A Knowledge-Guided Hybrid Learning Framework for Semantic Constraint Integration in Time Series Models

Simon Burbach

Professorship of Data Engineering, Helmut Schmidt University, Holstenhofweg 85, 22043 Hamburg, Germany

Abstract

Current time series models often operate solely on sensor data, lacking the contextual understanding that domain knowledge provides. This limitation particularly exists in domains like maritime operations or medical monitoring, where sensor data are often noisy, incomplete, or ambiguous. To address this gap, this doctoral research proposes a hybrid learning framework that integrates semantic knowledge from ontologies, domain texts, and expert-defined rules into the modeling process as formal constraints. The framework comprises three main building blocks: (1) learning joint representations from heterogeneous sources such as time series, structured knowledge, and unstructured text; (2) extracting and formalizing semantic knowledge into symbolic or functional constraints; and (3) fusing these components into a hybrid framework, where formal constraints complement machine-learned patterns. Initial work has been conducted in the maritime domain and will be extended to medical datasets for cross-domain evaluation.

Keywords

Hybrid Learning Framework, Semantic Constraints, Time Series Modeling, Knowledge Integration, Knowledge-Guided AI, Context-Aware Representation Learning

1. Problem Statement

Current machine learning (ML) models for time series analysis, particularly in domains such as maritime operations and medical monitoring, rely heavily on dynamic sensor data. This data is typically collected in real-time and processed by ML models such as Recurrent Neural Networks (RNNs) [1, 2] or Convolutional Neural Networks (CNNs) [3]. While these models have achieved notable success, they also exhibit critical limitations in practical applications.

In many real-world scenarios, sensor data is prone to intermittency, incompleteness, and distortion due to operational or environmental noise [4, 5]. For example, wearables may temporarily lose signal due to movement artifacts, or shipboard sensors may malfunction due to mechanical stress or weather interference [6, 7]. As a result, the training datasets for such systems often lack completeness and reliability. Although data preprocessing techniques mitigate various data quality issues, the resulting preprocessed data often still lacks *natural* completeness that reflects the inherent richness and context of real-world information [8]. This data sparsity and noise inevitably degrade model generalization, reducing performance in downstream tasks like anomaly detection and predictive maintenance [9, 10, 11].

This existing lack of contextual inclusion in data processing highlights a critical blind spot in many model architectures: the underutilization of static, semantically rich knowledge available in textual or structured formats. This includes, for instance, operational manuals, domain-specific ontologies, technical specifications, and guidelines created by experts. Although these sources encode essential relationships, behavioral patterns, and domain logic, they are rarely incorporated into ML pipelines [12]. Consequently, models must learn context purely from the available data samples, often without any external guidance. This limitation hinders their ability to make informed predictions in ambiguous or data-deficient situations [13].

ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan

burbachs@hsu-hh.de (S. Burbach)
https://www.hsu-hh.de/dataeng/en/ (S. Burbach)
0009-0009-8373-6622 (S. Burbach)

In contrast, physics-informed neural networks (PINNs) [14] as part of the emerging field of physics-enhanced machine learning (PEML) [15] have demonstrated the value of embedding domain knowledge directly into model training. PEML strategies use features derived from physical laws and expert knowledge to constrain and guide the model's learning process, resulting in improved interpretability, robustness, and performance, especially in low-data regimes [15, 16].

Despite progress in both data-driven and knowledge-informed modeling, a central unresolved problem remains: there is currently no systematic approach for jointly modeling heterogeneous knowledge sources, such as dynamic sensor data, domain texts, ontologies, and structured guidelines, within a unified learning framework. Existing methods often isolate these sources, failing to capture the complex interactions between temporal patterns and domain logic. This disconnect limits the ability of models to leverage rich semantic context during learning, particularly in scenarios where data is sparse, ambiguous, or noisy. As a result, ML systems struggle to generalize beyond surface-level patterns and lack the interpretability required for critical decision-making environments.

An approach to fill this gap could significantly advance the field of Prognostics and Health Management (PHM) [17] by enabling more context-aware and interpretable anomaly detection systems. Such systems would be particularly valuable in domains like maritime operations, where PHM tools are most effective when used in collaboration with human operators [18]. Embedding semantic knowledge could help bridge the gap between data-driven models and expert reasoning, supporting better decision-making. In practice, this may lead to earlier detection of failures, reduced maintenance costs, and fewer unexpected downtimes. For the research community, the proposed direction offers a novel pathway toward integrating symbolic knowledge with temporal learning models, which paces the way for more robust, human-aligned AI systems.

2. Related Work

Until the introduction of the Transformer architecture by Vaswani et al. [19], (RNNs) [1, 20, 21] and their gated versions, such as Long Short-Term Memory networks (LSTMs) [22, 23] were the state of the art for modeling temporal dependencies in ML [24, 25]. However, they suffer from inefficiency, as they require complex computational efforts for recurrence [24]. Transformers, on the other hand, emerged in 2017 as a more efficient way to learn long-range dependencies. Although their original design and primary application were in natural language processing and large language models (LLMs), they have also shown strong performance on time series tasks [26, 27, 28].

