

Involving before Evolving: A Vision for Trustworthy Enterprise Digital Twin Engineering

Kérian Fiter
Polytechnique Montréal
Montréal, Canada
kerian.fiter@polymtl.ca

Franck Dervault
Michelin Canada
Montréal, Canada
franck.dervault@michelin.com

Adil Lagrou
Polytechnique Montréal
Montréal, Canada
adil.lagrou@polymtl.ca

Bentley Oakes
Polytechnique Montréal
Montréal, Canada
bentley.oakes@polymtl.ca

Abstract

Enterprise Digital Twins (EDTs) promise data-driven decision support at organizational scale, but realizing them requires navigating siloed departments, tacit knowledge, and high-stakes decisions with long-horizon consequences. Existing approaches involve domain experts during model development but focus less on early organizational buy-in in EDTs. We present a vision for trustworthy EDT engineering grounded in an ‘involving before evolving’ paradigm: rapidly involving stakeholders through a working prototype before evolving toward federation and full interoperability. Our three-stage approach combines foundation models for rapid prototyping, an ontological backbone for federated interoperability, and observability tooling for stakeholder trust. We ground our vision in an ongoing collaboration with Michelin, a multinational manufacturer, where an initial prototype has helped support stakeholder buy-in.

CCS Concepts

• **Information systems** → **Enterprise information systems**; • **Computing methodologies** → *Ontology engineering*.

Keywords

Enterprise Digital Twin, What-if Analysis, Uncertainty, Trust, Federation, Ontologies, Foundation Models, Large Language Models

ACM Reference Format:

Kérian Fiter, Adil Lagrou, Franck Dervault, and Bentley Oakes. 2026. Involving before Evolving: A Vision for Trustworthy Enterprise Digital Twin Engineering. In *ACM/IEEE 29th International Conference on Model Driven Engineering Languages and Systems (MODELS Companion 2026)*, October 04–09, 2026, Málaga, Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3837062.3839094>

1 Introduction

A Digital Twin (DT) is a virtual representation of a real system that enables its monitoring, prediction, and control through bidirectional data exchange [30]. DTs can have vastly different scales, from simple twinned systems to complex Systems of Twinned Systems (SoTS) [1]. Here, we consider an Enterprise Digital Twin (EDT) twinning an enterprise system, defined as “a set of purposive analyzable and simulatable models representing the enterprise in order to mimic real-world phenomena,” supporting “what-if scenario[s] playing in virtual space” [27]. Our EDT is a blend of an “Acknowledged SoTS” and “Directed SoTS” [1], as there exist models, goals, and overall independence at both department- and enterprise-level.

While DTs are defined by *automatic* bidirectional data exchange [26], automatic control in EDTs is neither feasible nor desirable for high impact business decisions, implying the loop can only be closed at a long time-scale. We term this a *long-loop* DT, where a human-in-the-loop affects the real system based on the DT, but at a longer time-scale (days, months, years) and with potentially irreversible consequences. Thus, EDTs remain within the DT paradigm: they receive enterprise data, use modeling, simulation, and AI techniques to predict behavior and prescribe actions, and close the loop through human-mediated organizational decisions. In particular, we focus on the lessons learned for the interoperable engineering of EDTs.

Context and Challenges. We ground our vision in an ongoing collaboration with Michelin, a multinational manufacturer involved in world-wide supply chains, marketing, and sales, which serves as our motivating context. Their business professionals operate across a complex structure with specialized Domain Knowledge (DK) and responsibilities. They must navigate a siloed data and processes ecosystem to model business scenarios and assess cross-functional impacts. In particular, they rely on *tacit knowledge* to perform what-if analysis [19] and make decisions, which lacks explainability, data-based grounding, and is at risk of disappearing when individuals leave the organization. This fragmented information environment limits the company’s ability to rapidly evaluate future scenarios, quantify uncertainty, and generate robust long-term forecasts for informed decision-making and organizational resilience.

