



AI-driven digital twins in grinding: architecture, data pipelines, and initial industrial demonstrations

Downloaded from: <https://research.chalmers.se>, 2026-10-02 13:00 UTC

Citation for the original published paper (version of record):

Polak, M., Šenica, D., Brajak, V. et al (2026). AI-driven digital twins in grinding: architecture, data pipelines, and initial industrial demonstrations. *Procedia CIRP*, 146: 397-402. <http://dx.doi.org/10.1016/j.procir.2026.03.256>

N.B. When citing this work, cite the original published paper.

CIRP Conference on Manufacturing Systems (CIRP CMS)

AI-driven digital twins in grinding: architecture, data pipelines, and initial industrial demonstrations

Marko Polak^a, David Šenica^a, Viktor Brajak^a, Magnus Wahlgård^{b,c}, Tomas Gustavsson^c, Peter Krajnik^{b,*}^a*Medius, Tehnološki park 21, SI-1000 Ljubljana, Slovenia*^b*Chalmers University of Technology, Department of Mechanical Engineering, Hörsalsvägen 7B, SE-412 96 Gothenburg, Sweden*^c*AB SKF, Sven Wingquists Gata 2, 415 05 Gothenburg, Sweden** Corresponding author. Tel.: +46(0)31 772 13 11. E-mail address: peter.krajnik@chalmers.se**Abstract**

An event-driven digital-twin platform for grinding (AI-IGC4.0-DT) was developed with modular services for ingestion, storage, AI/ML operations, and visualization. Two complementary data pipelines were designed: for high-frequency, real-time (HFRT) signals, and large-file/REST exchange for contextual manufacturing data – unified via process-instance linkage (workpiece, wheel, material batch, time). The platform is deployable as DTaaS (cloud-hosted) and DTaaP (on-premises) under industrial integration and cybersecurity constraints. Feasibility was demonstrated in two industrial pilots: (i) bearing-ring grinding, where predicting deviation from the programmed size enabled high-precision runtime adaptation; and (ii) crankshaft grinding, where fixed dressing intervals were re-engineered into dynamic, quality-aware targets. Data linkage and governed, secure deployment were identified as primary enablers of industrial AI in grinding.

© 2026 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

Peer-review under responsibility of the scientific committee of the 2025 CIRP CMS

Keywords: Grinding; Machine learning; Artificial intelligence; Digital twin**1. Introduction**

Grinding is often considered operator-dependent; however, in high-end manufacturing companies, process plans permit only minor adjustments on the shop floor. Process parameters are typically determined using models [1] rather than subjective choices. Because these parameter sets are applied statically (i.e., not adjusted to the current material batch, wheel specification, or thermal state), the process windows are intentionally conservative to maintain low scrap rates and high process capability index, C_{pk} , which limits responsiveness to variability and potential efficiency gains.

Despite these fixed process plans, modern CNC systems expose interfaces for adaptive control (e.g., dressing intervals, thermal compensation). In practice, these capabilities are seldom utilized. Recent advances in AI make it feasible to

operationalize run-time parameter adaptation at the machine, providing continuous availability for adjustments under known conditions – whereas the success of such adaptation depends on the amount and quality of data available [2]. Grinding machines generate a wide range of signals – from CNC drive measurements (e.g., current/power) [3], temperatures, to external sensors such as acoustic emission [4] and size gauges. In parallel, contextual manufacturing data – material certificates, wheel- and grinding-fluid records, as well as quality measurements – are stored in adjacent systems but are not reliably linked to the corresponding process instance (i.e., the specific component/material batch, and wheel running at that time). In practice only a small fraction of these data are used for optimization; the majority are only temporarily logged or lost (e.g., overwritten at each CNC interpolation cycle).

This motivates the development of an event-driven digital platform for grinding (AI-IGC4.0-DT), which facilitates the end-to-end development and deployment of AI/ML models within the platform architecture. The platform is deployable as Digital Twin as a Service (DTaaS) or Digital Twin as a Platform (DTaP) to support a broad range of grinding DT applications. In this paper, two industrial pilots are reported as proof-of-concept demonstrations of AI-enabled, control-relevant outputs in grinding (predictions and recommendations), enabled by process-instance-linked integration of streaming and contextual data. Section 3 describes the platform architecture and data pipelines, and Section 4 reports the two pilots.

