GALACTICA: An Interactive System for Local and Global Counterfactual Explanations for Time-series Clustering

Christos Fragkathoulas
University of Ioannina
Archimedes, Athena Research Center
ch.fragkathoulas@athenarc.gr

Eleni Psaroudaki
Archimedes, Athena Research Center
h.psaroudaki@athenarc.gr

Themis Palpanas
Université Paris Cité, LIPADE, F-75006
Archimedes, Athena Research Center
themis@mi.parisdescartes.fr

Evaggelia Pitoura
University of Ioannina
Archimedes, Athena Research Center
pitoura@uoi.gr

ABSTRACT

Time-series clustering is widely used in practice, yet the resulting clusters and assignments are often hard to interpret. Moreover, there is a lack of tools that explain both how a specific series can move across cluster boundaries and how such transitions can be summarized at the cluster level. In this paper, we present GALACTICA, an interactive visual system for exploring local and global counterfactual explanations in time-series clustering. The system is built on GALACTIC, a unified framework that generates minimal, structurally meaningful perturbations for individual series while also producing compact summaries of cluster-to-cluster transitions using the Minimum Description Length (MDL) principle. GALACTICA allows users to inspect the structure of a clustered time-series collection, analyze internal cluster organization, explore instance-level counterfactual explanations, examine global summaries of cluster transitions, and compare GALACTIC against state-of-the-art baselines. By coupling the algorithmic internals of GALACTIC with an interactive visual interface, GALACTICA turns abstract cluster transitions into interpretable and auditable explanation workflows.

PVLDB Reference Format:
Christos Fragkathoulas, Eleni Psaroudaki, Themis Palpanas, and Evaggelia Pitoura. GALACTICA: An Interactive System for Local and Global Counterfactual Explanations for Time-series Clustering. PVLDB, 19(12): 4774 - 4777, 2026.
doi:10.14778/3827998.3828119

PVLDB Artifact Availability:
The source code, data, and/or other artifacts have been made available at https://github.com/xristosfrag/GALACTICA.

1 INTRODUCTION

Time-series clustering is a fundamental tool for exploratory analysis of temporal signals. Such clustering pipelines are widely used in the analysis of temporal data across many domains [12]. However, once the clusters have been formed, understanding why a series belongs to a given cluster and *how* it could transition to another remains difficult. This limitation is particularly important in domains where clustering is used not only for summarization, but also for downstream interpretation and decision support.

Time-series clustering pipelines are often difficult to interpret in practice due to proprietary components, complex non-linear distance measures (e.g., DTW), and multi-stage feature transformations that obscure the assignment logic [5, 14]. Existing explainability approaches generally fall into two categories: inherently interpretable models and post-hoc surrogates. The former integrates interpretability directly into specific model architectures [e.g., 4, 6], while the latter relies on surrogate models or static summaries [e.g., 3, 11], which rarely bridge instance-level reasoning with concise global cluster insights. Counterfactual explanations (CEs), which identify minimal input changes that alter a prediction [15], offer a promising alternative. However, existing CE methods for time-series classification [e.g., 1, 2, 9, 10] often disrupt temporal coherence, overwrite characteristic patterns, or fail to focus on truly discriminative regions [2, 9, 10]. More critically, methods designed for supervised labels do not address the unsupervised setting of discovering transitions between latent time-series clusters, nor the challenge of summarizing large numbers of local counterfactuals into compact global explanations.

GALACTIC was recently introduced to bridge this gap by providing a unified framework for both local and global counterfactual explanations in time-series clustering [8], agnostic to the underlying clustering procedure. At the local level, it computes minimal perturbations that move an individual series across cluster boundaries while respecting structurally important temporal regions. At the global level, it selects representative perturbations that summarize the dominant transition patterns across the entire cluster using a Minimum Description Length (MDL)-based objective [13].