We have extracted the following challenges for this EDT project through informal discussion with company stakeholders:

- C1 *Decision-making support* to explore business strategies
- C2 *What-if simulations* of financial scenarios



This work is licensed under a Creative Commons Attribution 4.0 International License. *MODELS Companion 2026, Málaga, Spain*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2903-4/2026/10
<https://doi.org/10.1145/3837062.3839094>

- C3 *Easy modeling* by non-technical stakeholders
- C4 *Interoperability* across (systems-of-)systems
- C5 *Observability* of the solution and the scenarios

2 Our Vision and Approach

Our vision is to develop a *trustworthy* and *interoperable* EDT for these departments to improve their prediction and decision-making capabilities. Large organizations can favor stability over innovation, and introducing a global EDT could face resistance from stakeholders wary of disrupting established workflows. We therefore argue for a socio-technical approach centered on an *involving before evolving* paradigm to secure organizational buy-in, before evolving the EDT toward higher levels of maturity and automation [2]. Our vision complements agent-based techniques [5–7] by focusing on the earlier socio-technical problem of securing organizational buy-in.

This *involving before evolving* vision is driven by the *long-loop* nature of EDTs. In most DTs, there is a tight feedback loop between the Physical Twin (PT) and DT [17], where errors have immediate consequences. In a long-loop DT, a prediction could drive a decision with severe long-term consequences, such as restructuring a supply chain or opening a new market, which raises the bar for stakeholders to understand and trust the EDT and its insights.

Proposed Approach. The three-stage approach is seen in Figure 1:

- (1) **Prototype Creation for Quick Involvement** (*Targets C1, C2, C3*). Our vision is to quickly collect requirements from Subject Matter Experts (SMEs), choose a formalism, and build a prototype (a low-maturity early instance of the EDT) using Foundation Models (FMs) to involve SMEs as soon as possible in the development process. We integrate SMEs’ DK into the simulation inputs and facilitate their validation of the low-maturity EDT.
In our approach, we propose using Large Language Models (LLMs), Time-Series Foundation Models (TSFMs) [25], and Monte Carlo simulation integrating uncertainty (ranges and distributions) to perform what-if financial simulations.
- (2) **Federation for Interoperable Evolution** (*Targets C4*). Soon after prototype validation, our vision is to federate [8] these models using an ontological backbone to formalize departmental DK. These predictions should be based on models which are federated and owned by each department. This will enable a federated prototype with organizational buy-in (adoption and trust), even across departments with entrenched silos and differing agendas.
Our approach proposes the Ontological Modeling Language (OML) [15] and related methodologies [24, 34] to support federation, semantic alignment, and interoperability.
- (3) **Observability for Promoting Trust** (*Targets C5*). Finally, building on the federated models, we envision an ecosystem where data and predictions are explainable, traceable, versioned, access controlled, consistent, and observable across silos to promote trustworthiness in the resulting ecosystem. Our approach proposes observability metrics and development of analysis tools suited to SMEs (such as DTInsight [16]) to increase the trustworthiness of the EDT.

Currently, we have created a decision support prototype for the marketing and finance departments. Our next step is to build a federated ontological backbone, before modeling other departments and making the system trustworthy for decision support.

3 Prototype Creation for Quick Involvement

We propose to quickly capture requirements in an informal textual manner and build a minimum viable product to support early validation before investing in full verification and formalization [9, 10]. While continuous *evolving* of the DT remains key [2], sustaining it requires buy-in from stakeholders in the form of a tangible prototype that immediately is *involving* them and their knowledge.

To achieve rapid deployment, we propose to develop User Interfaces (UIs) and models that would otherwise require expert DK with LLMs, and use FMs such as TSFMs [25] or action-conditioned World Models (WMs) [14], as rapid forecasting or simulation placeholders where expert-developed models are not yet available. FMs are broadly trained, large-scale models that can be adapted to many downstream problems: LLMs are FMs specialized for language understanding and generation; TSFMs are FMs specialized for temporal or sequential data using transformer architectures; and action-conditioned WMs are FMs specialized for learning compact predictive representations of environments and forecasting future states given current observations and planned actions. In particular, we go further than Giabbanelli *et al.*, who proposed LLMs as tool translators between models and simulation tools [20], and directly build models, tools, and interfaces with various FMs.