2. Background

Our preliminary research indicated that incorporating edge computing into existing manufacturing systems facilitated handling of HFRT data [5], and it investigated integration of edge devices into existing architectures to potentially enable AI/ML operations on HFRT streams. The validation focused on interoperability, ease of configuration, and maintainability as the primary criteria and included indicative latency and throughput measurements. The main challenges identified were connectivity reliability, end-to-end latency, and rapid growth in data volumes at scale. Additional risks identified included avoiding interference with legacy controls and protecting sensitive data, mitigated by local (on-premises) storage and processing. The study concluded that careful data management is necessary to avoid wasting resources through indiscriminate collection. Clear purpose definitions and lean data usage principles were recommended, as well as AI/ML for data-driven modeling [1]. Based on these findings, the focus shifted from single-machine edge devices to edge servers that can aggregate data from multiple machines and support orchestrated workloads while maintaining data processing close to the machines within existing security constraints.

Cyber-physical systems (CPS), including public and private clouds, are essential for standardized provisioning, data retention, large-scale AI/ML training, and distributed analytics. To achieve these capabilities, containerization and orchestration are required to standardize deployment and automate scaling. Here, heterogeneity poses a persistent challenge because architectures, operating systems, communication protocols, and computational capacities vary widely among vendors and machines. While some CNCs support containerized workloads, others require proprietary runtimes with limited interfaces. This complicates workload placement and necessitates abstraction via standardized interfaces and orchestration policies. These heterogeneity and security constraints motivated a common design that can operate in both deployment modes – DTaaS (cloud-hosted) and DTaP (on-premises) – without reauthoring core services.

CPS implementation is often considered on five levels: (i) connection (shop-floor data acquisition from machines), (ii) conversion (signal processing and feature extraction), (iii) cyber (digital twin representation and services), (iv) cognition (data analytics and AI/ML-based decision support), and (v) configuration (bidirectional, adaptive feedbacks to the physical processes). In this work, the AI-IGC4.0-DT platform

principally addresses levels (i)–(iv); while configuration is deferred to future work.

Digital twins (DTs) can be classified by maturity in terms of their integration architecture and functional capability: (i) a digital model, in which data exchange between physical and digital artifacts is manual; (ii) a digital shadow, which is characterized by an automated, unidirectional flow of data from the physical asset to its digital representation; and (iii) a digital twin, which entails automated, bidirectional communication between the physical system and its digital counterpart, enabling both monitoring and control [6]. This work focuses on enabling shadow-to-twin capabilities through reliable data acquisition and contextual linkage. It also focuses on serving these capabilities across DTaaS and DTaP modes. Although comprehensive DT maturity assessment frameworks that cover the lifecycle across different categories and characteristics exist [7], DT maturity assessment is outside the scope of this study.

Cybersecurity emerged as the primary constraint. Much of the necessary data cannot be exported to third-party AI developers, such as the SME developing the AI-IGC4.0-DT platform here. Authenticated and encrypted communication, least-privilege access, device and workload attestation, data minimization, and encryption in transit and at rest are all necessary, as well as auditable controls. In response, a modular, event-driven architecture was developed, with layers for ingestion, storage, visualization, and AI/ML operations—deployable as DTaaS or DTaP under governed access. Section 3 details the resulting architecture and pipelines.

Prior work on DT for machining [2] and grinding [8] has commonly emphasized (i) conceptual frameworks [9], (ii) edge/cloud connectivity for data acquisition [10], and (iii) data-driven monitoring or modeling for specific process indicators [11]. Fewer studies, however, are concerned with an integrated architecture that (a) combines HFRT streaming and contextual manufacturing data into a single process-instance-linked representation, (b) supports DTaaS/DTaP deployment under industrial cybersecurity constraints, and (c) operationalizes AI/ML models lifecycle management (training, versioning, and deployment/serving) as part of the DT platform. The present work addresses this gap by specifying and demonstrating an event-driven DT platform for grinding in two industrial pilots.

3. AI-IGC4.0-DT digital platform

3.1. Architecture

Applications of AI in grinding require a cloud-native, event-driven architecture; therefore, components were selected based on the use-cases' latency and domain requirements. The resulting architecture is shown in Fig. 1. It was designed for deployment across public/private clouds and on-premises Kubernetes, supporting modular horizontal scaling and interoperability via standardized industrial and enterprise interfaces (OPC UA, MQTT, AMQP/HTTPS, REST, S3, FTP/SFTP), which is essential for integrating brownfield equipment and legacy enterprise systems. The layered, event-driven design provides the basis for AI/ML operations and validation against production data.




Fig. 1. AI-IGC4.0-DT platform architecture.

