2026 NDIA MICHIGAN CHAPTER
GROUND VEHICLE SYSTEMS ENGINEERING
AND TECHNOLOGY SYMPOSIUM
OPEN TECHNICAL SESSION
AUG. 11-13, 2026 - NOVI, MICHIGAN

DECISION-CENTRIC VIRTUAL METROLOGY FOR GEOMETRIC DEVIATION VIA FOUNDATION-EMBEDDING RETRIEVAL

Danny Hoang[1], Ryan Matthiessen[2], Christopher Miller[2], Nasir Mannan[2], Ruby ElKharboutly, PhD[3], David Gorsich, PhD[4], Matthew P. Castanier, PhD[4], and Farhad Imani, PhD[1]

[1]School of Mechanical, Aerospace, and Manufacturing Engineering, University of Connecticut, Storrs, Connecticut, USA
[2]Connecticut Center for Advanced Technology, East Hartford, Connecticut, USA
[3]Quinnipiac University, Hamden, Connecticut, USA
[4]US Army DEVCOM Ground Vehicle Systems Center, Warren, Michigan, USA

ABSTRACT

Modern defense manufacturing and sustainment require timely engineering decisions (e.g., inspection triage, rework/accept decisions, and process adjustment) based on the as-built geometric quality. Advanced machining systems generate rich multichannel controller and sensor streams, yet accurate geometric deviation labels remain costly and delayed because they depend on downstream metrology. This paper introduces ChronosGD, a retrieval-based virtual metrology framework. ChronosGD predicts pointwise geometric deviation from multichannel time series data by: (1) retrieving the most similar historical process windows in a frozen Chronos-2 embedding space, and (2) transferring deviation information through similarity-weighted aggregation. ChronosGD avoids plant-specific gradient retraining during deployment; adaptation is achieved by refreshing a labeled historical memory as new inspected parts become available, while preserving traceability through explicit neighbor provenance.

Citation: D. Hoang, R. Matthiessen, C. Miller, N. Mannan, R. ElKharboutly, D. Gorsich, M. Castanier, F. Imani, "Decision-Centric Virtual Metrology for Geometric Deviation via Foundation-Embedding Retrieval," In *Proceedings of the 2026 Ground Vehicle Systems Engineering and Technology Symposium (GVSETS)*, NDIA, Novi, MI, Aug. 11-13, 2026.

1. INTRODUCTION AND RESEARCH BACKGROUND

Advanced manufacturing for aerospace and defense increasingly depends on the ability to produce complex, high-value components with exceptional geometric precision. Multi-axis CNC machining, robotic machining cells, and integrated metrology systems have made it possible to manufacture free-form parts under tight tolerance requirements, but maintaining that accuracy remains a persistent challenge [1]. In these environments, geometric deviation is both a dimensional concern and directly influences assembly fit, aerodynamic behavior, structural reliability, inspection burden, and overall production efficiency. Recent studies in precision machining [2, 3] and aerospace assembly [4, 5] continue to show that geometric and spatial errors remain major barriers when tight tolerances must be sustained across complex tool paths and multi-stage processes. In defense manufacturing such deviations increase rework, extend inspection cycles, and delay return-to-service for mission-critical systems.

At the same time, modern manufacturing environments are increasingly data-rich but label-poor [6, 7]. High-frequency multichannel process signals can be collected directly from machine controllers, yet reliable geometric quality labels remain expensive and delayed because they depend on downstream metrology and inspection. This imbalance has made predictive quality modeling and virtual metrology important research directions, but their practical deployment remains difficult under changing production conditions. Recent work suggests that concept drift caused by tool wear, machine condition, and shifting operating regimes continues to limit the robustness of data-driven models [8, 9]. Although conventional supervised approaches such as decision trees, support vector regression, and neural networks can learn relationships between process signals and quality outcomes, they often require repeated retraining to remain effective as manufacturing conditions evolve [10, 11].