The outcome of this approach is to lift user interactions from error-prone and inconsistent spreadsheet calculations to automated and visual simulations to prioritize business-level decisions.

Motivating Example. We use our industrial context to illustrate this stage: after capturing informal requirements from multiple departments, we have built a light-weight tool whose UI is showcased in Figure 2. It takes historical time-series data as input and proposes a spline trajectory to the user based on the predictions of a locally-hosted TSFM [13] (TimesFM 2.5¹, 10/90th percentiles in green). The user can manually edit the spline points (in blue) and select the uncertainty range (10/90th percentiles in light red, 50th percentile in red) to create their trajectory prediction. This provides users a visual and real-time interface to easily embed their DK predictions. We then use Monte Carlo simulation to sample from these spline-defined uncertainty ranges and propagate uncertainty through downstream expert-developed models, producing distributions over predicted business outcomes.

This functional prototype, rapidly built using agentic LLMs hosted on company infrastructure, satisfies SME requirements and demonstrates a novel way to easily model scenarios from historical data with uncertainty control. This prototype focuses on clear visual understanding to support financial and marketing decision-making over the short and long term.

Despite this simple formalism, it has been effective in securing buy-in and motivating further discussion across marketing and finance departments. It led to the refinement and creation of requirements, such as how the tool could represent multiple datasets

¹<https://github.com/google-research/timesfm/>




Figure 1: Outline for our EDT vision, reflecting our current progress and planned steps.




Figure 2: UI prototype for integrating historical data, TSFMs, and uncertain trajectory predictions

to detect correlations, and how the TSFM could be fine-tuned or improved to deliver more precise predictions. This prototype will be extended to the ontological backbone described in the next section.

4 Federation for Interoperable Evolution

Interoperability is “the ability of systems to exchange information and use the information that has been exchanged” [28]. To realize this across departmental silos, we adopt a federated approach that reflects the organization’s social structure, following “business and architecture isomorphism” [35], as demonstrated in organizational modeling approaches such as OrgML [3]. Rather than a federated DT in the sense of Navarro *et al.* (“a composition of autonomous DTs” [33]), we pursue a DT characterized by coordinated autonomy.

We aim for federation to let teams retain ownership of their data, models and services, while committing to shared semantic contracts that guarantee interoperability. Here, ontologies act as the backbone of these contracts in three main roles: (1) *structural wiring* of components and their relations, ensuring that when departments connect their DTs, the ontology defines how each component can legally be composed, preventing structurally invalid federations

from being created in the first place; (2) *typed interfaces* so that data exchanged between departments carries an ontologically grounded type to reduce ambiguity [22] (e.g., when marketing sends a demand signal to the supply chain, both sides are aligned on what that signal represents); (3) *domain semantics* expressed as constraints and rules to enable automated consistency checks and inference (e.g., the supply chain may only act on a demand forecast if it has been approved by finance, enabling automated reasoning to catch any violation of this sort without human auditing).

Beyond federated DTs [29], our Federated EDT approach mirrors the enterprise’s organization that is split into departmental silos with distinct agendas, sets of tools, and DK. These departments operate over large-scale datasets that are intentionally segregated to preserve confidentiality and ensure accountability. As the silos evolve independently over time, it is essential to specify and maintain stable versioning and well-defined interfaces across their models and systems to ensure interoperability and reliability for cross-cutting, multi-department what-if analysis.

Our approach, based on the proposal from Humphries *et al.* [24], suggests that SMEs from each department express their DK as an owned ontology, interoperable with the others through shared contracts and rules. The key insight is in how we use the ontologies: using them merely as *taxonomies* would resolve terminology clashes between departments but would not support reasoning. This is performed by encoding domain semantics as inference rules within each department’s own vocabulary.

