DOI: 10.34749/3061-1466.2026.47

Conceptual-Model-Guided Physics-Inspired Feature Engineering: A 3D-Printer Case Study

Martin Paczona
Greenformance
Feldbach, Austria
martin@greenformance.com

Abstract

Conceptual-model-guided physics-inspired feature engineering (CPFE) a lightweight representation-level strategy that formalizes the incorporation of domain knowledge into machine-learning pipelines. Instead of embedding physics in the loss or model architecture, the proposed approach structures physical and operational knowledge at the representation level through traceable domain concepts. The method is demonstrated in a 3D-printer case study, where the resulting physics representation achieves better performance than a compact baseline and automatically extracted features using tsfresh.

1 Introduction

Physics-informed and knowledge-guided machine learning integrate prior domain structure in multiple ways, including feature engineering and data preprocessing, training objectives and constraints, and architecture-level design [1–6]. Within this landscape, representation-level integration offers a lightweight and practically attractive route to hybrid modeling, particularly in power-only monitoring settings characterized by limited observability, low-cost sensing, and edge deployment. This paper presents conceptual-model-guided physics-inspired feature engineering (CPFE) for machine-state estimation from active-power measurements. Starting from a conceptual model of the monitored device and its operation, CPFE formalizes the derivation of interpretable feature concepts. This paper makes three contributions: a conceptual meta-model for organizing traceable domain concepts, a CPFE pipeline for deriving proxy features from power measurements, and an empirical case study with a 3D-printer.

2 Related Work

Physics-informed and knowledge-guided machine learning integrates prior domain structure into data-driven models to improve robustness, interpretability, and generalization [1]. Recent literature distinguishes several pathways for incorporating such knowledge, including theory-inspired feature engineering and data preprocessing, model selection, regularization, and architecture-level design [2–6]. In the energy domain, non-intrusive load monitoring (NILM) studies inference from electrical measurements, often under limited observability and overlapping signatures [7–9]. These challenges are particularly relevant in power-only settings, where only low-frequency active-power signals are available. Standardized datasets and toolkits such as UK-DALE and NILMTK support reproducible evaluation [10, 11], while recent industrial reviews emphasize the additional complexity of machine-level industrial scenarios [12]. Recent work has also begun to incorporate physics-informed principles into NILM during model learning [13]. More broadly, representation choice strongly affects learnability and inductive bias [14]. Automated feature engineering explores large transformation spaces to construct useful representations [15, 16], whereas theory-guided data science advocates

The Third Austrian Symposium on AI, Robotics, and Vision (AIROV26).

the controlled incorporation of prior knowledge into data-driven pipelines [4, 6, 17]. Building on this line of research, the present work formalizes a structured and traceable derivation of feature concepts from an explicit conceptual model of the device, thereby providing a lightweight and interpretable representation-level strategy for machine-state estimation from low-frequency active-power measurements.

3 Conceptual-Model-Guided Physics-Inspired Feature Engineering (CPFE)

Under power-only sensing, many relevant internal variables of the observed device remain latent, and different physical mechanisms may produce similar load signatures. Purely statistical feature generation is therefore often difficult to interpret and may fail to capture task-relevant structure. To address this challenge, we present CPFE, where feature generation is grounded in an explicit conceptualization of the observed system, the available raw data, and the intended prediction task. The resulting representation is treated as a purpose-driven abstraction that must satisfy two requirements: *ontological clarity*, i.e., an explicit definition of the relevant concepts and their relations, and *operationalizability*, i.e., a traceable mapping from those concepts to computable descriptors. To make this derivation explicit, CPFE is organized through the meta-model shown in Fig. 1. The meta-model distinguishes between the *system side* and the *representation side*. On the system side, a *SystemType* defines the class of device under consideration, while a *RawDataStream* represents the measured signal through which the system is observed. On the representation side, the central modeling element is the *DomainConcept*. A domain concept captures one interpretable aspect of the system that is relevant for the learning task and is linked to four elements: it *applies to* a given system type, *consumes* raw data, *has* parameters, and *produces* concrete features. In addition, each domain concept is explicitly *grounded in* physical principles, *restricted by* constraints, and *characterized by* operational concepts. This creates a transparent abstraction chain from the observed system and its measurements to implemented feature families. The CPFE procedure can be summarized in five steps. First, the prediction task and its temporal scope are specified, since the relevant abstraction level depends on whether the objective is, for example, state classification, forecasting, or anomaly detection. Second, the dominant physical principles, system constraints, and operational concepts that shape the measured power signal are identified. Third, these considerations are translated into a set of domain concepts, each representing one meaningful aspect of system behavior that is expected to be informative for the task. Fourth, each domain concept is parameterized and operationalized into one or more proxy features computable from the raw data stream and, where required, from available system information such as technical specifications. Fifth, the resulting feature pipeline is evaluated with respect to the target prediction task using explicit criteria such as classification or forecasting performance, and refined through comparison and ablation. Physical principles, constraints, and operational assumptions are encoded a priori in the feature space through structured domain concepts and their associated proxy variables. The machine-learning model then learns the remaining task-specific mapping on top of this concept-guided representation. This separates system understanding from statistical learning.