To enhance the performance of downstream tasks following a Transformer-encoder by involving semantic information, a model must embed the semantic data with the underlying time series inputs. One approach that addresses the challenge of integrating semantic information into sequential data is GraphCare [29]. This framework focuses on healthcare prediction using electronic health records (EHRs) by constructing personalized knowledge graphs for each patient. To do so, GraphCare uses both external biomedical knowledge bases and LLMs to extract concept-specific subgraphs, which are then composed into individual graphs that capture temporal and relational context. GraphCare demonstrates how semantic structures from LLMs and knowledge graphs can enhance the predictive performance of time-aware models in a real-world use case [29]. In addition to approaches like GraphCare, a growing number of methods have explored the integration of LLMs into time series modeling to inject semantic context into learning pipelines [30, 31, 32, 33]. One such approach is presented below to demonstrate how to embed semantics with time series data in detail.

The Language Time series Model (LTM) framework [30] enhances multiple tasks of time series analysis (e.g., forecasting, imputation, anomaly detection) by combining temporal data with semantic context from large language models (LLMs). First, user prompts are enriched using external knowledge sources via a knowledge graph and GraphRAG [34] retrieval, forming contextualized instructions.

These prompts are embedded and fused with time series patches through the Fusion-Aware Temporal Module (FATM). The alignment between natural language prompts and time series patches is not learned through supervision but rather through a combination of the chosen model architecture and extension of the loss function. To enable this, the model initially assumes that the prompt is relevant to all time series patches. Without labeled assignments, it learns to distinguish which prompt-patch relationships are significant through training step by step. Hao et al. [30] are adding the cosine similarity term to the loss function, which encourages the model to combine prompts and embedded patches semantically. This setup allows the model to implicitly learn which parts of the prompt are most relevant to different segments in the time series while being guided by the performance of the selected task.

The LTM framework demonstrates how domain knowledge can be effectively integrated with time series data through semantic prompts and learning based on a shared latent space. It shows that a meaningful connection between language and temporal patches can be created without explicit supervision, but only guided by task performance and the architecture of the model. Their approach also shows how the manipulation of the loss function with contextual similarity information can increase the overall performance [30].

However, the potential of the loss function extends beyond learning contextual relationships. It also offers an effective approach for embedding external knowledge that can actively regulate and constrain model outputs. This becomes clear when comparing the approach to Physics-Informed Neural Networks (PINNs) [14], which oftentimes embed physical laws directly into the training process by incorporating differential equations or boundary conditions into the loss function. Instead of learning only from data, PINNs are guided by domain-specific constraints, which serve as a form of supervision and improve generalization, especially in data-scarce scenarios [14]. Inspired by this principle, semantic rules, expert-defined guidelines, or logic-based structures that are derived from language or symbolic reasoning could similarly serve as a form of prior knowledge to guide model behavior. This motivates the idea of generating explicit equations or constraints from semantic sources and incorporating them into time series learning, similar to how physical constraints guide PINNs.

3. Research Questions

Building on the identified challenges in integrating contextual knowledge into time series modeling, this early-stage research aims to answer the following guiding question:

Research Question: How can a knowledge-guided hybrid learning framework for time series data be designed to jointly analyze heterogeneous data sources and semantic knowledge?

This main research question can be broken down (see Figure 1) to guide the research process:

RQ1 – Semantic Representation. How can heterogeneous input types, such as time series, unstructured text, and structured knowledge, be embedded together within a transformer-based model to enable integrated contextual representation learning?

This part explores methods for fusing multimodal sources into a unified semantic representation and forms the necessary basis for the work on RQ2. The aim is to encode domain logic and contextual cues alongside raw temporal signals to improve downstream learning performance. The main challenge lies in integrating a highly diverse range of knowledge sources, each of which is typically handled by a separate solution. As illustrative examples, the SHIP ontology [35], which extends SSN/SOSA [36, 37], and together with the MontoFlow framework [35], can be used to instantiate both static and dynamic sensor data. OTMKGRL [38] provides a strategy for embedding visual and textual inputs into a shared space. RDF2Vec [39] applies language modeling techniques to RDF graph walks to produce embeddings of structured knowledge. OWL2Vec [40] enables the transformation of ontological structures into vector representations. JOIE [41] allows for the joint embedding of ontological concepts and instance-level data within a unified latent space.

Figure 1 visualizes the approach by processing time series ($\mathbf{x}$) and semantic data ($\mathbf{y}$) in parallel, by first patching the time series data and then embedding both parts. After that, the time series need to be combined with their matching semantics. This can be achieved by a classification or similar task. The unified representation is then fed into a transformer encoder to jointly embed the input to learn relations and give time series data a semantic context.