For example, historical data can require different fidelity depending on the domain in which it is applied. One market may exhibit long cycles and therefore require ten years of historical data, whereas another only needs six months. Capturing this DK in ontologies lets us reason over models that use multiple historical inputs and determine whether the overall fidelity is adequate. If a model draws on two high-fidelity inputs and one low-fidelity input, the reasoning layer could flag the model’s output as degraded.

Crucially, vocabulary owners must be able to author and understand their own vocabularies. We therefore adopt OML, a Domain-Specific Language (DSL) over OWL for defining concepts, relationships, constraints, and basic rules, as well as associated tooling [15]. This reduces the friction required to develop ontologies in OWL or Protégé [32], yet preserves the semantic rigor needed for typed interfaces and automated checks. To further aid development, we can layer a DSL or generation tool on top of OML as needed.

5 Observability for Promoting Trust

SMEs rely on their experience to make business decisions. Therefore, they are wary of decision-support tooling unless they trust its predictions. We define trustworthiness as the ability of stakeholders to inspect, understand, trace, and challenge EDT outputs, while preserving data sovereignty and organizational accountability.

In our industrial example, we support trustworthiness through *observability* [31] by employing an explainable formalism based on probability (Monte Carlo simulation) that can be analyzed through techniques such as sensitivity analysis. We keep the stakeholders in control by allowing them to define inputs through curve editing (see Figure 2), which aligns with their DK understanding, i.e., market tendencies and seasonal trends. Furthermore, our prediction module will feature versioning, ensuring stakeholders understand which tool version and datasets were used to generate the final prediction.

Beyond observability and human oversight, EDT trustworthiness also depends on the privacy and safety of the supporting tooling, since the EDT processes critical business data. This requires local FMs or models that ensure the organization’s data sovereignty. We employ agentic LLMs hosted on the company’s infrastructure to develop the tool, and a local TSFM instance runs in a container to generate time-series forecasts on historical and sensitive data.

To broaden this trust to the entire EDT development life-cycle, we plan to enhance our DT observability tool DTInsight [16] with lightweight OML modeling capabilities. This aligns with our vision of continuously generating a “constellation” (a conceptual visualization of a DT) [21] alongside a dynamic reporting page that tracks the EDT’s evolution, thereby enabling stakeholders to actively collaborate on and discuss the system as it is being built.

6 Discussion

Expected benefits. The expected benefits of our vision are (1) rapid and increased buy-in in EDTs from stakeholders through early involvement with a prototype; (2) reduced errors and better collaboration between departments through ontology-backed semantic alignment; (3) increased agility due to the federated approach that allows departments to focus on and control their data and models; (4) trusted probabilistic what-if simulations that are observable, traceable, and privacy-respecting; (5) increased overall resilience for the enterprise due to data-driven decision-making support.

Research agenda alignment. Our model-based DT engineering survey [36] found an open research challenge in making Model-Based Engineering tooling accessible to non-experts, and we extend that concern to organizational adoption. Our vision further aligns with Combemale *et al.*’s DT engineering roadmap [12]:

RQ_P2 “What communication tools and practices can enhance collaboration between data scientists, engineers, and operations teams in a DT engineering environment?”

RQ_P3 “What communication strategies can ensure that all stakeholders are informed and engaged throughout the DT models lifecycle? [...]”

Related approaches. Our work relates to Franzén’s exploration of design and trade spaces in early Systems-of-Systems (SoS) development [18]. They capture needs in an ontology that outputs a reduced design space through reasoning. SoS simulations generate performance measures to identify the best architecture, which can then be translated into interactive visual analytics. Conversely, we focus on EDTs and propose to involve stakeholders early through a prototype that supports iterative design-space exploration and organizational adoption, using ontologies to federate the SoS.