In this paper, we present GALACTICA, an interactive system that operationalizes GALACTIC. GALACTICA enables users to inspect the structure of clustered time-series collections, explore local counterfactual transitions for individual series, analyze global summaries of cluster transitions, and compare GALACTIC with alternative counterfactual generation methods through an interactive visual interface. By combining algorithmic counterfactual generation of GALACTIC with visual analytics, GALACTICA supports human-in-the-loop exploration and interpretation of large time-series clustered collections.

Proceedings of the VLDB Endowment, Vol. 19, No. 12 ISSN 2150-8097.
doi:10.14778/3827998.3828119

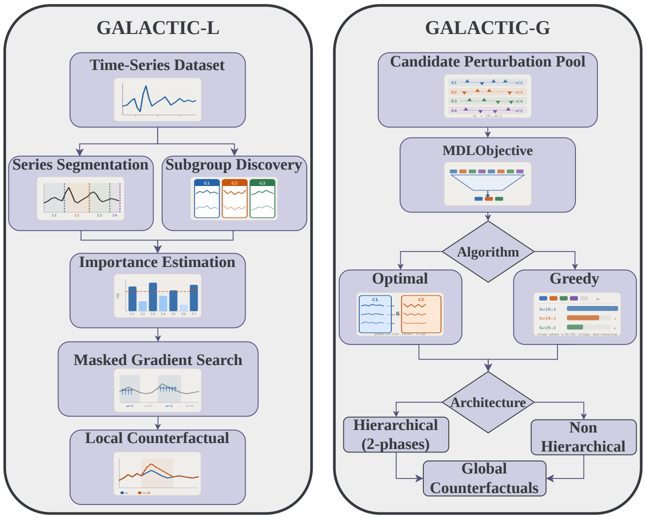

Figure 1: GALACTIC Suite

2 PROPOSED APPROACH: GALACTIC

We now describe GALACTIC [8], the algorithmic core of our system for generating local and global counterfactual (CF) explanations in clustered time-series data. As shown in Figure 1, GALACTIC is organized into two complementary components: GALACTIC-L generates counterfactual explanations for individual time series, and GALACTIC-G summarizes cluster-level transition patterns through a compact set of representative perturbations.

[Local Counterfactual Explanations] GALACTIC-L explains how a specific time series can move from its current cluster to another through a small, meaningful modification. As shown in Figure 1, the local pipeline first analyzes the dataset's temporal structure via series segmentation and subgroup discovery: each series is partitioned into contiguous segments, and the resulting gap vectors are clustered via K_{seg}-medoids to identify representative structural subgroups within each cluster. It then estimates each segment importance via a block-wise permutation test, measuring the drop in model accuracy when values of a segment are shuffled across the cluster population, and propagates resulting scores into a binary mask that highlights the most discriminative regions. This importance mask guides a constrained gradient search that restricts edits to the discriminative segments, producing counterfactuals that preserve the overall shape of the series while altering only regions responsible for the cluster transition.

[Global Counterfactual Explanations] While local counterfactuals explain individual series, users often seek higher-level explanations describing how entire clusters transition to others without generating multiple counterfactuals. GALACTIC-G addresses this need by selecting a small set of representative perturbations that summarize the dominant transition patterns of a cluster. As shown in Figure 1, the global pipeline begins with a candidate perturbation pool from GALACTIC-L or other local CF baselines, and selects the subset that best explains how cluster instances can transition to other clusters via the Minimum Description Length (MDL) principle, balancing explanation coverage and solution complexity. In simplified form, the objective can be expressed as: $M^\star = \arg\min_M L(D,M), \quad L(D,M) = L(M) + L(D \mid M)$, where D denotes the set of time series in the cluster and M the selected set of representative perturbations. Here, $L(M)$ captures the complexity of the explanation set, while $L(D \mid M)$ measures how many instances remain unexplained by the selected perturbations. The framework supports four variants: *Optimal*, which exhaustively finds the best MDL-minimizing subset, *Greedy*, which selects perturbations with the largest MDL gain for scalability, and their *Hierarchical* extensions, which operate within subgroups before the cluster level, improving scalability for high-variability clusters.