4 Case Study: 3D-Printer State Estimation

The proposed CPFE methodology is demonstrated on a filament 3D-printer using active-power measurements at 10 s resolution for state estimation. The dataset used is available at `https://doi.org/10.34740/kaggle/dsv/15995622`. Training labels were obtained from the printer's G-code interface. Following the meta-model introduced in the previous section, the 3D-printer case study is instantiated through an ordered set of domain concepts. Each domain concept defines one interpretable aspect of the system, together with its corresponding feature outputs, parameterization, and traceability links to physical principles, constraints, and operational concepts. In the present implementation, nine domain concepts are represented by the twenty-six features, one example feature per domain concept is shown in Table 1. Thermal behavior is represented through accumulation, slow-fast envelope separation, baseline-relative heating proxies, and a first-order RC-inspired latent thermal state. Transition behavior is represented through slope-based measures and transition density. Motion-related behavior is represented through residual high-frequency variability, a dedicated motion surrogate, and an additional low-load actuation concept intended to isolate actuation-dominated transients during quiescent thermal phases. The representation constructs a set of observable proxy variables that make the effects of latent mechanisms more distinguishable in the measured power

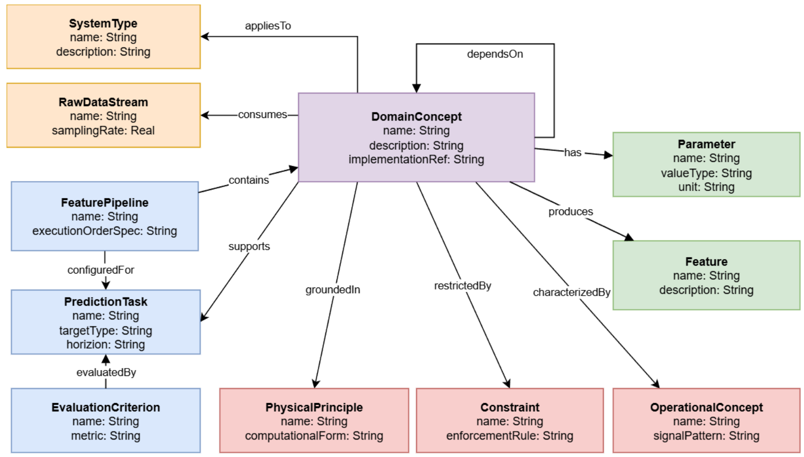

Figure 1: Meta-model of the proposed feature conceptualization method.

Table 1: Implemented domain concepts and example features.

Domain concept	Physical principle	Feature example
Thermal load accumulation	thermal	`energy_window`: accumulated recent power over a short time window; increases during sustained heating
Multi-timescale envelope	power dynamics	`ma_gap`: difference between short- and long-horizon moving-average levels; indicates whether load is currently above or below the longer trend
Regime-transition intensity	transition	`transition_density`: fraction of recent samples with unusually strong power changes; highlights switching and transition-heavy periods
Operating variability	power dynamics	`volatility_ratio`: ratio of short-term to long-term variability; higher values indicate more irregular operation
Thermal memory deviation	thermal	`ewma_cross`: difference between fast and slow exponentially weighted averages; captures short-term departure from the thermal trend
Baseline-to-active decomp.	thermal	`delta_p`: power above the estimated idle baseline; isolates active thermal load from background consumption
Motion surrogate	motion	`motion_like`: composite proxy for non-thermal actuation based on short-term variability, derivative magnitude, and residual fluctuations
RC thermal state	thermal	`stored_heat_proxy_rc`: recursively filtered proxy for retained heat; represents thermal state with gradual decay over time
Low-load actuation	motion	`low_load_actuation_energy`: transient actuation activity accumulated specifically during periods of low thermal load

signal. The implementation operationalizes these concepts, e.g., through derivative-based operators and recursive RC models applied to the active-power time series, as illustrated by Eq. (1):

$$\text{stored\_heat\_proxy\_rc}_k = \alpha_{\mathrm{RC}}\, x_{k-1} + (1-\alpha_{\mathrm{RC}})\,\Delta P_k, \quad \alpha_{\mathrm{RC}} = e^{-\Delta t/\tau_{\mathrm{RC}}} \tag{1}$$