Following the definition of digital twins for cutting processes by Bergs et al. [2], the platform architecture was designed to separate the digital shadow (DS) – i.e., instance-linked data acquisition and contextualization – from DT services that operate on this shadow. In the architecture, dual data pipelines and process-instance linkage (e.g., workpiece ID, wheel ID, material batch, timestamps) establish DS capabilities, whereas the DT is realized by model-enabled services that consume these linked data to provide monitoring and decision-support functions. The AI/ML models are treated as model components within DT services rather than as the DT itself; their outputs are integrated into DT views/APIs. Fully automated bidirectional feedback (closed-loop actuation of process parameters) was not yet implemented and constitutes future work.

The platform architecture consists of four functional domains (color-coded in Fig. 1): (i) event-driven data engineering (grey), providing ingestion and preprocessing (cleaning, semantic joining, transformation) for data streams and contextual data; (ii) multi-modal data storage (blue), combining time-series/relational storage with object storage to accommodate structured sensor data and contextual datasets; (iii) DT visualization (red), aggregating multi-source data into DT views with model outputs; and (iv) AI model orchestration (yellow), managing the model lifecycle from development and retraining to deployment and serving. The AI domain was separated into training and inference environments – on the same or separate clusters – to prevent resource-intensive optimization workloads from interfering with low-latency production inference.

Kubernetes served as the foundation for container orchestration, automating deployment, scaling, and management of open-source components and custom microservices, ensuring portability across cloud and on-premises infrastructure. The platform was designed for modularity and flexibility so that, for a given use case, only the required components need to be deployed, while the modular structure enables straightforward integration of additional components and AI tools as requirements evolve.

3.2. Data pipelines

Two complementary data pipelines were specified to address the different characteristics of grinding-related data. One pipeline handled HFRT process data, which required low-latency processing. The other pipeline handled contextual manufacturing data, which required extraction, validation, and traceable storage. HFRT signals were ingested at the edge and forwarded to the Azure IoT Hub. A dedicated IoT API client then consumed the signals and streamed events to Apache Kafka. The records were cleaned and transformed and then stored as timestamped and windowed samples in PostgreSQL/TimescaleDB. To shorten the data path for on-premises or more-demanding latency scenarios, edge-to-Kafka or edge-to-model paths were defined. Operationally, ingestion was decoupled through the event bus: raw HFRT data were normalized and enriched prior to persistence, and downstream consumers accessed either the event stream or the stored time-series for analytics and model inference.

The platform exposed S3 and REST endpoints for external ingestion and included S3/REST clients to consume enterprise data. Large files (e.g., PDFs, CSVs) were written to MinIO (S3); a Data Extractor parsed/validated content and wrote structured records to the databases while retaining originals in MinIO for traceability. For enterprise APIs, a REST-to-Kafka interface was developed to channel data into the event stream, enforcing uniform downstream transformation and storage across sources. Together, these pipelines unify sensor streams with contextual records while maintaining process-instance linkage (e.g., component, wheel, material batch, and time). This linkage was necessary for subsequent AI/ML operations.

3.3. AI/ML operations

The AI/ML operations were divided into two parts. The first focused on model development, including training and retraining; MLflow served for managing the model lifecycle, ensuring versioning, and enabling reproducibility. The second part was a custom-built ML operations solution providing a Model Serving Facade, i.e., a unified interface for deploying and running models in production environments, with runtime model switching and monitoring of model performance. These capabilities were integrated into the DT custom user interface and supported deployment to edge devices, centralized Kubernetes clusters, or other target environments. Model versions could be staged in testing and production and compared side-by-side on live streams prior to promotion, reducing deployment complexity and operational effort for administrators and data scientists.

4. Industrial use cases and proof-of-concept demonstration

Two industrial use cases were selected to demonstrate the AI-IGC4.0-DT pipelines for HFRT signals and large-file/REST contextual data and to show that the same event-driven platform supports both workloads under industrial integration and cybersecurity constraints.

4.1. From in-process gauging to data-driven quality control in OD grinding of bearing rings

In grinding of thin-walled bearing rings, dimensional accuracy is mainly influenced by the thermal state of the machine and the normal forces, with thermally induced structural drift being the predominant limiting factor [12]. Accordingly, the SKF use case adopted a data-centric approach to managing thermal stability, relying on data-driven inference rather than a fully coupled, bidirectional DT with real-time feedback control [13]. Process and machine parameters, including temperatures, were continuously logged to monitor the thermal state over rolling two-hour windows. The objective was to achieve “sizematic” control with “gaugematic” quality, thereby reducing dependence on the in-process gauge and enabling greater automation. In a traditional configuration, diameter control relies on the in-process gauge as the master and the machine slide position as the backup, which prevents fully automatic reset / bearing-type changeover. Since the used machine type exhibits high positional repeatability once thermal equilibrium is reached, it is feasible to rely on the final slide position for size control, except during thermal transients or significant variations in specific material removal rate.