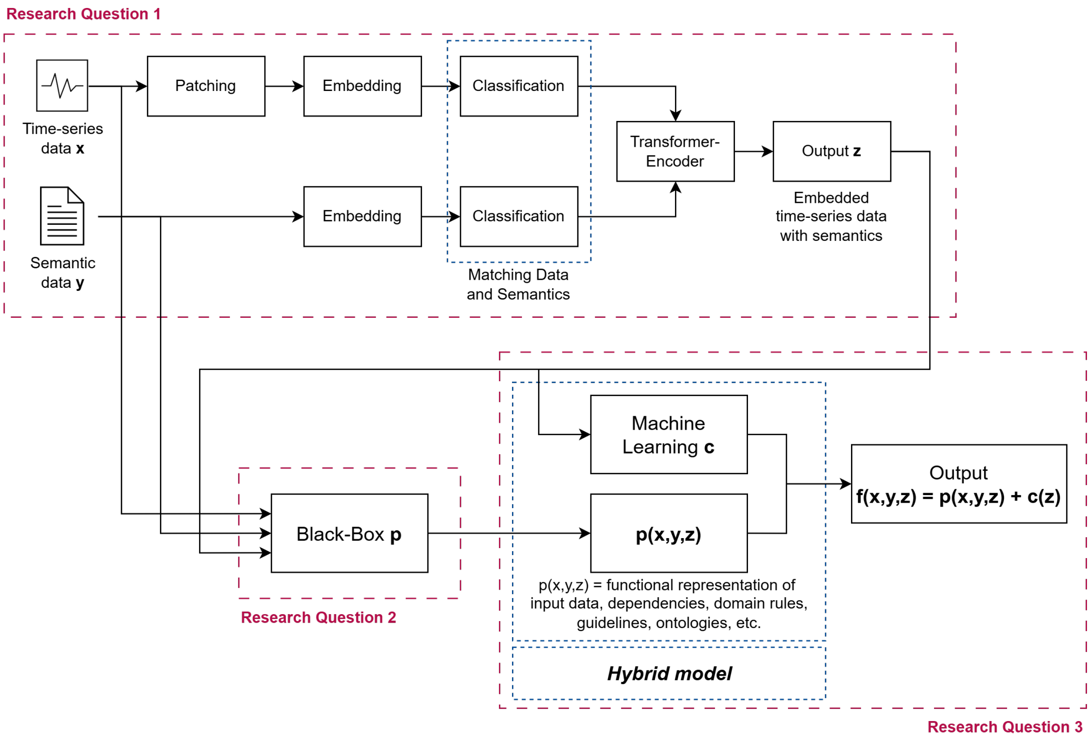

Figure 1: A graphical representation of the PhD thesis. Directional arrows show the data flow through the model. Red dashed boxes divide the model into the individual research questions.

RQ2 – From Time Series and Semantics to Formal Constraints (main focus). How can time series data and semantic knowledge extracted from domain texts be translated into formalized constraints that guide the learning process in time series models?

RQ2 is the main focus of this PhD work and explores how to extract rules, dependencies, or constraints from the input data and formalize them (e.g., differential equations or symbolic logic rules). It remains an open design decision to determine which parts of the input can be used for the equation. Hence, Figure 1 shows time series (**x**), semantics (**y**), and joint embeddings (**z**) as possible inputs. It still needs to be investigated which method is most suitable for generating a formal representation. As for now, we call the formal constraint learning the "black-box" (**p**), which is further discussed below. Possible approaches include the use of regressors, symbolic mathematical models, physically inspired formulations, or hybrid strategies that combine these elements.

Building on these general strategies, recent advances in symbolic regression (SR), such as Sparse Identification of Nonlinear Dynamics (SINDy) [42], Python Symbolic Regression (PySR) [43], and their improved variants, provide a principled starting point for discovering candidate formulas and constraints that capture dependencies in multivariate time series. In our framework, leveraging state-of-the-art techniques from other domains, these methods can discover compact, interpretable relations among observed variables and learned embeddings, which can be encoded to guide the predictive model [44],[45]. One of the ways to build closed-form symbolic relations from the data is to pair SR with neural components in a hybrid setup (for example, neural features or derivative estimates feeding PySR, SINDy, or neural guidance over operator libraries), so that the network learns rich representations, while the symbolic module recovers closed-form structure.

In parallel, we aim to formalize **z** (semantic knowledge from domain text units, invariants, bounds, and monotonicities or temporal relations), as machine-checkable constraints (algebraic equalities or inequalities, dimensional-consistency rules, or temporal logics). Thus, the integration of SR, neural representation learning, and semantics fulfills RQ2 by converting time-series and domain semantics into formal constraints and improving generalization while yielding interpretable inductive biases.

RQ3 – Fusion to Hybrid Model. How can a machine learning model and a formalized representation of system behavior be integrated within a hybrid architecture to complement each other and enhance generalization performance?