Time series foundation models provide a promising alternative as they separate large-scale pretraining from downstream task adaptation. Chronos demonstrated strong zero-shot and transfer performance across diverse forecasting benchmarks, while Chronos-2 extends this capability to univariate, multivariate, and covariate-informed settings, making it more suitable for industrial sensor streams [12, 13]. Motivated by these developments, this paper proposes ChronosGD, a retrieval-based framework for geometric deviation prediction that combines frozen Chronos-2 embeddings with similarity search over historical process windows. Instead of learning a single global regression function, ChronosGD retrieves similar historical signal patterns and transfers deviation information from these local analogs. This shifts adaptation from repeated parameter retraining to memory refresh, while also providing a traceable prediction pathway through explicit historical neighbors. Accordingly, ChronosGD is not absolutely training-free, since Chronos-2 is pretrained upstream, but it is deployment-level no-fine-tuning for the plant-specific deviation prediction task.

2. METHODOLOGY

We propose ChronosGD, a retrieval-based framework for geometric deviation prediction from multichannel process signals. An overview of the framework is shown in Fig. 1. The model estimates deviation at each inspection location by comparing local signal behavior from a query part against a historical memory of previously observed parts and transferring deviation information from the most similar historical patterns. An overview of the proposed framework is shown in Fig. 1. The method comprises two stages: an online prediction path and an offline memory path. In the online path, query signals are segmented into local windows, encoded into latent representations, and prepared for inference time retrieval. In the offline path, historical process windows and their associated deviation records are stored in a searchable database. The two

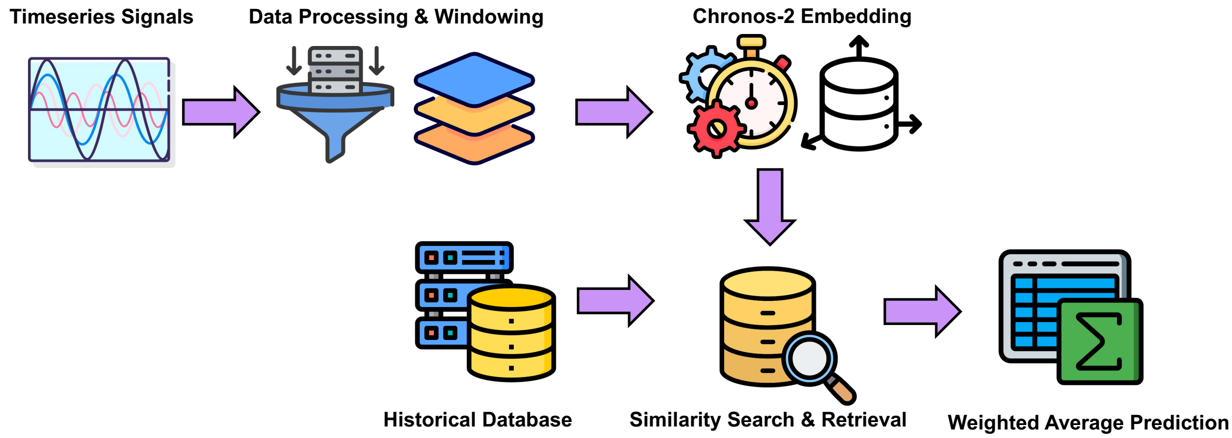

Figure 1: Overview of ChronosGD. Query signal windows are encoded and matched against a historical retrieval memory of prior process windows and deviation records. Retrieved neighbors are then aggregated to produce the final geometric deviation prediction.

paths meet at the retrieval stage, where the nearest historical analogs are identified and their deviation information is aggregated to generate the final prediction.

To keep the methodology general, we describe the framework independently of any specific encoder or retrieval engine. ChronosGD is then obtained as a particular instantiation in which the representation model is Chronos-2 and the retrieval backend is FAISS [14].

2.1. Problem Formulation

Let $\mathcal{P}$ denote the set of parts and let $\mathcal{S}$ denote the set of selected process signals. For each part $p \in \mathcal{P}$ and signal $s \in \mathcal{S}$, we observe a time series $x_{p,s} \in R^{L_{p,s}}$, where $L_{p,s}$ is the signal length.

We assume all parts are registered to a common finite set of inspection locations $\mathcal{T}$. For each part $p \in \mathcal{P}$ and inspection location $t \in \mathcal{T}$, the corresponding ground truth geometric deviation is denoted by $y_{p,t} \in R$.