Together, GALACTIC-L and GALACTIC-G provide a unified framework for interpreting time-series clustering at multiple levels, from individual instance changes to cluster-level transition summaries.

3 GALACTICA: SYSTEM OVERVIEW

In this section, we describe GALACTICA, an interactive system designed to generate and visualize counterfactual explanations for time-series clustering via the GALACTIC algorithmic suite. The application provides a transparent visual interface for auditing the internals of the framework, enabling users to investigate the structural logic and decision boundaries that drive cluster transitions. By visualizing counterfactual explanations, GALACTICA transforms abstract perturbations into interpretable changes, allowing the evaluation of local and global explanations. GALACTICA is built as a standalone, easily extensible, web application using Python 3.13 and the Streamlit framework. The datasets are from the UCR Time Series Archive [7], comprising univariate, fixed-length time series of discrete snapshots, spanning diverse domains such as ECG, image contours, and sensor signals, with varying series counts, sequence lengths, and cluster cardinalities. The backend implements a library of counterfactual generation algorithms. The front end allows users to interact with the system. As shown in Figure 2, it follows a dual-pane layout: a persistent sidebar on the left that facilitates dataset selection and navigation, and a primary workspace which dynamically renders one of four specialized frames: *Clustering*, *Local*, *Global Counterfactual Exploration*, and *Benchmark* views.

[*Frame 1: Clustering Analytics*] This view is organized in three components: the Dataset Statistics, Cluster, and Cluster Internals Panels. The Statistics Panel presents a high-level summary of the loaded dataset, reporting key metadata such as the number of time-series, the sequence length, and cluster cardinality. The Cluster Panel summarizes the clusters by reporting the number of series, coresets (representative instances), and detected structural groups within each cluster. For a more granular inspection, the Cluster Internals Panel decomposes a user-selected cluster into structural subgroups. By rendering these subgroups as collections of time-series, the user uncovers the intra-cluster variance that will guide the counterfactual search process in subsequent steps.

[*Frame 2: Local Counterfactual Exploration*] This view is organized into two functional components: the Configuration Panel and the Results Panel. In the Configuration Panel, the user selects a source cluster to browse and specific instances for counterfactual generation. The Results and Internals Panel renders the original time-series (solid line) with its generated counterfactual (dashed

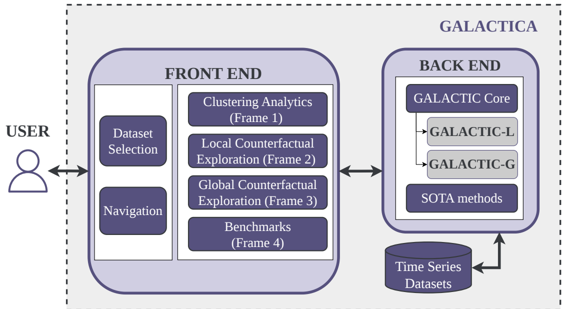

Figure 2: GALACTICA

line). To emphasize that the changes should be targeted, we color-code the temporal segments according to their importance scores. The underlying perturbation is isolated and visualized in a detailed subplot, allowing for finer inspection of the specific modifications. Finally, the Explanations Insights button triggers a detailed pop-up that explains the generation logic, revealing how the framework prioritizes segments for modification based on the cluster.

[*Frame 3: Global Counterfactual Exploration*] This view has three functional modules. The Configuration Panel allows users to specify the cluster of interest and the number of desired perturbations. The Results Panel renders the output for the selected framework variation, i.e., Flat and Hierarchical variants of Optimal and Greedy MDL Selection algorithms. This panel also dynamically adapts to reveal algorithmic internals: for Greedy methods, the progression of the MDL objective; for Hierarchical methods, the multi-stage selection hierarchy. Candidate perturbations are rendered as background trajectories, while selected representatives are highlighted. As the number of time-series in a cluster grows, generating local-level counterfactuals becomes computationally and cognitively intractable. Abstracting individual counterfactuals into a concise global summary lets users view population-wide perturbations without manually inspecting every series.