Barat *et al.* have extensive work on EDTs [4–6]. However, they focus on agent-based techniques using the Enterprise Simulation Language (ESL) [11]. Instead, we propose a top-down approach that considers the EDT as a digital wrapper over the social structure of the organization with simulation placeholders. We provide the formalism (Monte Carlo simulation) and interfaces for SMEs to work within their DK in a federated manner. Barat *et al.* also touch on the need for SMEs to increasingly focus on validation due to improving LLM modeling capabilities [7]. Our vision goes further into prototyping and trustworthiness concerns.

We also see potential in action-conditioned WMs as placeholders for simulating business behavior and finding optimal scenarios. These models are currently developed mostly on video data, but we expect their generalized ability to explore possible futures will be useful for businesses [23].

7 Conclusion

This paper has explored our vision of *involving before evolving*, where we target non-technical stakeholders divided among organizational silos. Currently, we have a working EDT prototype with stakeholder buy-in across Michelin departments.

The next step is to connect this prototype to the ontological backbone, using the methodological framework of Humphries *et al.* [24]. This will support versioned releases of models and datasets that can be composed into a trustworthy EDT. We also plan to make requirements gathering more incremental by progressively formalizing requirements against the emerging ontology. Although this reverses the order adopted in some ontology-first approaches [18], it enables discussions to start from a validated prototype and gradually converge toward a shared formal language.

Open challenges. First, it is unclear whether ontology-backed contracts can handle real semantic heterogeneity at scale, or whether automated consistency checks remain feasible as the federation grows. Second, scenario provenance, i.e., labeling, sharing, and comparing trajectory scenarios across departments over time, is important for EDTs but is not yet handled in our approach. Third, we assume that EDT-informed decisions will produce better organizational outcomes. However, the long-loop time-scale makes this methodologically difficult to evaluate without a longitudinal study across multiple deployments, which is not feasible in our project.

Acknowledgments

Feedback on the *long-loop DT* concept came from discussions with Istvan David (McMaster University), Houari Sahraoui (Université de Montréal), and Maxime Lamothe (Polytechnique Montréal). We also sincerely thank our partner Michelin, Mitacs Accelerate (IT38322), and NSERC Discovery (RGPIN-2024-05622) for funding support.