Given a query part $q \in \mathcal{P}$, the goal is to predict $\hat{y}_{q,t}$ for each inspection location $t \in \mathcal{T}$, using the multichannel process signals of q together with a historical memory built from previously observed parts. Performance for query part q is evaluated by mean-squared error:

$$\mathrm{MSE}_q = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\hat{y}_{q,t} - y_{q,t})^2 . \qquad (1)$$

2.2. Historical Memory Construction

The proposed framework uses a fixed representation map $\phi : R^W \to R^d$, which transforms a local signal window of length W into a d-dimensional latent representation.

For each historical part $p \in \mathcal{P} \setminus \{q\}$ and signal $s \in \mathcal{S}$, sliding windows of length W and stride Δ are extracted from the corresponding signal. Using one-based indexing, define the valid window-index set

$$\mathcal{J}_{p,s} = \{j \in N : 1 + (j-1)\Delta + W - 1 \leq L_{p,s}\} . \qquad (2)$$

For each $j \in \mathcal{J}_{p,s}$, let $\tau_j = 1 + (j-1)\Delta$, and define the corresponding signal window as

$$w_{p,s,j} = x_{p,s}[\tau_j : \tau_j + W - 1] \in R^W . \qquad (3)$$

Each window is mapped into the latent space by

$$z_{p,s,j} = \phi(w_{p,s,j}) \in R^d . \qquad (4)$$

For each signal s, the historical retrieval memory excluding the query part is defined as the set of tuples

$$\mathcal{I}_s^{(-q)} = \{(z_{p,s,j}, p, j) : p \in \mathcal{P} \setminus \{q\},\ j \in \mathcal{J}_{p,s}\}. \quad (5)$$

Each memory entry stores the latent representation and metadata identifying the source part and window index. Since all parts are registered to the common inspection set $\mathcal{T}$, a retrieved historical window from source part p can be associated with the historical deviation $y_{p,t}$ at the same inspection location t.

2.3. Point-wise Query Window Construction

Prediction is performed independently for each inspection location $t \in \mathcal{T}$. To connect geometric inspection locations to the process signals, we assume that the regional construction defines, for each query part q, signal s, and inspection location t, a local signal-index region $\mathcal{R}_{q,s,t} \subseteq \{1, \ldots, L_{q,s}\}$. This region specifies the portion of signal s associated with inspection location t.

From the local region $\mathcal{R}_{q,s,t}$, up to R valid windows of length W are extracted:

$$\mathcal{W}_{q,s,t} = \{w_{q,s,t,i}\}_{i=1}^{m_{q,s,t}}, \quad 0 \le m_{q,s,t} \le R, \quad (6)$$

where $m_{q,s,t}$ is the number of valid query windows extracted for signal s around inspection location t.

Each query window is embedded into the same latent space:

$$u_{q,s,t,i} = \phi(w_{q,s,t,i}) \in R^d, \quad i = 1, \ldots, m_{q,s,t}. \quad (7)$$

For later aggregation, define the set of signals that provide at least one valid query window at inspection location t:

$$\mathcal{S}_{q,t}^{+} = \{s \in \mathcal{S} : m_{q,s,t} > 0\}. \quad (8)$$

2.4. Retrieval-Based Deviation Transfer

For each query representation $u_{q,s,t,i}$, the method retrieves a nonempty set of nearest historical neighbors from $\mathcal{I}_s^{(-q)}$. Let $\mathcal{N}(q,s,t,i) \subseteq \mathcal{I}_s^{(-q)}$ denote this neighbor set.

For each retrieved neighbor $n \in \mathcal{N}(q,s,t,i)$, write $n = (z_n, p_n, j_n)$, where z_n is the retrieved latent representation, p_n is the source part, and j_n is the corresponding historical window index. Let

$$d_n = D(u_{q,s,t,i}, z_n) \quad (9)$$

denote the distance or dissimilarity between the query representation and the retrieved historical representation, where $D : R^d \times R^d \to R_{\ge 0}$ is a chosen distance or dissimilarity measure.