The approach was demonstrated on a Lidköping SGP 320 shoe-centerless grinding machine (Fig. 2) where cycle-synchronous process data and size measurements were acquired via an OPC UA/BFC Gateway (Brownfield Connectivity Gateway). Data ingestion and storage were configured on a private cloud, with validation and visualization performed through Grafana dashboards to support data exploration, feature/label verification, and operational monitoring during model development.



Fig. 2. Lidköping SGP 320 shoe-centerless grinding machine configured for high-volume external cylindrical grinding of bearing outer rings.

For this use case, three primary data types were utilized: (i) time-series data captured during grinding (e.g., slide-table command feed, wheel speed), (ii) static setup data defined prior to grinding (e.g., workpiece diameter and width), and (iii) key performance indicator (KPI) data recorded after each cycle (e.g., actual grinding time). The objective was to develop a data-driven model capable of predicting the final slide table position for the subsequent grinding cycle. This prediction represents a high-precision regression task, where even a relative error of 0.01% is unacceptable.

To meet this requirement, a hybrid neural network capable of jointly modeling temporal features and static covariates was developed. The architecture comprised a Temporal Convolutional Network (TCN) to extract temporal dependencies from the time-series inputs, and linear layers combined with a Gated Residual Network (GRN) to encode the static covariates. The outputs of these two streams were merged through an Alpha Gate layer, which applied adaptive weighting to balance the contributions of each feature set. The fused representation was then passed through a final linear layer to produce the table slide-position prediction. The resulting network architecture is illustrated in Fig. 3.

Two predictive strategies were evaluated. The first aimed to predict the absolute wheel position. Although the resulting prediction accuracy was high (within the 0.01%), the residual deviation remained above the threshold acceptable for precision control. Consequently, a second strategy was developed, in which the model predicted the deviation (error) of the programmed position rather than the position itself. Because these target error values exhibited a narrower dynamic range, the approach produced more stable and accurate predictions across validation sets. Model performance was evaluated using standard regression metrics (e.g., mean absolute error and coefficient of determination) under cross-validation; the specific tolerance limits and target precision levels are proprietary but fall within the range required for high-precision manufacturing of bearings.




Fig. 3. Architecture of the hybrid TCN-GRN neural network used: time-series branch (4 TCN blocks, kernel=3, dilations 1/2/4/8) and static-feature branch (fully connected layers + GRN), fused via an Alpha Gate.

4.2. Intelligent dressing intervals in grinding of heavy-duty automotive crankshafts

While the crankshaft grinding cycle can be optimized using physical models that capture the geometry, kinematics, and temperature of the process [14], the wear behavior of cBN grinding wheels cannot yet be predicted analytically with sufficient accuracy for reliable dressing-interval planning [15]. Consequently, in industrial crankshaft grinding, dressing intervals are usually set at fixed intervals. However, this practice is inefficient because wheel wear is strongly influenced by variations in incoming material batches [16], which directly affect the optimal timing of dressing.

In this Volvo Group Trucks Operations pilot, the use case focused on automating extraction and harmonization of multi-source data to enable AI-assisted optimization of dressing intervals. Information from various sources was aggregated, including material certificates (PDF/OCR), heat treatment and grinding fluid data (CSV), grinding wheel and machine data, and quality data (Q-DAS exports). Schemas were defined, and datasets were validated to address inconsistencies caused by manually entering batch numbers, wheel specifications, and inconsistent Barkhausen noise measurements. Due to internal data-governance constraints, data sharing and deployment were arranged within an enterprise-controlled environment, enabling ingestion, storage, and DT-enabled monitoring and analytics.

The resulting integrated dataset (quality, heat treatment, and workpiece material data) was primarily static, enabling construction of one feature vector per crankshaft and a tabular dataset for modeling. Predicting the preset dressing interval served as an exploratory baseline but provided limited insight because the target was largely fixed per machine. The target was therefore re-engineered into a dynamic interval derived from historical data, increasing when quality metrics were well within tolerance and decreasing as limits were approached.

The dataset was split into training (≈ 2500 samples) and test sets (≈ 1500 samples). Several ML algorithms were trained on the training set and evaluated on the test set. As shown in Table 1, the XGBoost Regressor produced the most accurate predictions for the optimized dressing interval, outperforming the other models due to its ability to capture nonlinear feature interactions in heterogeneous process variables.

Table 1. Performance comparison of machine learning models for predicting dressing intervals. Results are expressed as mean absolute percentage error (MAPE) and coefficient of determination (R^2) on the validation set.