References

- [1] Feyi Adesanya, Kanan Castro Silva, Valdemar V Graciano Neto, and Istvan David. 2025. Systems of twinned systems: A systematic literature review. *arXiv preprint arXiv:2505.19916* (2025).
- [2] Tarek Alskaf, Onder Babur, Francis Bordeleau, Loek Cleophas, Benoît Combemale, Joachim Denil, Øystein Haugen, Judith Michael, Phu Nguyen, Tiberiu Secleanu, et al. 2025. Evolution at the core of digital twin engineering. In *2025 ACM/IEEE 28th International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. IEEE, 210–216.
- [3] Souvik Barat, Balbir Barn, Tony Clark, and Vinay Kulkarni. 2020. OrgML - A Domain Specific Language for Organisational Decision-Making. In *The Practice of Enterprise Modeling*, Jánis Grabis and Dominik Bork (Eds.). Springer International Publishing, Cham, 155–170.
- [4] Souvik Barat, Vinay Kulkarni, and Kaustav Bhattacharya. 2022. Enterprise digital twins for risk free business experimentations. In *2022 Winter Simulation Conference (WSC)*. IEEE, 2864–2875.
- [5] Souvik Barat, Vinay Kulkarni, Tony Clark, and Balbir Barn. 2019. An actor based simulation driven digital twin for analyzing complex business systems. In *2019 Winter Simulation Conference (WSC)*. IEEE, 157–168.
- [6] Souvik Barat, Dushyanthi Mulpuru, Himabindu Thogaru, Abhishek Yadav, and Vinay Kulkarni. 2025. Synergic Use of Modeling & Simulation, Digital Twins and Large Language Models to Make Complex Systems Adaptive and Resilient. In *2025 Winter Simulation Conference (WSC)*. IEEE, 474–485.
- [7] Souvik Barat, Dushyanthi Mulpuru, Abhishek Yadav, Reshma Korabu, Himabindu Thogaru, and Vinay Kulkarni. 2025. Constructing Enterprise Digital Twins by Augmenting LLMs with MDE. In *Proceedings of the 18th Innovations in Software Engineering Conference*. 1–11.
- [8] Luca Berardinelli, Vittorio Muttillio, Hugo Bruneliere, Romina Eramo, Alessio Buaioni, Benoît Combemale, Sashko Ristov, Philipp Gritsch, Djamel Eddine Khelladi, Dao Khanh Duy, et al. 2026. Barriers & Enablers for Traceability and Federation in Model-Based Digital Twin Engineering. *Information and Software Technology* (2026).
- [9] Johan Cederbladh, Antonio Cicchetti, and Robbert Jongeling. 2025. A road-map to readily available early validation and verification of system behaviour in model-based systems engineering using software engineering best practices. *ACM Transactions on Software Engineering and Methodology* 34, 5 (2025), 1–30.
- [10] Johan Cederbladh, Antonio Cicchetti, and Jagadish Suryadevara. 2024. Early validation and verification of system behaviour in model-based systems engineering: A systematic literature review. *ACM Transactions on Software Engineering and Methodology* 33, 3 (2024), 1–67.
- [11] Tony Clark, Vinay Kulkarni, Souvik Barat, and Balbir Barn. 2017. ESL: an actor-based platform for developing emergent behaviour organisation simulations. In *International conference on practical applications of agents and multi-agent systems*. Springer, 311–315.
- [12] Benoît Combemale, Pascale Vicat-Blanc, Arnaud Blouin, Hind Bril El Haouzi, Jean-Michel Bruel, Julien Deantoni, Thierry Duval, Sébastien Gérard, and Jean-Marc Jézéquel. 2025. Engineering Digital Twins: A Research Roadmap. In *2025 ACM/IEEE 28th International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. IEEE, 181–188.
- [13] Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. 2023. A decoder-only foundation model for time-series forecasting. *arXiv preprint arXiv:2310.10688* (2023).
- [14] Jingtao Ding, Yunke Zhang, Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan Su, Nian Li, Nicholas Sukiennik, et al. 2025. Understanding world or predicting future? a comprehensive survey of world models. *Comput. Surveys* 58, 3 (2025), 1–38.
- [15] Maged Elaasar, Nicolas Rouquette, David Wagner, Bentley James Oakes, Abdelwahab Hamou-Lhadj, and Mohammad Hamdaqa. 2023. openCAESAR: Balancing agility and rigor in model-based systems engineering. In *2023 ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. IEEE, 221–230.
- [16] Kérian Fiter, Louis Malassigné-Onfroy, and Bentley Oakes. 2025. DTInsight: A Tool for Explicit, Interactive, and Continuous Digital Twin Reporting. In *2025 ACM/IEEE 28th International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. 139–143. doi:10.1109/MODELS-C68889.2025.00030
