

# TFSM in-context learning for time-series classification of bearing health status

Michel Tokic<sup>1,2</sup>, Slobodan Djukanović<sup>1,3</sup>, Anja von Beuningen<sup>1</sup>, Cheng Feng<sup>1</sup>

1- Siemens AG, Data & Artificial Intelligence, Otto-Hahn-Ring 6,  
81739 Munich, Germany

2- Faculty of Mathematics, Informatics and Statistics,  
Ludwig-Maximilians-University Munich, Munich, Germany

3- Faculty of Electrical Engineering, University of Montenegro,  
Podgorica, Montenegro

**Abstract.** We introduce a classification method based on in-context learning using time-series foundation models (TFSMs). We demonstrate how data not included in the TFSM training can be classified without fine-tuning the foundation model or training a traditional classification model. Examples are represented as targets (class labels) and covariates (data matrices) within the TFSM prompt, enabling the classification of unknown covariate data patterns alongside the forecast horizon through in-context learning. We apply this method to vibration data to assess the health state of a bearing within a servo-press motor. The method transforms frequency-domain reference signals into pseudo time-series patterns, generates aligned covariate and target signals, and uses the TFSM to predict class-membership probabilities for predefined labels. Leveraging the scalability of pre-trained models, the proposed method demonstrates effectiveness across varying operational conditions. This represents significant progress beyond traditional, custom AI solutions towards broader AI-driven maintenance systems that could potentially be provided as Model- or Software-as-a-Service applications.

## 1 Introduction

In recent years, industries such as manufacturing, energy, transportation, and healthcare have increasingly recognized the importance of anomaly detection to prevent failures of critical machines and systems. The traditional maintenance approach involves reactive strategies that can lead to costly downtime, production losses, and safety risks. As a result, many companies are shifting towards predictive and preventive maintenance strategies, aiming to optimize operational efficiency by preemptively addressing signs of equipment deterioration.

Despite advancements in predictive maintenance, current solutions are largely tailored for specific applications and assets, relying heavily on domain-specific knowledge and handcrafted rules [1]. Although effective in controlled contexts, these systems often lack flexibility and require substantial effort to adapt to new environments or equipment.

The rise of artificial intelligence (AI) has introduced AI-driven maintenance strategies that rely on data-driven approaches. However, existing implementations primarily depend on traditional AI solutions, trained on asset-specific

datasets, and typically struggle to scale across different machines, industries, or failure patterns [1]. This limitation is due to their inability to generalize beyond predefined conditions, forcing companies to rely on custom-built solutions that hinder the full potential of AI-driven maintenance.

This work addresses these challenges by introducing a novel approach for the classification of data, specifically vibration data, using a time-series foundation model (TSFM). Unlike traditional AI models, TSFMs offer the generalization capabilities required for scalable and flexible classification solutions across different assets and operating conditions. Notably, they enable reliable data classification without extensive fine-tuning of large reference datasets or the need to train a traditional AI classification model [2]. This represents a significant step towards adaptive and generalizable AI models [3] and could potentially be offered to customers as Model- or Software-as-a-Service applications.

## 2 Method


This chapter details the proposed methodology for classifying the health state of a motor bearing using spectral data.

### 2.1 Time Series Foundation Models

TSFMs have emerged as a transformative paradigm for time-series forecasting, delivering strong zero-shot and few-shot performance across a broad spectrum of datasets and application domains [4–6]. Recent studies demonstrate that TSFMs can match or even surpass specialized statistical models and earlier task-specific models on unseen benchmarks. Our method is centered on the *General Time Transformer* (GTT) architecture [7] that captures complex temporal patterns and multivariate dependencies through alternating temporal and channel attentions in Transformer encoder blocks. By adding a learnable sink token to the end of channels for target variates similar to [4], and replacing the original point forecast head with a probabilistic forecast head that uses a four-component Gaussian Mixture Model to calibrate the output distribution, we adapt the GTT architecture for multivariate probabilistic forecast with covariates that are available in the forecast horizon [8]. Pretrained on a large-scale cross-domain dataset, GTT is a general-purpose zero-shot forecaster that takes a look-back window of  $L$  time points for one or more target variables (optionally with covariates) and predicts their next  $H$  values. Let  $\mathbf{x}_{1:L}$  and  $\mathbf{c}_{1:L+H}$  be the target variates and corresponding covariates. Targets are available only in the context window, whereas the covariates can be available also in the forecast horizon. The model is a function that predicts the distribution of  $\mathbf{x}_{L+1:L+H}$ , such that  $P(\mathbf{x}_{L+1:L+H}) = f(\mathbf{x}_{1:L}, \mathbf{c}_{1:L+H})$ .