It is unlikely that a complete and precise formulation of system dynamics can be derived purely from the available data. Therefore, RQ3 merges the joint embeddings (**z**) and formal components (**p(x,y,z)**) developed in RQ1 and RQ2 into a hybrid architecture with a complementary ML model, needed to learn what cannot be explicitly defined (**c(z)**) to capture subtleties, exceptions, and context-dependent variations. The constraints, learned from the black box (**p**), would act as functional regulators within the model (**c**) that guide the learning process and constrain predictions in a way that aligns with expert reasoning and domain logic. Combining both techniques enables a hybrid system that benefits from the strengths of explicit reasoning and data-driven adaptability. To enable this integration, the formal expressions may be incorporated into the model either as architectural components in the form of constraint-aware layers or through modifications to the loss function. Loss-function–based approaches, as seen in Section 2, enforce domain-aligned behavior, which allows the formal representation to actively guide the learning process alongside the data-driven components.

The motivation behind this approach is taken from PEML. PEML integrates domain knowledge and physical principles into data-driven models to overcome the limitations of relying solely on observational data, particularly the inability to generalize well to unseen scenarios. These kinds of hybrid approaches incorporate various forms of knowledge, from domain expertise and empirical observations to first principles and mathematical formulations, by embedding observational, learning, inductive, model form, and discrepancy biases [46]. Such biases enable models to achieve the "inductive leap," guiding them toward physically meaningful generalizations beyond what is strictly inferred from training data [47]. Depending on the strategy, physics can be integrated in different ways [15]. For example, 1) Physics-Guided ML leverages detailed physics-based models as the backbone, refining them with data to identify latent parameters and improve predictive accuracy. Examples include probabilistic model updating strategies [48]. 2) Physics-Informed ML constrains data-driven models with physics-based laws and biases, ensuring solutions remain physically plausible, with examples including models like PINNs [49]. 3) Physics-Encoded ML embeds physics directly into the structure of the algorithms, examples including PhI-SINDy [46]. By combining these strategies, PEML accelerates training, enhances generalization under limited data conditions, and ensures that models remain consistent with established scientific understanding while retaining the flexibility of modern ML techniques.

4. Preliminary Results

As a foundational step toward the hybrid framework, we have already gained access to maritime sensor datasets collected from multiple Search and Rescue (SAR) vessels. These datasets contain time series covering hundreds of onboard sensors, which have been recorded over several years with high temporal resolution. Based on these data and domain requirements, we developed an extensible semantic modeling framework called MontoFlow, which is centered around the SHIP Ontology [35].

The SHIP Ontology is a domain-specific extension of the W3C SSN/SOSA standard [36, 37], which provides a rich semantic model of ship sensors, components, and operational context. It introduces concepts such as anomaly observations, value thresholds, and sensor classifications that describe the structure and behavior of maritime systems. MontoFlow connects this ontology with dynamic sensor data using a dual instantiation pipeline: MontoFlow-Static allows for structured ABox population from tabular configuration files, while MontoFlow-Dynamic enables real-time semantic querying over telemetry streams via virtual RDF mappings using Ontop.

Together, SHIP and MontoFlow support **RQ1** by enabling the semantic integration of time series data with structured domain knowledge. For **RQ2**, SHIP provides a vocabulary for identifying domain rules that can be formalized into constraints. Initial results show that the framework allows for real-time semantic enrichment of sensor data, laying the foundation for embedding expert knowledge into time series models.

5. Evaluation

This work implements an evaluation strategy that regards multiple levels to assess the validity of the proposed hypotheses and to answer the research questions systematically. Each research question will be addressed through targeted experiments in both use cases: maritime operations and medical monitoring. The perspective from both domains allows us not only to analyze the behavior of the proposed methods across different data modalities but also their generalizability across different application contexts.

Evaluation of RQ1 – Semantic Representation. As a foundational step, we test whether the integrated model can learn context-aware representations from heterogeneous inputs, including time series, structured knowledge, and unstructured text, in order to evaluate the effectiveness of the semantic representation strategy. This will be done by comparing downstream task performance (e.g., anomaly detection or forecasting) with and without the integration of semantic components based on a compact ship sensor dataset and structured sources, like the SHIP ontology.

After successfully proving the effect, a large dataset of ship sensors will be combined with extended unstructured content like maintenance manuals. The expectation is that the added context will improve the model's ability to identify abnormal engine states or faulty processes. Finally, in the medical domain, wearable glucose sensor data will be combined with medical guidelines, ICD-10 codes, and patient records. In this case, we expect a more accurate detection of early hypoglycemic events due to the added semantic context.

The results from this step are crucial for enabling the main investigations in RQ2.

Evaluation of RQ2 – From Time Series and Semantics to Formal Constraints. To evaluate whether time series data and semantic knowledge can be translated into formalized constraints that effectively guide model learning, we conceptualize a two-step evaluation process:

As the central focus of this work, we first compare the solutions against known ground truth of established mathematical models, like fuel efficiency in maritime systems [50] or the glucose-insulin dynamics in medical monitoring [51]. This comparison allows us to assess symbolic and semantic similarity, as well as to validate whether the extracted constraints capture meaningful relationships recognized by domain experts.

Next, we treat the extracted constraints as predictive functions and test their behavior against real-world sensor data. Their accuracy can be assessed by using standard metrics (e.g., RMSE, MAE, R^2). This helps determine the extent to which the formalized representation approximates real-world system behavior under varying data conditions.