Each neighbor contributes the historical deviation of its source part at the same registered inspection location t, namely $y_{p_n,t}$. The local prediction for query window $w_{q,s,t,i}$ is obtained by a similarity-weighted aggregation:

$$\tilde{y}_{q,s,t,i} = \sum_{n \in \mathcal{N}(q,s,t,i)} \alpha_n y_{p_n,t}. \quad (10)$$

The normalized retrieval weights are defined by

$$\alpha_n = \frac{\kappa(d_n)}{\sum_{r \in \mathcal{N}(q,s,t,i)} \kappa(d_r)}, \quad n \in \mathcal{N}(q,s,t,i), \quad (11)$$

where $\kappa : R_{\ge 0} \to R_{\ge 0}$ is a nonnegative nonincreasing kernel that assigns larger weights to more similar historical neighbors. We assume that $\sum_{r \in \mathcal{N}(q,s,t,i)} \kappa(d_r) > 0$. It follows that

$$\sum_{n \in \mathcal{N}(q,s,t,i)} \alpha_n = 1. \quad (12)$$

2.5. Regional Aggregation Within Each Signal

A single inspection location may produce multiple query windows within the same signal region. For each signal $s \in$

$\mathcal{S}_{q,t}^{+}$, the local window-level predictions are aggregated into a signal-specific estimate:

$$\hat{y}_{q,s,t} = A_s\left(\tilde{y}_{q,s,t,1}, \ldots, \tilde{y}_{q,s,t,m_{q,s,t}}\right), \quad (13)$$

where $A_s(\cdot)$ is an aggregation operator over the regional windows for signal s.

In the simplest case, A_s is the arithmetic mean:

$$\hat{y}_{q,s,t} = \frac{1}{m_{q,s,t}} \sum_{i=1}^{m_{q,s,t}} \tilde{y}_{q,s,t,i}, \quad s \in \mathcal{S}_{q,t}^{+}. \quad (14)$$

2.6. Cross-Signal Fusion

Since deviation is inferred from multiple process signals, the signal-specific predictions are combined to form the final estimate:

$$\hat{y}_{q,t} = F_t\left(\{\hat{y}_{q,s,t}\}_{s\in\mathcal{S}_{q,t}^{+}}\right), \quad (15)$$

where $F_t(\cdot)$ denotes the cross-signal fusion rule at inspection location t.

A common choice is convex fusion:

$$\hat{y}_{q,t} = \sum_{s\in\mathcal{S}_{q,t}^{+}} \beta_{q,s,t}\hat{y}_{q,s,t}, \quad (16)$$

with weights satisfying

$$\beta_{q,s,t} \geq 0, \quad \sum_{s\in\mathcal{S}_{q,t}^{+}} \beta_{q,s,t} = 1. \quad (17)$$

Here, $\beta_{q,s,t}$ denotes the contribution of signal s to the final prediction at inspection location t.

2.7. ChronosGD Instantiation

ChronosGD is obtained as a specific realization of the general framework above. In this work, the representation map $\phi(\cdot)$ is implemented by a frozen Chronos-2 encoder, the retrieval memory $\mathcal{I}_s$ is implemented using FAISS, the dissimilarity function $D(\cdot,\cdot)$ is squared Euclidean distance in the embedding space, and the kernel $\kappa(d)$ is inverse-distance weighting $\kappa(d) = (d+\varepsilon)^{-1}$ with $\varepsilon > 0$ for numerical stability. Under this instantiation, adaptation is achieved by refreshing the historical memory with newly observed windows and their associated deviations, without gradient-based retraining of the encoder.

3. EXPERIMENTAL DESIGN AND RESULTS

The ChronosGD framework was evaluated using a rotor blade manufacturing experiment conducted at the Connecticut Center for Advanced Technology (CCAT). Figure 2 summarizes the overall fabrication, signal acquisition, and inspection workflow used in the study. Grade 5 titanium rods with a diameter of 1 1/4 in (31.75 mm) were used to manufacture rotor blades with a length of 3 in (76.2 mm) designed at CCAT. Siemens NX CAM was used to generate the baseline tool paths required to manufacture the blades on a 4-axis Trak VMC10si.

