

RocketLens: Interactive Visual Analytics for Explaining MiniRocket Time Series Classifiers

Gonzalo Gabriel Méndez^{1,2}, Luis Galárraga² ✉,
Paul Viallard², and Rodne Quijije³

¹ Universitat Politècnica de València, Valencia, Spain, ggmenco1@upv.es

² Inria, Rennes, France, {luis.galarraga, paul.viallard}@inria.fr

³ Escuela Superior Politécnica del Litoral, Guayaquil, Ecuador,
rodanqui@espol.edu.ec

Abstract. MiniRocket enables fast, competitive time-series classification by projecting signals through large banks of 1D convolutional kernels and training classifiers on pooled responses. Its speed makes it attractive for iterative modeling, but its internal logic is difficult to inspect: model behavior emerges from thousands of kernel responses, pooling operations, and linear weights, creating a gap between high accuracy and actionable understanding. We present *RocketLens*, an interactive visualization tool that supports end-to-end inspection of MiniRocket-based models by exploiting feature-attribution explanation, kernel structure and parameters, and convolution-response evidence for selected features.

Keywords: Time series classification · XAI · Feature attribution

1 Introduction

Time series classification underpins decision-making in settings such as cardiac monitoring, industrial condition monitoring, and other sensor-driven pipelines [2]. ROCKET-based methods are a popular approach to address time series classification. They work by transforming time series data into a set of embeddings using a large bank of randomly drawn convolutional kernels. The convolution response is then summarized via a pooling operator before fitting a classifier [1]. MiniRocket refines this idea by reducing randomness through a fixed set of sparse kernels and proportion of positive values (PPV) pooling, while maintaining competitive accuracy and enhancing speed [1].

However, the same design choices that make ROCKET-based classifiers practical also create interpretability friction. A trained classifier operates over thousands of latent features whose values are not meaningful on their own: a practitioner can see that a prediction is confident, or apply post-hoc explainability techniques to conclude that a particular feature is highly attributed, without being able to *(i)* interpret what pattern that feature represents (kernel shape), *(ii)* determine the scale at which it operates (dilation), or *(iii)* locate the evidence on the time axis that drives the pooled feature (threshold exceedance and receptive fields) [1]. This disconnection is especially visible in forensics: when

predictions fail, users need to know whether errors concentrate in specific kernel regimes (e.g., coarse dilations or narrow threshold bands), but typical feature-attribution explainability methods do not surface such structure directly.

Reviews of explainable AI approaches for time series likewise highlight gaps in how well explanations align with a model’s underlying mechanisms [4]. We address this gap with *RocketLens*, a visual analytics tool for MiniRocket-based models that links attributed latent features to their generating kernels (tap pattern, dilation, threshold) and convolution-responses evidence on the original time axis. Beyond single features, *RocketLens* supports exploring which kernel templates and parameter regimes recur in prediction across instances.

2 *RocketLens*: Visual Inspection of MiniRocket Models

*RocketLens*⁴ is a web visual analytics tool for inspecting MiniRocket time-series classifiers. A lightweight Python pipeline precomputes the artifacts required for interactive inspection, including predictions, MiniRocket transform features, kernel parameters, convolution traces, and explainer outputs. The interface is organized around a sequence of interpretive questions. It helps users establish what the model predicted, which latent features support that prediction, what those features compute, and where their evidence appears in the original time series.

Figure 1 uses the abnormal-heartbeat dataset as a running example to describe *RocketLens* and the kinds of inspection it supports. The dataset is a binary classification task that distinguishes normal from abnormal heartbeat recordings. As shown on the top bar of the interface, the selected analytical configuration uses the instance #47 of the dataset, a `RandomForestClassifier`, the `stratoshap-k1` feature-attribution explainer (a stable approximation of SHAP [3]), and the `opposite medoid` reference strategy. We remark that SHAP explains the prediction difference w.r.t. a reference instance, in this case, the medoid from a different class than the one predicted for the instance. The *Reference vs Instance* panel establishes this contrast. In the example, the reference is predicted as *Normal*, while the selected instance is predicted as *Abnormal*. The panel shows how far apart the two cases are in the model output space.

The *Embeddings* view translates this contrast into a feature-level explanation. It presents MiniRocket features as a browsable overview, with bar heights encoding embedding values and an aligned heatmap encoding attribution. Users can sort the view and select features through either representation. In Figure 1, selecting feature #88 updates the linked views, allowing users to trace one positive feature from attribution score to kernel behavior and raw-signal evidence.