The insights gained here are expected to play a key role in shaping the hybrid architecture evaluated in RQ3.

Evaluation of RQ3 – Fusion to Hybrid Model. The third research question focuses on how well the formal knowledge representations and ML components can be fused into a hybrid architecture. To evaluate this, we will benchmark the hybrid model against standard baselines such as Transformers and LSTMs that operate without semantic integration. The evaluation begins with maritime sensor data, where the hybrid model will be tested in scenarios such as engine diagnostics or anomaly detection. In the second step, we transfer the approach to medical datasets that include patient profiles and glucose sensor data – data that is often incomplete or noisy. This progression allows us to assess how well the hybrid model generalizes across domains and data qualities. Indicators to evaluate the performance will include predictive accuracy, robustness to missing values, and adaptability across the two mentioned domains.

Baseline Comparison with LLMs. Recent advances in LLMs have made them a valuable tool for a wide range of tasks [52, 53], and we intend to leverage this opportunity by using them as a common baseline throughout our research. By systematically comparing our methods against LLM-based baselines, we can track and quantify the improvements achieved over each task section.

We assume that LLMs offer a promising approach for extracting formal constraints, thanks to their ability to capture contextual relationships in complex domains, thereby positioning them as a natural point of reference across the research questions and the pipeline as a whole.

For **RQ1**, LLMs could serve as a baseline for semantic representation and context integration by matching time-series data patches to their affiliated semantic descriptions. As an additional baseline, structured knowledge may be linearized into natural language statements and embedded with an LLM, allowing a direct comparison with specialized graph-based embedding methods such as RDF2Vec, OWL2Vec, or JOIE. For **RQ2**, we will employ LLMs to extract formal constraints from text and assess their effectiveness relative to alternative hybrid PySR and SINDy approaches. For **RQ3**, LLMs will be used as a baseline model for predicting use-case related events (e.g., hypoglycemia, due dates for ship maintenance) based on text. The results are compared against those obtained from a model using only time series data, as well as from the hybrid approach. This strategy allows us to evaluate the added value of integrating formal knowledge representations with machine learning components.

Overall, the evaluation will assess whether semantic integration improves performance, interpretability, and robustness. Success is defined by outperforming baselines, leveraging meaningful constraints, and maintaining reliability under real-world conditions. Ultimately, the goal is to support more transparent and context-aware AI across domains.

6. Reflection and Future Work

Despite the potential of the proposed hybrid framework, there are limitations that must be taken into account. First, the alignment between the semantic knowledge and time series data remains challenging. Time series data rarely comes with labels. As a result, assigning relevant semantic annotations to specific segments of a time series requires manual effort. This limits the automation and scalability of the approach.

Second, the availability and quality of domain knowledge vary across domains. While maritime systems offer relatively accessible technical documentation and structured sensor descriptions, medical knowledge often exists in unstructured, proprietary, or fragmented formats. If they are available, they are not machine-readable or formalized most of the time.

Third, although the hybrid model aims to increase interpretability by incorporating structured knowledge, the data-driven model still operates like a black-box. This means, in practice, that this part of the framework remains difficult to interpret.

These limitations do not invalidate the approach, but they highlight areas where further methodological refinement is needed. The next phase of this work will focus on the first steps of the proposed framework (see Figure 1). A first step involves developing a method for segmenting time series data into meaningful patches. This process is either based on fixed intervals or event-driven. The latter is the more promising approach. Further investigation into semantic representation techniques will be done. These include the proposed options in Section 3, among others that we identify during research.

To strengthen the formal components of the model, PINNs will be further explored for techniques that reveal how to derive formal representations from input data. With regard to the data side, maritime datasets will be prepared with annotated events, while public medical datasets, like MIMIC-III or OhioT1DM, will be reviewed with a focus on glucose signal availability and quality.

Acknowledgments

This work is under the supervision of Prof. Maria Maleshkova. I thank Vaibhav Gupta for the inspiring discussions.

Declaration on Generative AI

During the preparation of this work, the author used ChatGPT and Grammarly in order to: Grammar and spelling check, paraphrase, and reword. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] Z. Che, S. Purushotham, K. Cho, D. Sontag, Y. Liu, Recurrent Neural Networks for Multivariate Time Series with Missing Values, Scientific Reports 8 (2018). URL: http://dx.doi.org/10.1038/s41598-018-24271-9. doi:10.1038/s41598-018-24271-9.

[2] P. B. Weerakody, K. W. Wong, G. Wang, W. Ela, A review of irregular time series data handling with gated recurrent neural networks, Neurocomputing 441 (2021) 161–178. URL: http://dx.doi.org/10.1016/j.neucom.2021.02.046. doi:10.1016/j.neucom.2021.02.046.

[3] A. P. Wibawa, A. B. P. Utama, H. Elmunsyah, U. Pujianto, F. A. Dwiyanto, L. Hernandez, Time-series analysis with smoothed Convolutional Neural Network, Journal of Big Data 9 (2022). URL: http://dx.doi.org/10.1186/s40537-022-00599-y. doi:10.1186/s40537-022-00599-y.