Blade fabrication consisted of ten machining operations: (1) tip engage, (2) top half rough, (3) top half semi-finish, (4) bottom half rough, (5) bottom half semi-finish, (6) platform outer rough, (7) platform outer finish, (8) platform top finish, (9) blade finish, and (10) tip finish. For the ChronosGD study, the process signals associated with the blade finishing stage were used as the input time series data for prediction. During machining of each part, the corresponding time series signals were recorded at 500 Hz directly from the SINUMERIK controller on the Trak machine and stored for subsequent analysis. Specifically, signals were collected for the X, Y, Z, and A axes, as well as the spindle, yielding a total of 65 channels. For each axis/component, 13 signal variables were recorded: Load, Encoder 1 Position, Control Position, Torque, Velocity Feed

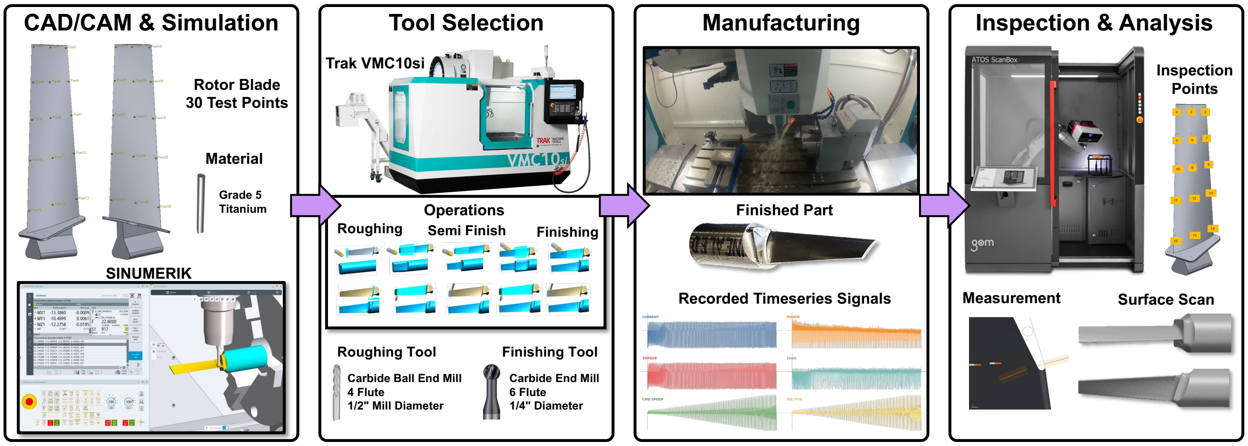

Figure 2: Illustration of the experimental design from part design, tool selection, manufacturing, and inspection.

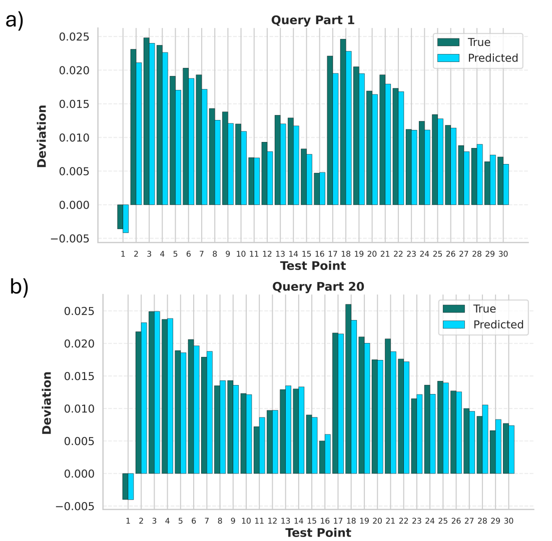

Figure 3: Example showcasing the prediction of ChronosGD model for a) Part 1 and b) Part 20.

Forward, Command Speed, Destination Position, Control Differential, Control Differential 2, Current, Contour Deviation, Power, and Torque Feed Forward. These recorded signals form the multichannel process traces used to construct the historical retrieval memory and the query inputs for ChronosGD.

A total of 20 rotor blades were manufactured and inspected for geometric deviation relative to the nominal CAD model using a GOM ATOS ScanBox. For each blade, the scanned surface was registered to the nominal geometry and deviation values were extracted at labeled inspection locations. We defined 30 inspection points distributed across the blade surface and arranged them into matched pressure-side and suction-side pairs using a consistent pair-key convention of the form $i + (i + 15)$ (e.g., `2+17`, `3+18`, ..., `16+31`) so that corresponding blade locations could be analyzed consistently across parts. The inspection results for all blades were exported to CSV and aggregated into a deviation database. These measured deviations serve as the ground truth targets for ChronosGD, while the recorded finishing operation signals serve as the time series inputs used to predict deviation at each inspection point.