Finally, the Global-to-Local Panel demonstrates the transition from global logic to local application. Here, the original time-series is depicted as a solid line, while the counterfactuals derived from global (GALACTIC-G) and local (GALACTIC-L) perturbations are represented by dashed and dotted lines, respectively, enabling a direct qualitative comparison of their alignment.

[*Frame 4: Benchmarks*] This view is organized into two functional components: the Configuration Panel and the Benchmarks Panel. In the Configuration Panel, users select a cluster of interest to browse and isolate specific instances for counterfactual generation. The Benchmarks Panel visually overlays the original time-series (solid line) with the counterfactuals (dashed lines) produced by GALACTIC and various state-of-the-art baseline methods, allowing for a visual comparison of the counterfactual explanations. To quantify the quality of the proposed counterfactuals, a comparative metrics table tracks key indicators (i.e., the number of modified time steps, altered segments, perturbation cost, and execution runtime). This summary exposes the inherent trade-offs between competing methods, enabling users to evaluate, for example, the balance between cost-efficiency and temporal sparsity or the relationship between solution quality and real-time responsiveness (runtime).

4 DEMONSTRATION SCENARIOS

In this demo, we present a GUI designed to (i) facilitate the interpretation of time-series clusters; (ii) visualize instance-level perturbations that correspond to local counterfactuals; (iii) compare global-level perturbations across different GALACTIC variations; and (iv) benchmark GALACTIC outputs against state-of-the-art methods. All scenarios start with the user loading a dataset (Figure 3-A).

[*Scenario 1: Understanding the clustered collection*] The user is invited to navigate to the *"Overview"* screen (Figure 3-A), to investigate the underlying cluster structure of the loaded dataset (Figure 3-B1). The interface provides high-level dataset statistics while allowing a deeper dive into the cluster structure via the "Cluster Details" panel. There, the user selects a cluster of interest to explore its decomposition into subgroups. The initial exploration establishes the necessary context for interpreting explanations, ensuring users understand how the dataset instances are organized before moving to counterfactual analysis.

[*Scenario 2: Analyzing Local Counterfactuals*] Upon navigating from the sidebar (Figure 3-A) to the *"Local Counterfactuals"* frame (Figure 3-B2), the user is invited to select a cluster and an instance x within that cluster. The interface visualizes the original series x with the counterfactual explanation x_{cf} generated by GALACTIC-L, and the perturbation $\delta = x_{cf} - x$ that led to x_{cf}. The user is invited to repeat the process for multiple series. To uncover the logic behind the search process, the user is invited to toggle the "Explanations Insights" view (Figure 3-B3), which reveals how the instance relates to its parent subgroup, displaying the temporal segments and importance scores that drove the modification. By connecting the perturbation to these structural features, the system explains why certain regions were altered.

[*Scenario 3: Evaluating Global Counterfactuals*] Upon navigating from the sidebar (Figure 3-A) to the *"Global Counterfactuals"* frame, the user is invited to select a cluster to analyze cluster-level perturbations (Figure 3-B4) and to choose the maximum number of explanations for the cluster of interest. The system automatically generates the MDL-based optimal summary, showing the selected perturbations and their respective coverage across the cluster. The user can toggle between GALACTIC-G algorithm variations to compare solutions. Then, the user is invited to compare the generated solution to another non-MDL one (Figure 3-B6) by pressing the "Visualize non-MDL selected ones" button. Finally, the user is invited to select an individual instance to assess the effectiveness of a global perturbation compared to a custom local one (Figure 3-B4), effectively demonstrating the utility of global explanations in capturing population-wide perturbations. Since scenario 4 follows scenario 3, the user can now see the bottlenecks of local explanations as the number of series increases.

[*Scenario 4: Benchmarking GALACTIC against SOTA methods*] Finally, the demo allows users to position GALACTIC-L explanations against SOTA approaches by choosing the Benchmarks view in the Navigation sidebar (Figure 3-A). The user is invited to select a cluster and a corresponding instance. Comparing GALACTIC with alternative perturbations from baseline methods that do not exploit subgroup structure or segment importance highlights how GALACTIC concentrates edits on meaningful, short temporal regions while maintaining low perturbation cost.