[4] H. Y. Teh, A. W. Kempa-Liehr, K. I.-K. Wang, Sensor data quality: a systematic review, Journal of Big Data 7 (2020). URL: http://dx.doi.org/10.1186/s40537-020-0285-1. doi:10.1186/s40537-020-0285-1.

[5] A. B. Sharma, L. Golubchik, R. Govindan, Sensor faults: Detection methods and prevalence in real-world datasets, ACM Transactions on Sensor Networks 6 (2010) 1–39. URL: http://dx.doi.org/10.1145/1754414.1754419. doi:10.1145/1754414.1754419.

[6] D. Pollreisz, N. TaheriNejad, Detection and Removal of Motion Artifacts in PPG Signals, Mobile Networks and Applications 27 (2019) 728–738. URL: http://dx.doi.org/10.1007/s11036-019-01323-6. doi:10.1007/s11036-019-01323-6.

[7] F. Sezgin, D. Vriesman, D. Steinhauser, R. Lugner, T. Brandmeier, Safe Autonomous Driving in Adverse Weather: Sensor Evaluation and Performance Monitoring, 2023. URL: https://arxiv.org/abs/2305.01336. doi:10.48550/ARXIV.2305.01336.

[8] A. W. Blocker, X.-L. Meng, The potential and perils of preprocessing: Building new foundations, Bernoulli 19 (2013). URL: http://dx.doi.org/10.3150/13-BEJSP16. doi:10.3150/13-bejsp16.

[9] G. Pang, C. Shen, L. Cao, A. V. D. Hengel, Deep Learning for Anomaly Detection: A Review, ACM Computing Surveys 54 (2021) 1–38. URL: http://dx.doi.org/10.1145/3439950. doi:10.1145/3439950.

[10] M. A. Bansal, D. R. Sharma, D. M. Kathuria, A Systematic Review on Data Scarcity Problem in Deep Learning: Solution and Applications, ACM Computing Surveys 54 (2022) 1–29. URL: http://dx.doi.org/10.1145/3502287. doi:10.1145/3502287.

[11] L. Alzubaidi, J. Bai, A. Al-Sabaawi, J. Santamaría, A. S. Albahri, B. S. N. Al-dabbagh, M. A. Fadhel, M. Manoufali, J. Zhang, A. H. Al-Timemy, Y. Duan, A. Abdullah, L. Farhan, Y. Lu, A. Gupta, F. Albu, A. Abbosh, Y. Gu, A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications, Journal of Big Data 10 (2023). URL: http://dx.doi.org/10.1186/s40537-023-00727-2. doi:10.1186/s40537-023-00727-2.

[12] M. Kaden, S. Saralajew, T. Villmann, Domain Knowledge Integration in Machine Learning Systems: An Introduction, in: ESANN 2024 Proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, i6doc.com, Bruges (Belgium) and online, 2024, pp. 405–412. URL: https://www.esann.org/sites/default/files/proceedings/2024/ES2024-5.pdf, accessed: 29.05.2025.

[13] L. von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, M. Walczak, J. Pfrommer, A. Pick, R. Ramamurthy, J. Garcke, C. Bauckhage, J. Schuecker, Informed Machine Learning - A Taxonomy and Survey of Integrating Prior Knowledge into Learning Systems, IEEE Transactions on Knowledge and Data Engineering (2021) 1–1. URL: http://dx.doi.org/10.1109/TKDE.2021.3079836. doi:10.1109/tkde.2021.3079836.

[14] M. Raissi, P. Perdikaris, G. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, Journal of Computational Physics 378 (2019) 686–707. URL: http://dx.doi.org/10.1016/j.jcp.2018.10.045. doi:10.1016/j.jcp.2018.10.045.

[15] A. Cicirello, Physics-Enhanced Machine Learning: a position paper for dynamical systems investigations, Journal of Physics: Conference Series 2909 (2024) 012034. URL: http://dx.doi.org/10.1088/1742-6596/2909/1/012034. doi:10.1088/1742-6596/2909/1/012034.

[16] S. A. Faroughi, N. Pawar, C. Fernandes, M. Raissi, S. Das, N. K. Kalantari, S. K. Mahjour, Physics-Guided, Physics-Informed, and Physics-Encoded Neural Networks in Scientific Computing, 2022. URL: https://arxiv.org/abs/2211.07377. doi:10.48550/ARXIV.2211.07377.

[17] K. L. Tsui, N. Chen, Q. Zhou, Y. Hai, W. Wang, Prognostics and Health Management: A Review on Data Driven Approaches, Mathematical Problems in Engineering 2015 (2015) 1–17. URL: http://dx.doi.org/10.1155/2015/793161. doi:10.1155/2015/793161.

[18] T. Hiruta, T. Uchida, S. Yuda, Y. Umeda, A design method of data analytics process for condition based maintenance, CIRP Annals 68 (2019) 145–148. URL: http://dx.doi.org/10.1016/j.cirp.2019.04.049. doi:10.1016/j.cirp.2019.04.049.