ChronosGD was implemented in Python with a window size of 250 and stride of 50. Figure 3 showcases example prediction results of ChronosGD for two query parts. Figure 3a) shows the predicted and true geometric deviations for Query Part 1 across 30 inspection points, while Fig. 3b) shows the corresponding results for Query Part 20. In both cases, ChronosGD captures the overall deviation profile across the blade, including the major peaks in the early and middle inspection regions and the gradual reduction in deviation toward later test points. The predicted values generally follow the spatial trend of the measured deviations with relatively small local discrepancies, indicating that the retrieval-based formulation preserves the dominant part-specific variation pattern. These examples illustrate that ChronosGD can recover the overall geometric deviation distribution of unseen parts from recorded finishing operation signals while maintaining point wise fidelity across the inspection sequence.

We further evaluate ChronosGD predictive performance against conventional machine learning baselines and deep learning methods under matched data construction and segmentation settings. The conventional models included decision tree regression and support vector regression (SVR). Deep learning baselines included multilayer perceptron (MLP), time series Convolutional Neural Network (CNN) [15], Fully Convolutional Neural Network (FCN) [16], Residual Network (ResNet) [16], and InceptionTime [17]. To preserve fairness, all compared methods used the same test point segmentation logic, window size, and regional window construction. The ChronosGD and conventional regression baselines operated on the same per-signal query windows extracted from each test point region, while the deep learning baselines used multivariate windows formed from the same region boundaries across all selected signals. Under this protocol, observed performance differences primarily reflect the modeling approach rather than inconsistencies in preprocessing or feature construction.

Figure 4 summarizes the aggregate prediction error across all evaluated query parts. ChronosGD achieves the lowest error among all compared methods in both root mean squared error (RMSE $= 0.00224$) and mean absolute error (MAE $= 0.00156$), indicating that the

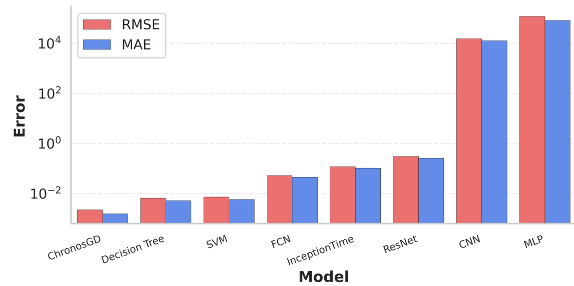

Figure 4: Aggregate error comparison of ChronosGD and baseline models using RMSE and MAE.

retrieval-based formulation provides the most accurate overall deviation estimates under the shared evaluation protocol.

Among the conventional baselines, decision tree regression (RMSE $= 0.00661$, MAE $= 0.00523$) and SVR (RMSE $= 0.00730$, MAE $= 0.00585$) yield the next-lowest errors, although both remain above ChronosGD. The FCN, InceptionTime, and ResNet baselines exhibit substantially higher error levels (FCN: RMSE $= 0.05248$, MAE $= 0.04550$; InceptionTime: RMSE $= 0.11841$, MAE $= 0.10427$; ResNet: RMSE $= 0.30648$, MAE $= 0.26028$), suggesting reduced effectiveness under the present training and evaluation setting.

CNN and MLP produce errors several orders of magnitude larger than the other methods (CNN: RMSE $= 1.55 \times 10^4$, MAE $= 1.29 \times 10^4$; MLP: RMSE $= 1.20 \times 10^5$, MAE $= 8.27 \times 10^4$) suggesting these methods are not effective for high multi-channel timeseries prediction.