The *Feature* panel reports the selected feature’s attribution, embedding value, kernel template, dilation, and threshold. For feature #88, it shows $\alpha \approx 0.044$, $\text{emb} \approx 0.979$, $d = 2$, and $\text{thr} \approx -0.908$. The depicted kernel has positive weights on early taps and negative weights on later taps, making it a short-range contrast

⁴ The tool is available at <https://rocketlens.webs.upv.es>, with a video demo at https://youtu.be/rl4yZ3o_fI8 and source code at <https://github.com/lgalarraxai-for-minirocket/tree/rocketlens/RocketLens>




Fig. 1. *RocketLens* main interface. Linked views support inspection from prediction contrast, to attributed MiniRocket features, to kernel behavior and activated subsequences in the original time series.

detector. Its small dilation compares nearby observations within a local window that can be inspected directly in the heartbeat signal.

The *Time series* view grounds this hypothesis in the original recording and lets users slide the selected kernel along the time axis. In the example, the kernel is at $t = 6546$, on a small descending rise. The positive and negative taps can be compared with the curve, showing how the kernel responds to this short local contrast. The top strip shows where back-propagated β attribution concentrates, while the lower strip projects the selected feature’s threshold exceedances onto the input timeline. Together, these tracks separate time points that receive attribution from subsequences where the selected convolutional feature is active.

The *Convolution response* view plots the selected kernel’s response over the same temporal region and overlays the MiniRocket PPV threshold. For feature #88, most responses exceed this threshold ($\text{PPV}(v > \text{thr}) = 97.9\%$), so the feature is active across most alignments and mainly excludes strong negative-response segments. In the zoomed view, the selected alignment, marked by the vertical line, remains above the threshold; the deeper negative dip that follows crosses below it and therefore does not contribute to the PPV count. Linking these response patterns to receptive fields connects the MiniRocket computation to concrete heartbeat subsequences. This linked inspection shows that the abnormal prediction is partly supported by a convolutional feature that captures short oscillatory rises and declines while discounting deeper valley-like configurations.

To move beyond single-feature inspection, *RocketLens* includes a *Top embeddings* view (Figure 2) that aggregates the top- k features by $|\alpha|$ and displays their attribution, threshold, and dilation values. In the abnormal-heartbeat example, feature #88 ranks highest, but the evidence extends beyond a single feature. Several positive features share related contrast templates, including repeated leading-edge and trailing-edge kernels with small dilations. The top- k set also includes a very large-dilation feature, indicating that the prediction combines local shape evidence with broader temporal evidence.




Fig. 2. Top embeddings view. Top- k features ranked by $|\alpha|$ with attribution, threshold, and dilation, plus a grouped summary by kernel template.

The grouped kernel view makes this pattern easier to diagnose by showing recurring templates and their parameter regimes. Leading-edge contrast kernels with low dilations and negative thresholds suggest sensitivity to local oscillatory structure, while the high-dilation feature signals evidence is harder to interpret directly from the raw signal. The view helps users judge whether the explanation rests on interpretable local patterns, diffuse long-range configurations, or a mixture of both.

3 Conclusion

RocketLens shows that MiniRocket models can be inspected through their computational mechanisms, not just attribution scores. By linking features to kernel templates, dilations, thresholds, and time-localized responses, it shows how evidence is formed, where it concentrates, and which kernel regimes recur across predictions. Although *RocketLens* targets MiniRocket, it points to a broader direction for time-series XAI: explanations that expose the structure of model transformations alongside their attributions.

References

1. Dempster, A., Schmidt, D.F., Webb, G.I.: Minirocket: A very fast (almost) deterministic transform for time series classification. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD). pp. 248–257 (2021). <https://doi.org/10.1145/3447548.3467231>
2. Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., Muller, P.: Deep learning for time series classification: A review. *Data Mining and Knowledge Discovery* **33**(4), 917–963 (2019). <https://doi.org/10.1007/s10618-019-00619-1>
3. Kelodjou, G., Rozé, L., Masson, V., Galárraga, L., Gaudel, R., Tchuente, M., Termier, A.: Shaping Up SHAP: Enhancing Stability through Layer-Wise Neighbor Selection. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 13094–13103 (2024). <https://doi.org/10.1609/aaai.v38i12.29208>
4. Theissler, A., Spinnato, F., Schlegel, U., Guidotti, R.: Explainable ai for time series classification: A review, taxonomy and research directions. *IEEE Access* **10**, 100700–100724 (2022). <https://doi.org/10.1109/ACCESS.2022.3207777>