honolulu, Hawaii, USA, July; 2023. p. 7264–7302. https://proceedings.mlr.press/v202/daw23a.html.
[242] Krishnapriyan AS, Gholami A, Zhe S, et al. Characterizing possible failure modes in physics-informed neural networks. In: Proc. Advances in Neural Information Processing Systems (NeurIPS), December, 2021. p. 26548–26560. https://proceedings.neurips.cc/paper/2021/hash/df438e5206f31600e6ae4af72f2725f1-Abstract.html.
[243] Wang S, Teng Y, Perdikaris P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. SIAM J Sci Comput. 2021;43(5):A3055–A3081. URL . doi: 10.1137/20M1318043
[244] Cai S, Mao Z, Wang Z, et al. Physics-informed neural networks (pinns) for fluid mechanics: a review. Acta Mech Sin. 2021;37(12):1727–1738. doi: 10.1007/s10409-021-01148-1
[245] Cuomo S, Di Cola VS, Giampaolo F, et al. Scientific machine learning through physics–informed neural networks: where we are and what's next. J Sci Comput. 2022;92(3):88. doi: 10.1007/s10915-022-01939-z
[246] Huang B, Wang J. Applications of physics-informed neural networks in power systems-a review. IEEE Trans Power Syst. 2023;38(1):572–588. doi: 10.1109/TPWRS.2022.3162473
[247] Vowels MJ, Camgoz NC, Bowden R. D'ya like dags? a survey on structure learning and causal discovery. ACM Comput Surv. 2023;55(4):1–36. doi: 10.1145/3527154
[248] Pearl J. Causality. Cambridge, United Kingdom: Cambridge University Press; 2009.
[249] Glymour MM, Greenland S. Causal diagrams. Modern Epidemiology. 2008;3:183–209.
[250] Schölkopf B, Hogg DW, Wang D, et al. Modeling confounding by half-sibling regression. Proc Natl Acad Sci USA. 2016;113(27):7391–7398. doi: 10.1073/pnas.1511656113
[251] Thomas CK, Saad W, Xiao Y. Causal semantic communication for digital twins: a generalizable imitation learning approach. IEEE J Sel Areas Inf Theory. 2023;4:698–717. doi: 10.1109/JSAIT.2023.3336538
[252] Bongers S, Forré P, Peters J, et al. Foundations of structural causal models with cycles and latent variables. Ann Statist. 2021;49(5):2885–2915. doi: 10.1214/21-AOS2064
[253] Grieves M, Vickers J. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In: Transdisciplinary perspectives on complex systems: new findings and approaches. Cham, Switzerland: Springer; 2016. p. 85–113.
[254] Di Vito V, Mohammadian M, Baker K, et al. Learning to solve differential equation constrained optimization problems. In: International Conference on Learning Representations; 2025.
[255] Lueg LR, Alves V, Schicksnus D, et al. A simultaneous approach for training neural differential-algebraic systems of equations. arXiv preprint arXiv:2504.04665. 2025.
[256] Shapovalova M, Tsay C. Training neural odes using fully discretized simultaneous optimization. arXiv preprint arXiv:2502.15642. 2025.
[257] Brouillard P, Lachapelle S, Lacoste A, et al. Differentiable causal discovery from interventional data. Adv Neural Inf Process Syst. 2020;33:21865–21877.
[258] Addanki R, Kasiviswanathan S, McGregor A, et al. Efficient intervention design for causal discovery with latents. In: International Conference on Machine Learning. PMLR; 2020. p. 63–73.
[259] Yang J, Zhou K, Li Y, et al. Generalized out-of-distribution detection: a survey. Int J Comput Vis. 2024;132(12):5635–5662. doi: 10.1007/s11263-024-02117-4
[260] Shamanna P, Erukulapati RS, Shukla A, et al. One-year outcomes of a digital twin intervention for type 2 diabetes: a retrospective real-world study. Sci Rep. 2024;14(1):25478. doi: 10.1038/s41598-024-76584-7
[261] Guidotti R, Monreale A, Ruggieri S, et al. A survey of methods for explaining black box models. ACM Comput Surv. 2019;51(5):1–42. doi: 10.1145/3236009
[262] Jacod J. Discretization of processes. Berlin, Germany: Springer; 2012.