Figure 5 reports part-wise performance in terms of $\log_{10}(\mathrm{MSE})$ across the 20 query parts. ChronosGD maintains low error across nearly all parts and exhibits the most favorable overall heatmap profile (mean -5.627, range $[-6.270, -4.477]$), and achieving the lowest MSE on all $20/20$ query parts. Decision tree regression and SVR also show relatively stable behavior across parts (Decision Tree: mean -4.366, std 0.076; SVR: mean -4.276, std 0.051), but their

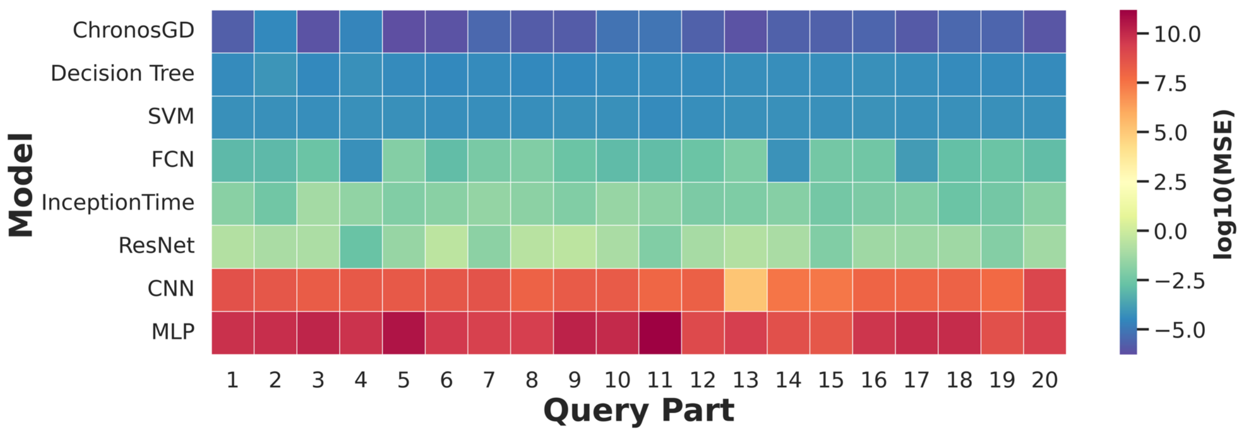

Figure 5: Part-wise prediction error comparison across 20 query parts.

errors remain consistently higher than those of ChronosGD. FCN, InceptionTime, and ResNet display greater variability across query parts, indicating stronger dependence on the specific query instance (FCN: mean -2.842, std 0.619; InceptionTime: mean -1.999, std 0.342; ResNet: mean -1.299, std 0.570). By contrast, CNN and MLP are associated with uniformly high error across most parts, with mean values of 8.069 and 9.647, respectively. Overall, these results suggest that ChronosGD improves both predictive accuracy and consistency across unseen query parts relative to the conventional regression baselines and deep learning models evaluated in this study, indicating robust generalization across diverse query conditions rather than an advantage driven by only a small subset of favorable cases.

4. CONCLUSION

This paper presented ChronosGD, a retrieval-based framework for geometric deviation prediction from multichannel manufacturing signals using frozen Chronos-2 embeddings and historical similarity search. ChronosGD is formulated as a deployment-level no-fine-tuning approach that replaces repeated task-specific retraining with memory refresh over historical process windows. Experimental evaluation on rotor blade finishing data showed that ChronosGD accurately recovered part-wise deviation profiles and achieved lower prediction error and more consistent performance across unseen parts than the conventional regression and deep learning baselines considered in this study. Future research will investigate real-time prediction during machining, including streaming inference from controller data, as well as edge deployment strategies that enable low-latency on-machine or near-machine decision support for quality monitoring and adaptive manufacturing applications.

5. ACKNOWLEDGMENTS

This work was supported under Cooperative Agreement W56HZV-21-2-0001 with the U.S. Army DEVCOM Ground Vehicle Systems Center (GVSC), through the Virtual Prototyping of Autonomy Enabled Ground Systems (VIPR-GS) program; the National Science Foundation (grant no. 2434519); and the Department of Energy (grant no. DE-EE0011029).

Disclaimer: Reference herein to any specific commercial company, product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the Department of the Army (DoA). The opinions of the authors expressed herein do not necessarily state or reflect those of the United States Government or the DoA and shall not be used for advertising or product endorsement purposes.