- [17] John Fitzgerald, Cláudio Gomes, and Peter Gorm Larsen. 2024. *The engineering of digital twins*. Springer.
- [18] Ludvig Knöös Franzén. 2023. *A System of Systems View in Early Product Development: An Ontology-Based Approach*. Linköping University.
- [19] Sneha Gathani, Zhicheng Liu, Peter J Haas, and Çağatay Demiralp. 2025. What-if analysis for business professionals: Current practices and future opportunities. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [20] Philippe J Giabbanelli, John Beverley, Istvan David, and Andreas Tolk. 2025. From over-reliance to smart integration: Using large-language models as translators between specialized modeling and simulation tools. In *2025 Winter Simulation Conference (WSC)*. IEEE, 1119–1130.
- [21] Santiago Gil, Bentley Oakes, Cláudio Gomes, Mirgita Frasheri, and Peter G. Larsen. 2025. Toward a Systematic Reporting Framework for Digital Twins: A Cooperative Robotics Case Study. *SIMULATION* 101, 3 (2025), 313–339. doi:10.1177/00375497241261406
- [22] Joe Gregory, Steven Jenkins, Visalakshi Iyer, and Alejandro Salado. 2026. Driving Optimization Through an Ontological Digital Engineering Workflow. In *Proceedings of the INCOSE International Symposium* (Yokohama, Japan).
- [23] Daniyar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2025. Mastering diverse control tasks through world models. *Nature* 640, 8059 (2025), 647–653.
- [24] Lewis Humphries, Gianmaria Bullegas, John Golledge, Campbell McCausland, Donald Taylor, Dimitris Kolovos, Antonio Garcia-Dominguez, and Simos Gerasimou. 2025. Integrated Knowledge Centric Engineering: Delivering next-generation aircraft projects at pace. In *2025 ACM/IEEE 28th International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. IEEE, 834–843.
- [25] Siva Rama Krishna Kottapalli, Karthik Hubli, Sandeep Chandrashekhara, Garima Jain, Sunayana Hubli, Gayathri Botla, and Ramesh Doddaiiah. 2025. Foundation Models for Time Series: A Survey. *arXiv:2504.04011 [cs.LG]* <https://arxiv.org/abs/2504.04011>
- [26] Werner Kritzing, Matthias Karner, Georg Traar, Jan Henjes, and Wilfried Sihm. 2018. Digital Twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine* 51, 11 (2018), 1016–1022.
- [27] Vinay Kulkarni, Souvik Barat, and Tony Clark. 2019. Towards adaptive enterprises using digital twins. In *2019 Winter Simulation Conference (WSC)*. IEEE, 60–74.
- [28] Gabriel S. S. Leal, Wided Guédria, and Hervé Panetto. 2020. Enterprise interoperability assessment: a requirements engineering approach. *International Journal of Computer Integrated Manufacturing* 33, 3 (2020), 265–286. doi:10.1080/0951192X.2020.1736636
- [29] Yang Liu, Jun Feng, Jiamin Lu, and Siyuan Zhou. 2024. A review of digital twin capabilities, technologies, and applications based on the maturity model. *Advanced Engineering Informatics* 62 (2024), 102592.
- [30] Azad M Madni, Carla C Madni, and Scott D Lucero. 2019. Leveraging digital twin technology in model-based systems engineering. *Systems* 7, 1 (2019), 7.
- [31] Charity Majors, Liz Fong-Jones, and George Miranda. 2022. *Observability Engineering: Achieving Production Excellence*. O'Reilly Media, Sebastopol, CA.
- [32] Mark A Musen and Protégé Team. 2015. The Protégé Project: A Look Back and a Look Forward. *AI Matters* 1, 4 (2015), 4–12. doi:10.1145/2757001.2757003
- [33] Elena Navarro, Pascual González, Jose Luis de la Vara, Francisco Montero, María T López, Victor López-Jaquero, Elena Pretel, and Clara Ayora. 2026. Towards Developing Federated Digital Twins for Society. In *Proceedings of the 20th International Conference on Research Challenges in Information Science (RCIS)*.
- [34] Bentley Oakes, Cláudio Gomes, Eduard Kamburjan, Giuseppe Abbiati, Elif Ecem Bas, and Sebastian Engelsgaard. 2024. Towards Ontological Service-Driven Engineering of Digital Twins. In *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*. 464–469. doi:10.1145/3652620.3688261
- [35] Johannes Thönes. 2015. Microservices. *IEEE Software* 32, 1 (2015), 116–116. doi:10.1109/MS.2015.11
- [36] Philipp Zech, Souvik Barat, Benjamin Nast, Bentley Oakes, Judith Michael, Steffen Zschaler, Balbir Barn, and Ruth Breu. 2026. Model-based digital twin engineering: insights, challenges, and future directions. *Software and Systems Modeling* (04 May 2026). doi:10.1007/s10270-026-01368-8