[263] Al-Badour F, Sunar M, Cheded L. Vibration analysis of rotating machinery using time–frequency analysis and wavelet techniques. Mech Syst Sig Process. 2011;25(6):2083–2101. doi: 10.1016/j.ymssp.2011.01.017
[264] Donzis DA, Yeung PK, Sreenivasan KR. Dissipation and enstrophy in isotropic turbulence: resolution effects and scaling in direct numerical simulations. Phys Fluids. 2008;20(4):1–16. doi: 10.1063/1.2907227
[265] Huang Y, Bárdossy A, Zhang K. Sensitivity of hydrological models to temporal and spatial resolutions of rainfall data. Hydrol Earth Syst Sci. 2019;23(6):2647–2663. doi: 10.5194/hess-23-2647-2019
[266] Shannon CE. Communication in the presence of noise. Proc IRE. 1949;37(1):10–21. doi: 10.1109/JRPROC.1949.232969
[267] Georgiou TT, Lindquist A. Dynamic relations in sampled processes. IEEE Control Syst Lett. 2019;3(1):144–149. doi: 10.1109/LCSYS.2018.2859481
[268] Chaurasia CS. In vivo microdialysis sampling: theory and applications. Biomed Chromatogr. 1999;13(5):317–332. doi: 10.1002/(SICI)1099-0801(199908)13:5<317::AID-BMC891>3.0.CO;2-I
[269] Almbauer RA, Piringer M, Baumann K, et al. Analysis of the daily variations of wintertime air pollution concentrations in the city of Graz, Austria. Environ Monit Assess. 2000;65(1-2):79–87. doi: 10.1023/A:1006464225727
[270] Gleaves JT, Ebner JR, Kuechler TC. Temporal analysis of products (TAP)—a unique catalyst evaluation system with submillisecond time resolution. Catal Rev Sci Eng. 1988;30(1):49–116. doi: 10.1080/01614948808078616
[271] Richecoeur F, Ducruix S, Scouflaire P, et al. Experimental investigation of high-frequency combustion instabilities in liquid rocket engine. Acta Astronaut (UK). 2008;62(1):18–27. doi: 10.1016/j.actaastro.2006.12.034
[272] Shang C-F, Fan Wang YU, Zhao M-T, et al. Real-time analysis of large-scale neuronal imaging enables closed-loop investigation of neural dynamics. Nat Neurosci. 2024;27(5):1014–1018. doi: 10.1038/s41593-024-01595-6
[273] Javh J, Slavič J, Boltežar M. High frequency modal identification on noisy high-speed camera data. Mech Syst Sig Process. 2018;98:344–351. doi: 10.1016/j.ymssp.2017.05.008
[274] He Y, Guo J, Zheng X. From surveillance to digital twin: Challenges and recent advances of signal processing for industrial internet of things. IEEE Signal Process Mag. 2018;35(5):120–129. doi: 10.1109/MSP.2018.2842228
[275] Mihai S, Yaqoob M, Hung DV, et al. Digital twins: a survey on enabling technologies, challenges, trends and future prospects. IEEE Commun Surv Tutorials. 2022;24(4):2255–2291. doi: 10.1109/COMST.2022.3208773
[276] Huang H, Yang L, Wang Y, et al. Digital twin-driven online anomaly detection for an automation system based on edge intelligence. J Manuf Syst. 2021;59:138–150. doi: 10.1016/j.jmsy.2021.02.010
[277] Stephen H, Kellert. In the wake of chaos: unpredictable order in dynamical systems. Chicago, IL, USA: University of Chicago Press; 1993.
[278] Boccaletti S, Grebogi C, Lai Y-C, et al. The control of chaos: theory and applications. Physics Reports. 2000;329(3):103–197. doi: 10.1016/S0370-1573(99)00096-4
[279] Farmer JD, Sidorowich JJ. Predicting chaotic time series. Phys Rev Lett. 1987;59(8):845–848. doi: 10.1103/PhysRevLett.59.845
[280] Brunton SL, Brunton BW, Proctor JL, et al. Chaos as an intermittently forced linear system. Nat Commun. 2017;8(1):19. doi: 10.1038/s41467-017-00030-8
[281] Xie X, Samaei A, Guo J, et al. Data-driven discovery of dimensionless numbers and governing laws from scarce measurements. Nat Commun. 2022;13(1):7562. doi: 10.1038/s41467-022-35084-w
[282] Atabekov A, Yastrebov O. Legal status of artificial intelligence across countries: legislation on the move. European Research Studies. 2018;21(4):773–782.
[283] Güneş Peschke S, Peschke L. Artificial intelligence and the new challenges for eu legislation. YBHD. 2022;2(2022):1267–1292. doi: 10.33432/ybuhukuk.1104344
[284] Parinandi S, Crosson J, Peterson K, et al. Investigating the politics and content of us state artificial intelligence legislation. Bus Polit. 2024;26(2):240–262. doi: 10.1017/bap.2023.40
[285] Lheureux F, Bürkle T, Gómez-González E, et al. Ethical and social issues of digital twins in health: a scoping review. Int J Med Inf. 2021;156:104610.
[286] Petrova-Antonova D, Ilieva S, Ivanova A. Ethical framework for development and implementation of a digital twin of a city. In: 2021 13th International Conference on-The "Challenge of Change" Digitalisation of the Economy and Society; 2021. . IEEE, 1–7.
[287] Diakopoulos N. Accountability in algorithmic decision making. Commun ACM. 2016;59(2):56–62. doi: 10.1145/2844110
[288] Kapteyn MG, Pretorius JVR, Willcox KE. A probabilistic digital twin for fatigue crack growth prediction. Nat Comput Sci. 2021;1(5):337–347. doi: 10.1038/s43588-021-00069-0
[289] McMahan B, Moore E, Ramage D, et al. Communication-efficient learning of deep networks from decentralized data. Artif Intell Stat. 2017;54:1273–1282.

[290] Dwork C. Differential privacy: a survey of results. In: International Conference on Theory and Applications of Models of Computation; 2008. Springer. p. 1–19.
[291] Mehrabi N, Morstatter F, Saxena N, et al. A survey on bias and fairness in machine learning. ACM Comput Surv. 2022;54(6):1–35. doi: 10.1145/3457607
[292] Deshpande A. Regulatory compliance and ai: navigating the legal and regulatory challenges of ai in finance. In: 2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS). Vol. 1. IEEE; 2024. p.1–5. doi: 10.1109/ICKECS61492.2024.10616752
[293] Brennan L. Ai ethical compliance is undecidable. Hastings Sci Tech LJ. 2023;14:311.
[294] Autischer G, Waxnegger K, Kowald D. Practical application and limitations of ai certification catalogues in the light of the ai act. arXiv preprint arXiv:2502.10398. 2025.