References

[1] D. Hoang, N. Mannan, R. ElKharboutly, R. Chen, and F. Imani, "Edge cognitive data fusion: From in-situ sensing to quality characterization in hybrid manufacturing process," in *International Manufacturing Science and Engineering Conference*, vol. 87240. American Society of Mechanical Engineers, 2023, p. V002T06A029.

[2] D. Hoang, R. Chen, G. Bollas, and F. Imani, "Hyperdimensional computing for explainable information fusion and multi-task adaptation in advanced manufacturing," *Information Fusion*, p. 103898, 2025.

[3] S. He and T. Shi, "Sensitivity analysis of surface contour error to geometric errors for four-axis ultra-precision machine tools," *Journal of Manufacturing Processes*, vol. 131, pp. 569–588, 2024.

[4] R. Yang, J. Huang, Z. Chen, D. Lian, S. Gao, X.-C. Zhong, J. Li, Y. Liu, and J. Tan, "Measurement and optimization method for aero-engine rotors based on binocular multi-line laser sensing and virtual assembly," *Measurement*, vol. 242, p. 115808, 2025.

[5] F. Guo, Y. Zhang, C. Song, and X. Sha, "Identification and precise optimization of key assembly error links for complex aviation components driven by mechanism and data fusion model," *Advanced Engineering Informatics*, vol. 64, p. 103059, 2025.

[6] E. Stathatos, E. Tzimas, P. Benardos, and G.-C. Vosniakos, "Convolutional neural networks for raw signal classification in cnc turning process monitoring," *Sensors*, vol. 24, no. 5, p. 1390, 2024.

[7] Q. Zhuang, N. Wan, Y. Guo, G. Zhu, and D. Qian, "A state-of-the-art review on the research and application of on-machine measurement with a touch-trigger probe," *Measurement*, vol. 224, p. 113923, 2024.

[8] P. Wiederkehr, F. Finkeldey, and T. Siebrecht, "Reduction of experimental efforts for predicting milling stability affected by concept drift using transfer learning on multiple machine tools," *CIRP Annals*, vol. 73, no. 1, pp. 301–304, 2024.

[9] A. Bilen, A. Grommes, F. Stamer, and G. Lanza, "Coping with concept drift in a virtual metrology application to predict part quality in micro gear manufacturing," *Measurement: Sensors*, vol. 38, p. 101797, 2025.

[10] A. K. Kausik, A. B. Rashid, R. F. Baki, and M. M. J. Maktum, "Machine learning algorithms for manufacturing quality assurance: A systematic review of performance metrics and applications," *Array*, vol. 26, p. 100393, 2025.

[11] C. Kubik, J. Hofmann, D. A. Molitor, M. Becker, and P. Groche,

"Reliable machine learning models for manufacturing processes," *The International Journal of Advanced Manufacturing Technology*, vol. 138, no. 11, pp. 5913–5927, 2025.

[12] A. F. Ansari, L. Stella, C. Turkmen, X. Zhang, P. Mercado, H. Shen, O. Shchur, S. S. Rangapuram, S. P. Arango, S. Kapoor *et al.*, "Chronos: Learning the language of time series," *arXiv preprint arXiv:2403.07815*, 2024.

[13] 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.

[14] M. Douze, A. Guzhva, C. Deng, J. Johnson, G. Szilvasy, P.-E. Mazaré, M. Lomeli, L. Hosseini, and H. Jégou, "The faiss library," *IEEE Transactions on Big Data*, 2025.

[15] B. Zhao, H. Lu, S. Chen, J. Liu, and D. Wu, "Convolutional neural networks for time series classification," *Journal of Systems Engineering and Electronics*, vol. 28, no. 1, pp. 162–169, 2017.

[16] Z. Wang, W. Yan, and T. Oates, "Time series classification from scratch with deep neural networks: A strong baseline," in *2017 International joint conference on neural networks (IJCNN)*. IEEE, 2017, pp. 1578–1585.

[17] H. Ismail Fawaz, B. Lucas, G. Forestier, C. Pelletier, D. F. Schmidt, J. Weber, G. I. Webb, L. Idoumghar, P.-A. Muller, and F. Petitjean, "Inceptiontime: Finding alexnet for time series classification," *Data Mining and Knowledge Discovery*, vol. 34, no. 6, pp. 1936–1962, 2020.