[19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, I. Polosukhin, Attention is All you Need, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 30, Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

[20] M. Hüsken, P. Stagge, Recurrent neural networks for time series classification, Neurocomputing 50 (2003) 223–235. URL: http://dx.doi.org/10.1016/S0925-2312(01)00706-8. doi:10.1016/s0925-2312(01)00706-8.

[21] J. Connor, R. Martin, L. Atlas, Recurrent neural networks and robust time series prediction, IEEE Transactions on Neural Networks 5 (1994) 240–254. doi:10.1109/72.279188.

[22] S. Hochreiter, J. Schmidhuber, Long Short-Term Memory, Neural Computation 9 (1997) 1735–1780. doi:10.1162/neco.1997.9.8.1735.

[23] F. A. Gers, D. Eck, J. Schmidhuber, Applying LSTM to Time Series Predictable through Time-Window Approaches, Springer Berlin Heidelberg, 2001, p. 669–676. URL: http://dx.doi.org/10.1007/3-540-44668-0_93. doi:10.1007/3-540-44668-0_93.

[24] Z. Wang, J. Lin, Z. Wang, Accelerating Recurrent Neural Networks: A Memory-Efficient Approach, IEEE Transactions on Very Large Scale Integration (VLSI) Systems 25 (2017) 2763–2775. doi:10.1109/TVLSI.2017.2717950.

[25] A. Pulver, S. Lyu, LSTM with working memory, in: 2017 International Joint Conference on Neural Networks (IJCNN), 2017, pp. 845–851. doi:10.1109/IJCNN.2017.7965940.

[26] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, W. Zhang, Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting, Proceedings of the AAAI Conference on Artificial Intelligence 35 (2021) 11106–11115. URL: http://dx.doi.org/10.1609/aaai.v35i12.17325. doi:10.1609/aaai.v35i12.17325.

[27] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, R. Jin, FEDformer: Frequency Enhanced Decomposed Transformer for Long-term Series Forecasting, in: K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, S. Sabato (Eds.), Proceedings of the 39th International Conference on Machine Learning, volume 162 of *Proceedings of Machine Learning Research*, PMLR, 2022, pp. 27268–27286. URL: https://proceedings.mlr.press/v162/zhou22g.html.

[28] H. Wu, J. Xu, J. Wang, M. Long, Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting, in: M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, J. W. Vaughan (Eds.), Advances in Neural Information Processing Systems, volume 34, Curran Associates, Inc., 2021, pp. 22419–22430. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/bcc0d400288793e8bdcd7c19a8ac0c2b-Paper.pdf.

[29] P. Jiang, C. Xiao, A. Cross, J. Sun, GraphCare: Enhancing Healthcare Predictions with Personalized Knowledge Graphs, 2024. URL: https://arxiv.org/abs/2305.12788. arXiv:2305.12788.

[30] S. Hao, J. Bao, C. Lu, A Time Series Multitask Framework Integrating a Large Language Model, Pre-Trained Time Series Model, and Knowledge Graph, 2025. URL: https://arxiv.org/abs/2503.07682. arXiv:2503.07682.

[31] Z. Zhao, P. Wang, H. Wen, S. Wang, L. Yu, Y. Wang, STEM-LTS: Integrating Semantic-Temporal Dynamics in LLM-driven Time Series Analysis, Proceedings of the AAAI Conference on Artificial Intelligence 39 (2025) 22858–22866. URL: http://dx.doi.org/10.1609/aaai.v39i21.34447. doi:10.1609/aaai.v39i21.34447.