Here, α_{RC} is the RC decay factor, ΔP_k is the power deviation at sample k, Δt the sampling interval, and τ_{RC} the thermal time constant. The resulting representation remains interpretable because each implemented feature can be traced back to a domain concept and to the underlying physical principles.

5 Results

This section evaluates whether the proposed method improves machine-state estimation and whether the resulting performance gains can be interpreted in terms of the underlying domain concepts. The data were partitioned once into chronological 60/20/20 train/validation/test splits, with purge gaps of 35 samples between partitions. The validation split was reserved for exploratory model-development outside the main comparison in this paper. To isolate the effect of the feature representation, the comparative results reported in Fig. 2 use one pre-specified Random Forest hyperparameter configuration for all three feature sets (`max_depth=3`, `min_samples_leaf=20`, `random_state=42`), and performance is reported on the held-out test set using accuracy, macro-F1, and balanced accuracy. The baseline features consisted of generic power features: rolling mean, rolling standard deviation, first-order difference, rolling maximum and raw power. Examples of `tsfresh`-derived features from the power signal include the `minimum`, the 0.1-`quantile`, `lempel_ziv_complexity` with 100 bins,

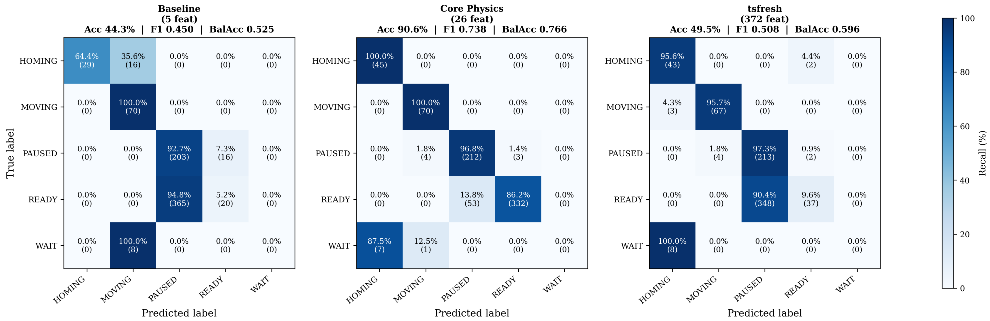

Figure 2: Confusion matrices for the compared feature sets, using identical RF hyperparameters.

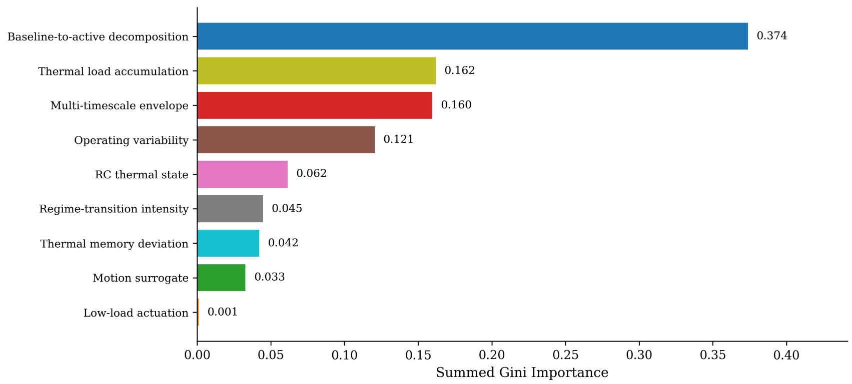

Figure 3: Aggregated feature importance by domain concept for the proposed approach.