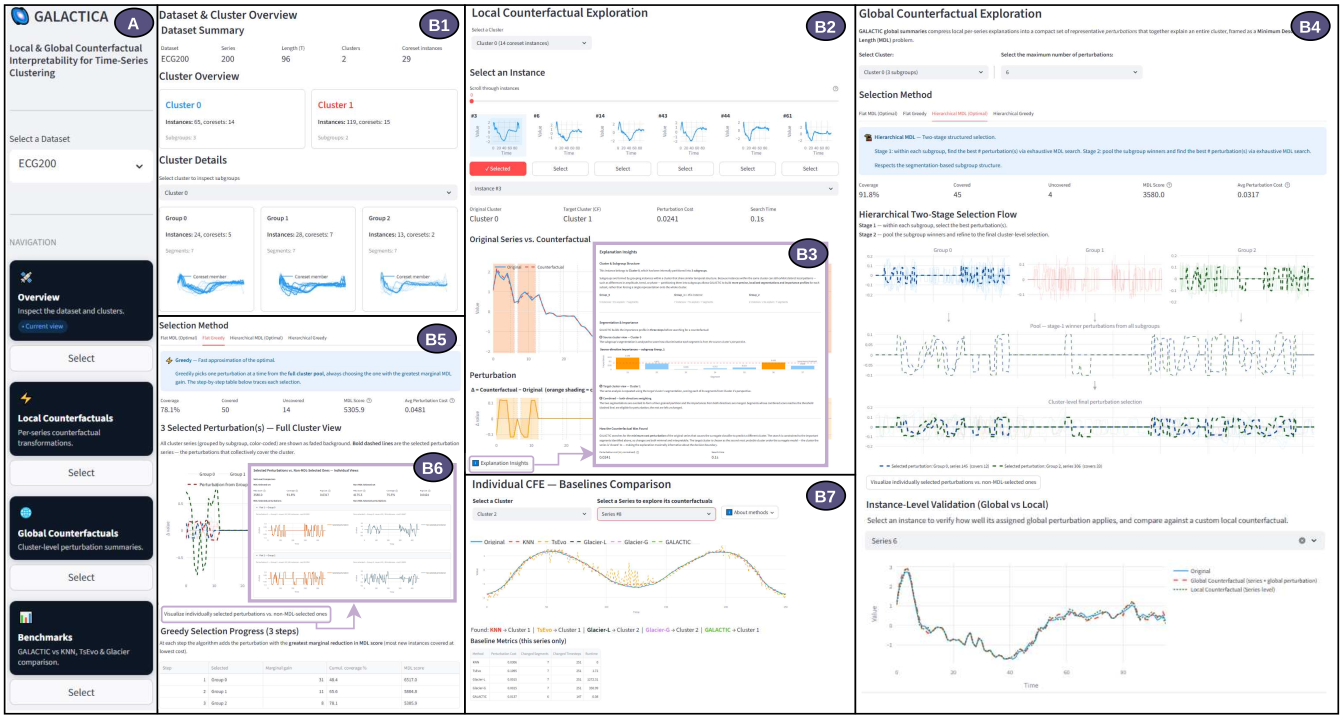

Figure 3: Galactica Scenarios

5 CONCLUSIONS

We demonstrate Galactica, a system designed to bridge the gap between complex counterfactual algorithms and human interpretability in time-series clustering. Galactica allows users to explore both instance-level (Galactic-L) and population-level (Galactic-G) explanations through a cohesive visual interface. Ultimately, Galactica provides an environment for benchmarking the trade-offs between local precision and global generalizability.

ACKNOWLEDGMENTS

This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program, Horizon projects TwinODIS (101160009) and DataGEMS (101188416), *ΥΠΑΙΘΑ* & NextGenerationEU project HARSH (*ΥΠ3ΤΑ* – 0560901) that is carried out within the framework of the National Recovery and Resilience Plan "Greece 2.0" with funding from the European Union – NextGenerationEU.