[32] P. Liu, H. Guo, T. Dai, N. Li, J. Bao, X. Ren, Y. Jiang, S.-T. Xia, CALF: Aligning LLMs for Time Series Forecasting via Cross-modal Fine-Tuning, 2025. URL: https://arxiv.org/abs/2403.07300. arXiv:2403.07300.
[33] T. Zhao, X. Chen, M. Sun, Enhancing Time Series Forecasting via Multi-Level Text Alignment with LLMs, 2025. URL: https://arxiv.org/abs/2504.07360. arXiv:2504.07360.
[34] D. Edge, H. Trinh, N. Cheng, J. Bradley, A. Chao, A. Mody, S. Truitt, D. Metropolitansky, R. O. Ness, J. Larson, From Local to Global: A Graph RAG Approach to Query-Focused Summarization, 2025. URL: https://arxiv.org/abs/2404.16130. arXiv:2404.16130.
[35] P. Ivanovic, S. Burbach, M. Maleshkova, MontoFlow/SHIP, https://doi.org/10.5281/zenodo.15390282, 2024. doi:10.5281/zenodo.15390282.
[36] A. Haller, K. Janowicz, S. Cox, D. L. Phuoc, K. Taylor, M. Lefrançois, Semantic Sensor Network Ontology, https://www.w3.org/TR/vocab-ssn/, 2017. W3C Recommendation, World Wide Web Consortium (W3C), accessed: 02.05.2025.
[37] K. Taylor, A. Haller, M. Lefrançois, S. Cox, K. Janowicz, R. García-Castro, D. Le-Phuoc, J. Lieberman, R. Atkinson, C. Stadler, The Semantic Sensor Network Ontology, Revamped, in: Proceedings of the Journal Track co-located with the 18th International Semantic Web Conference (ISWC 2019), volume 2576 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019, pp. 1–9. URL: https://ceur-ws.org/Vol-2576/paper11.pdf.
[38] T. Wang, B. Shen, OTMKGRL: a universal multimodal knowledge graph representation learning framework using optimal transport and cross-modal relation, Applied Intelligence 55 (2025). URL: http://dx.doi.org/10.1007/s10489-025-06459-5. doi:10.1007/s10489-025-06459-5.
[39] P. Ristoski, H. Paulheim, RDF2Vec: RDF Graph Embeddings for Data Mining, Springer International Publishing, 2016, p. 498–514. URL: http://dx.doi.org/10.1007/978-3-319-46523-4_30. doi:10.1007/978-3-319-46523-4_30.
[40] J. Chen, P. Hu, E. Jimenez-Ruiz, O. M. Holter, D. Antonyrajah, I. Horrocks, OWL2Vec*: embedding of OWL ontologies, Machine Learning (2021). URL: http://dx.doi.org/10.1007/s10994-021-05997-6. doi:10.1007/s10994-021-05997-6.
[41] J. Hao, M. Chen, W. Yu, Y. Sun, W. Wang, Universal Representation Learning of Knowledge Bases by Jointly Embedding Instances and Ontological Concepts, 2021. URL: https://arxiv.org/abs/2103.08115. arXiv:2103.08115.
[42] S. L. Brunton, J. L. Proctor, J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, Proceedings of the National Academy of Sciences 113 (2016) 3932–3937. URL: http://dx.doi.org/10.1073/pnas.1517384113. doi:10.1073/pnas.1517384113.
[43] M. Cranmer, Interpretable Machine Learning for Science with PySR and SymbolicRegression.jl, 2023. URL: https://arxiv.org/abs/2305.01582. arXiv:2305.01582.
[44] M. Á. de Carvalho Servia, I. O. Sandoval, K. K. M. Hii, K. Hellgardt, D. Zhang, E. Antonio del Rio Chanona, The automated discovery of kinetic rate models – methodological frameworks, Digital Discovery 3 (2024) 954–968. URL: http://dx.doi.org/10.1039/D3DD00212H. doi:10.1039/d3dd00212h.
[45] C. J. Soelistyo, A. R. Lowe, Discovering interpretable models of scientific image data with deep learning, 2024. URL: https://arxiv.org/abs/2402.03115. doi:10.48550/ARXIV.2402.03115.
[46] C. Lathourakis, A. Cicirello, Physics enhanced sparse identification of dynamical systems with discontinuous nonlinearities, Nonlinear Dynamics 112 (2024) 11237–11264. URL: https://doi.org/10.1007/s11071-024-09652-2. doi:10.1007/s11071-024-09652-2.
[47] T. M. Mitchell, The Need for Biases in Learning Generalizations, Technical Report CBM-TR-117, Department of Computer Science, Rutgers University, New Brunswick, NJ, 1980. URL: https://www.cs.cmu.edu/~tom/pubs/NeedForBias_1980.pdf.
[48] M. Haywood-Alexander, W. Liu, K. Bacsa, Z. Lai, E. Chatzi, Discussing the Spectrum of Physics-Enhanced Machine Learning; a Survey on Structural Mechanics Applications, 2023. URL: https://arxiv.org/abs/2310.20425. doi:10.48550/arXiv.2310.20425. arXiv:2310.20425, arXiv preprint.
[49] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nature Reviews Physics 3 (2021) 422–440. URL: https://doi.org/10.1038/

s42254-021-00314-5. doi:10.1038/s42254-021-00314-5.

[50] A. Fan, J. Yang, L. Yang, D. Wu, N. Vladimir, A review of ship fuel consumption models, Ocean Engineering 264 (2022) 112405. URL: http://dx.doi.org/10.1016/j.oceaneng.2022.112405. doi:10.1016/j.oceaneng.2022.112405.

[51] R. N. Bergman, Origins and History of the Minimal Model of Glucose Regulation, Frontiers in Endocrinology 11 (2021). URL: http://dx.doi.org/10.3389/fendo.2020.583016. doi:10.3389/fendo.2020.583016.

[52] D. H. Hagos, R. Battle, D. B. Rawat, Recent advances in generative ai and large language models: Current status, challenges, and perspectives, 2024. URL: https://arxiv.org/abs/2407.14962. doi:10.48550/ARXIV.2407.14962.

[53] S. Minaee, T. Mikolov, N. Nikzad, M. Chenaghlu, R. Socher, X. Amatriain, J. Gao, Large language models: A survey, 2024. URL: https://arxiv.org/abs/2402.06196. doi:10.48550/ARXIV.2402.06196.