and `binned_entropy` with 10 bins. Figure 2 compares the three feature representations. In this test the proposed core-physics representation achieves the best overall performance, reaching 90.6% accuracy, a macro-averaged F1 score of 0.738, and a balanced accuracy of 0.766. The persistent misclassification of the rare `WAIT` class is consistent with its very short duration and the lack of a distinct power signature in 10 s power-only data. While low prevalence may contribute, the main limitation is likely weak class separability rather than imbalance alone. Figure 3 shows the aggregated Gini feature importance by the nine domain concepts for the core-physics Random Forest model. The dominant contribution comes from *baseline-to-active decomposition*, followed by *thermal load accumulation*, *multi-timescale envelope*, and *operating variability*. This ranking is physically plausible for the considered system. The remaining concepts play a supportive but still interpretable role. The results support the main hypothesis of this work: CPFE can produce a compact and interpretable representation that outperforms both a simple baseline and a large automatic feature set while remaining more interpretable. These findings should be interpreted as evidence from a single-device case study focused on representation quality rather than broad cross-device generalization.

6 Conclusion and Outlook

This work shows that conceptual-model-guided feature-space hybridization can improve machine-state estimation from power-only measurements while remaining lightweight and interpretable. In this 3D-printer case study, the proposed core-physics representation outperformed both a compact baseline and automatic feature extraction. Overall, the findings suggest that conceptual-model-guided and traceable feature construction offers a practical complement to physics-informed learning in the loss or architecture, particularly in limited-observability settings. Future work will address transferability across devices, evaluation of the proposed feature conceptualization across different model architectures, automated concept-to-feature generation, and integration with end-to-end physics-inspired machine learning.

7 Acknowledgments and Disclosure of Funding

This work was carried out as part of the GreenformanceOptimizer project of the startup Greenformance. No external funding was received.

References

[1] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021. doi:10.1038/s42254-021-00314-5.

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

[3] J. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4):66:1–37, 2022. doi:10.1145/3514228.

[4] 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, and J. Schuecker. Informed machine learning—A taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):614–633, 2023. doi:10.1109/TKDE.2021.3079836.

[5] T. Dash, S. Chitlangia, A. Ahuja, and A. Srinivasan. A review of some techniques for inclusion of domain-knowledge into deep neural networks. *Scientific Reports*, 12:1040, 2022. doi:10.1038/s41598-021-04590-0.

[6] J. G. Hoffer, A. B. Ofner, F. M. Rohrhofer, M. Lovrić, R. Kern, S. Lindstaedt, and B. C. Geiger. Theory-inspired machine learning—towards a synergy between knowledge and data. *Welding in the World*, 66:1291–1304, 2022. doi:10.1007/s40194-022-01270-z.

[7] G. W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992. doi:10.1109/5.192069.

[8] M. Zeifman and K. Roth. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics*, 57(1):76–84, 2011. doi:10.1109/TCE.2011.5735484.

[9] A. Zoha, A. Gluhak, M. A. Imran, and S. Rajasegarar. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors*, 12(12):16838–16866, 2012. doi:10.3390/s121216838.

[10] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava. NILMTK: An open source toolkit for non-intrusive load monitoring. In *Proceedings of the 5th International Conference on Future Energy Systems (ACM e-Energy '14)*, pages 265–276. ACM, 2014. doi:10.1145/2602044.2602051.

[11] J. Kelly and W. Knottenbelt. The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. *Scientific Data*, 2:150007, 2015. doi:10.1038/sdata.2015.7.

[12] G. Tanoni, E. Principi, and S. Squartini. Non-intrusive load monitoring in industrial settings: A systematic review. *Renewable and Sustainable Energy Reviews*, 202:114703, 2024. doi:10.1016/j.rser.2024.114703.

[13] G. Huang, Z. Zhou, F. Wu, and W. Hua. Physics-informed time-aware neural networks for industrial nonintrusive load monitoring. *IEEE Transactions on Industrial Informatics*, 19(6):7312–7322, 2023. doi:10.1109/TII.2022.3211075.

[14] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. doi:10.1109/TPAMI.2013.50.

[15] X. He, K. Zhao, and X. Chu. AutoML: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212:106622, 2021. doi:10.1016/j.knosys.2020.106622.

[16] J. M. Kanter and K. Veeramachaneni. Deep feature synthesis: Towards automating data science endeavors. In *Proceedings of the 2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10. IEEE, 2015. doi:10.1109/DSAA.2015.7344858.

[17] A. Karpatne, G. Atluri, J. H. Faghmous, M. Steinbach, A. Banerjee, A. Ganguly, S. Shekhar, N. Samatova, and V. Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331, 2017. doi:10.1109/TKDE.2017.2720168.