### 2.2 Data preprocessing

Raw vibration signals sampled at 48 kHz from a servo-press motor are processed to compute the fast Fourier transform (FFT) [9], with recordings' length of



**Fig. 1:** Preprocessing: Signal spectrum (magnitude of FFT) is transformed into a  $60 \times 64$  matrix ( $N = 60$  data channels and  $M = 64$  frequency sub-bands per channel). *Note:* On the right graph we only depict five out of 60 possible channels.

$\approx 60$  seconds. As illustrated in Fig. 1, the preprocessing divides the whole spectrum into  $N = 60$  *data channels* (covariates). For each channel, the FFT is segmented into  $M = 64$  *frequency sub-bands*, with the mean value from each sub-band forming a feature. Overall, one raw vibration signal results in a covariate structure of size  $N \times M$ , used as input to the TSFM. No additional covariates have been added.

The system defines target variables for classification, corresponding to four health states (classes): 1. *Normal operation*, 2. *Outer ring fault*, 3. *Sand in bearing\**, and 4. *Inner ring fault*. Targets are one-hot encoded (1 if the class is present, 0 otherwise) over a forecast horizon of  $M = 64$  time steps, enabling sensitivity to subtle variations and emphasizing fault-indicative frequency components (see Figs. 1-3).


### 2.3 Few-shot prompting for in-context learning classification

We apply a few-shot prompting to adapt the pre-trained GTT for asset health classification tasks, as shown in Fig. 3.

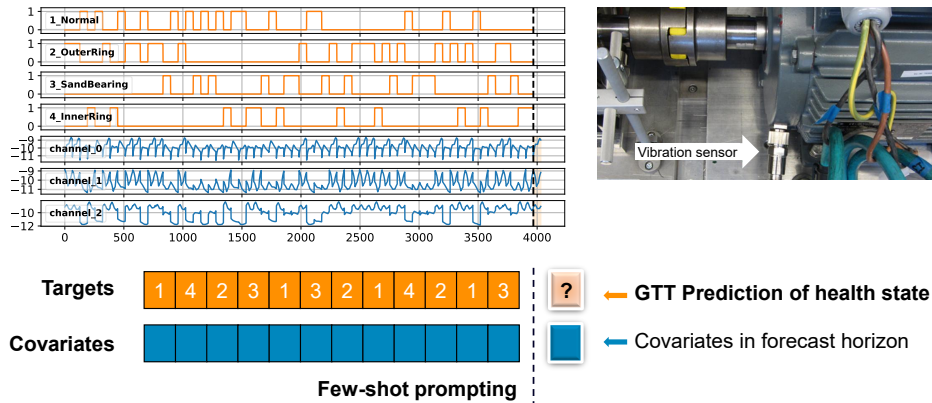
A small sequence of recorded target health states and their corresponding covariates forms the "context" or "prompt". This allows the model to learn about patterns associated with specific health states from a limited number of examples, which is known as *in-context learning* or *few-shot prompting*.

For a new covariate input, the model predicts the corresponding class as depicted in Fig. 4. Within 64 time steps, the model predicts the class of the tested covariate data. As a classification method, we apply a winner-takes-all rule at the last forecast step  $\mathbf{x}_{L+H}$ ; alternatively, a discrete probability distribution

\*Sand has been artificially added in the bearing.



**Fig. 2:** Example FFTs from the servo-press dataset. Note similarity between  $1\_Normal$  and  $3\_SandBearing$  and between  $2\_OuterRing$  and  $4\_InnerRing$ .



**Fig. 3:** A sequence of random samples from the servo-press dataset with one-hot encoded targets (orange) and corresponding covariates (blue) forms the few-shot prompting context. The model outputs a prediction of the targets w.r.t. covariates in the forecast horizon (blue block to the right of the dashed line). *Note:* Only the first three out of 60 covariate channels are depicted in the upper context.

could be obtained by applying a Softmax function to the predicted intensities.

### 3 Experiments


We conducted experiments with a GTT model with 750 million parameters pre-trained on a large-scale cross domain time series corpus with 124 billion data points. The model’s context accommodates 64 channels/variables, with a maximum of 4480 time steps, which is sufficient for up to  $4480/64 = 70$  servo-press data samples (vibration signals). Each sample occupies  $M = 64$  steps and  $N = 60$  covariate channels, as shown in Fig. 1 (left).

The dataset comprises 280 nearly class-balanced samples depicted in Fig. 2. Including four target variables, we have a total of 64 variables (targets + covariates) per class sample, which corresponds to the current limit of GTT.

A typical classification is depicted in Fig. 4, where the task was to predict class *2\_ OuterRing*. In the corresponding prompt, the last example was drawn from class *4\_ InnerRing*, explaining why the prediction for this class starts with high intensity (around 0.7) at the beginning, but develops down to almost 0 until step  $t = 63$ . In contrast, the intensity of class *2\_ OuterRing* rises towards almost 1. At  $t = 63$ , the applied winner-takes-it all rule classifies the data correctly.