REFERENCES

[1] Emre Ates, Burak Aksar, Vitus J Leung, and Ayse K Coskun. 2021. Counterfactual explanations for multivariate time series. In *ICAPAI*. IEEE, Halden, Norway, 1–8.

[2] Omar Bahri, Peiyu Li, Soukaina Filali Boubrahimi, and Shah Muhammad Hamdi. 2022. Temporal rule-based counterfactual explanations for multivariate time series. In *ICMLA*. IEEE, Nassau, Bahamas, 1244–1249.

[3] Jürgen Bernard, Tobias Ruppert, Maximilian Scherer, Tobias Schreck, and Jörn Kohlhammer. 2012. Guided discovery of interesting relationships between time series clusters and metadata properties. In *Proceedings of the 12th International Conference on Knowledge Management and Knowledge Technologies* (Graz, Austria). Association for Computing Machinery, New York, NY, USA, 8.

[4] Dimitris Bertsimas, Agni Orfanoudaki, and Holly Wiberg. 2021. Interpretable clustering: an optimization approach. *Machine Learning* 110, 1 (2021), 89–138.

[5] Angela Bonifati, Francesco Del Buono, Francesco Guerra, and Donato Tiano. 2022. Time2Feat: learning interpretable representations for multivariate time series clustering. *Proc. VLDB Endow.* 16, 2 (2022), 193–201.

[6] Paul Boniol, Donato Tiano, Angela Bonifati, and Themis Palpanas. 2025. *k*-Graph: A Graph Embedding for Interpretable Time Series Clustering. *IEEE Transactions on Knowledge and Data Engineering* 37, 5 (2025), 2680–2694.

[7] Hoang Anh Dau, Anthony Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, and Eamonn Keogh. 2019. The UCR time series archive. *IEEE/CAA Journal of Automatica Sinica* 6, 6 (2019), 1293–1305.

[8] Christos Fragkathoulas, Eleni Psaroudaki, Themis Palpanas, and Evaggelia Pitoura. 2026. GALACTIC: Global and Local Agnostic Counterfactuals for Time-series Clustering. arXiv:2603.05318 [cs.LG] https://arxiv.org/abs/2603.05318

[9] Jacqueline Höllig, Cedric Kulbach, and Steffen Thoma. 2022. TSEvo: Evolutionary Counterfactual Explanations for Time Series Classification. In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, Nassau, Bahamas, 29–36.

[10] Jana Lang, Martin A. Giese, Winfried Ilg, and Sebastian Otte. 2023. Generating Sparse Counterfactual Explanations for Multivariate Time Series. In *Artificial Neural Networks and Machine Learning – ICANN 2023: 32nd International Conference on Artificial Neural Networks, Heraklion, Crete, Greece, September 26–29, 2023, Proceedings, Part VI*. Springer-Verlag, Berlin, Heidelberg, 180–193.

[11] Ozan Ozyegen, Nicholas Prayogo, Mucahit Cevik, and Ayse Basar. 2022. Interpretable Time Series Clustering Using Local Explanations. arXiv:2208.01152

[12] John Paparrizos and Sai Prasanna Teja Reddy Bogireddy. 2025. Time-Series Clustering: A Comprehensive Study of Data Mining, Machine Learning, and Deep Learning Methods. *Proc. VLDB Endow.* 18, 11 (2025), 4380–4395.

[13] J. Rissanen. 1978. Paper: Modeling by shortest data description. *Automatica* 14, 5 (1978), 465–471.

[14] Udo Schlegel, Gabriel Marques Tavares, and Thomas Seidl. 2026. Towards Explainable Deep Clustering for Time Series Data. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, Irena Koprinska, João Mendes-Moreira, and Paula Branco (Eds.). Springer Nature Switzerland, Cham, 365–379.

[15] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2018. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology* 31 (2018), 841–887.