An extensive study has been carried out to analyze the dependency on the number of context examples. Figure 5 presents the classification accuracy, obtained over 1000 runs, for randomly selected few-shot examples with varying context sizes (number of examples). Random sampling enables varying the last sample in the context, just before the forecast is generated (see Fig. 4).

A classification accuracy of 97.5% is achieved when using the full context length of 4480 time-steps. All four classes were identified with high precision and recall, exhibiting only minor misclassifications. These results demonstrate



**Fig. 4:** Correct classification of class  $2\_OuterRing$  (upper right), based on corresponding covariate matrix available in the GTT forecast horizon. Note, the prediction intensity stays/develops towards 0 for all other classes.


that the TSFM-based approach is highly reliable and consistent in differentiating between normal and faulty bearing conditions, particularly considering that this dataset was excluded from the model’s training data.

As a baseline, a multilayer perceptron (MLP) classifier was used, consisting of two fully connected hidden layers with 256 and 32 neurons, each followed by ReLU activation and 0.4 dropout, and a final softmax four-class output layer. The input spectra were preprocessed using the same aggregation procedure as in the proposed GTT method, but provided as a single one-dimensional feature vector rather than separated into covariates. Performance was evaluated using repeated stratified random splits (70/15/15 train-validation-test) with different random seeds in each run. Averaged over 20 runs, the MLP achieved an accuracy of 97.9% (dashed red line in Fig. 5).

## 4 Results and Discussion

The presented methodology offers significant advantages over traditional AI approaches by combining pre-trained TSFMs with few-shot prompting. It enables rapid deployment, eliminates training time, minimizes expert dependency, and ensures scalability across diverse assets and failure modes. Our work highlights the broader potential of TSFMs to scale to entire new applications, such as classification as investigated in our work. The method’s flexibility makes it adaptable to diverse assets and fault types, or even to entire new classification tasks.

The results demonstrate that aligning covariates and target data as pseudo time-series patterns effectively utilized the TSFM’s generalization capacity. The



**Fig. 5:** Comparison of classification accuracy over varied amount of few-shot examples.

context with 69 few-shot examples yielded 97.5% accuracy, on-par performance with the MLP baseline (97.9%). Both consistently distinguish between normal and faulty conditions, although the servo-press dataset was not part of the GTT training data. The FFT-based approach was crucial in highlighting fault-indicative frequencies, and the presentation of more random examples in the context contributed to minimizing the confusion between classes.

While this method is promising, a current limitation is the application to only a few target classes, such as four classes as in our study. Since the context length in the examined GTT model is currently limited to 4480 time steps, a higher number of classes would result in less represented examples per class, with potentially negative impact on the classification accuracy.

## References

- [1] S. Qiu, X. Cui, Z. Ping, N. Shan, Z. Li, X. Bao, and X. Xu, “Deep learning techniques in intelligent fault diagnosis and prognosis for industrial systems: A review,” *Sensors*, vol. 23, no. 3, 1305 2023.
- [2] A. Auer, D. Klotz, S. Böck, and S. Hochreiter, *Pre-trained forecasting models: Strong zero-shot feature extractors for time series classification*, 2025. arXiv: 2510.26777 [cs.LG].
- [3] Y. Liang, H. Wen, Y. Nie, Y. Jiang, M. Jin, D. Song, S. Pan, and Q. Wen, “Foundation models for time series analysis: A tutorial and survey,” in *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, 2024, pp. 6555–6565.
- [4] G. Woo, C. Liu, A. Kumar, C. Xiong, S. Savarese, and D. Sahoo, “Unified training of universal time series forecasting transformers,” in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 53 140–53 164.
- [5] A. Das, W. Kong, R. Sen, and Y. Zhou, “A decoder-only foundation model for time-series forecasting,” in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 10 148–10 167.
- [6] A. F. Ansari, O. Shchur, J. Küken, A. Auer, B. Han, P. Mercado, S. S. Rangapuram, H. Shen, L. Stella, X. Zhang, et al., “Chronos-2: From univariate to universal forecasting,” *arXiv preprint arXiv:2510.15821*, 2025.
- [7] C. Feng, L. Huang, and D. Krompass, “General time transformer: An encoder-only foundation model for zero-shot multivariate time series forecasting,” in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 3757–3761.
- [8] M. Tokic, A. von Beuningen, C. Tietz, and H.-G. Zimmermann, “Handling missing data in recurrent neural networks for air quality forecasting,” in *Proceedings of the 28th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, 2020, pp. 351–356.
- [9] M. K. Steven, *Fundamentals of statistical signal processing*. PTR Prentice-Hall, Englewood Cliffs, NJ, 1993.