Model	MAPE	R^2
HistGradientBoostingRegressor	0.08	0.94
Random Forest	0.15	0.81
XGBoost	0.02	0.99

5. Discussion and lessons learned

Based on the deployment experience and the identified gaps, several guidelines emerge for AI-IGC4.0-DT: (i) enhanced temporal learning, (ii) explainable AI (XAI) for trust building, (iii) organizational adoption barriers, and (iv) quicker model development via agile data flow for faster deployment.

Sequence models that capture short-term dynamics may be inadequate for long-term, cross-cycle dependencies typical in bearing grinding (e.g., warm-up transients, gradual thermal drift, wheel-wear progression, and stops/changeovers). Future work should therefore explore state-aware temporal models – such as TCNs and state-space sequence models (SSMs). Since grinding parameters are typically selected using models, a hybrid approach is desirable: temporal networks can learn residual corrections to physics-based predictions and/or enforce analytical constraints, thereby complementing established understanding of process mechanics with data-driven inference. For use in production, models should (i) report uncertainty and keep recommendations within approved limits; (ii) use identical feature processing at training and runtime (via platform pipelines) to avoid train/serve mismatch; and (iii) be evaluated with time-aware validation (train earlier, test later) that preserves event order and known process/machine changes [17].

In the AI-IGC4.0-DT framework, learning is understood as inferring models that can plausibly explain observed data and support decision-making under uncertainty. Since many AI/ML models can fit the same observations – and both predicted outputs and the effects of actions are uncertain – explainability is essential for industry acceptance. The black-box nature of AI/ML models motivates the integration of explainable AI (XAI) to make model reasoning transparent to operators and process planners, as well as to map explanations to permissible control actions [18]. At the global level, Shapley Additive Explanations (SHAP) quantify how much each feature or family of features (e.g., temperature, wheel wear, drive signals) contributes to predictions across (process) operating windows, supporting verification of physical plausibility. At the local level, Local Interpretable Model-agnostic Explanations (LIME) provide cycle-specific attributions and simple what-if probes (local sensitivity checks) for individual recommendations.

Adoption barriers in real-world manufacturing stem less from algorithmic capability than from unclear objectives and insufficiently defined use cases, limited understanding of the data required to train ML models, and the significant investment needed to prepare the necessary data. This amplifies organizational and integration challenges. When the goals are vague, the necessary data, model scope, and OT/IT interfaces increase uncontrollably, placing an additional burden on legacy systems and operations. Furthermore, a significant constraint lies in effectively collaborating with external partners, specifically AI specialists and machine vendors who control machine data and control systems. The lack of in-house AI expertise widens the skills gap, increases dependence on vendors, and deepens mistrust of black-box models whose outputs cannot be linked to process understanding. Successful adoption therefore requires treating use-case definition as the key design step and data as the key ingredient, followed by a phased, governance-led deployment [19] with standardized interfaces, operational playbooks, and explainable models that safeguard process know-how.

A structured data strategy supported the development of the AI-IGC4.0-DT platform and facilitated its adoption. Standardized ingestion interfaces, dedicated pipelines for data, and a flexible storage layer tailored to use-case needs reduced

the data preparation overhead. A simplified DT interface usable by domain experts supported practical steps such as data gathering, quality assessment, and final preparation for model training. The on-premises deployment option keeps data and models close to machines, meeting low-latency requirements for time-sensitive monitoring and decision support while maintaining data within segmented OT networks to satisfy cybersecurity constraints. The platform abstracts data complexity at ingress and standardizes deployment at egress, reducing one-off engineering and shortening time-to-production. This on-premises mode enables robust data governance and access control, including role-based permissions and fine-grained, auditable sharing with external partners, eliminating the need for bulk exports. Crucially, stream-level control implements data minimization by activating only the signals required for a given use case, which lowers transfer and storage volume, reduces network load and serving latency, narrows the exposure surface, and allows streams to be enabled/disabled dynamically as needs change. In practice, these mechanisms support faster preparation and deployment while maintaining compliance with internal policies and regulatory requirements.

6. Conclusions

This work was concerned with the development of an event-driven digital platform for grinding (AI-IGC4.0-DT), with dual data pipelines (HFRT streaming, and large-file/REST), process-instance linkage, and a uniform AI/ML serving layer – deployable as DTaaS or DTaaSP under industrial cybersecurity constraints. Initial demonstrations established feasibility of addressing heterogeneity and interoperability issues in legacy manufacturing environments at both pilot sites. The SKF pilot demonstrated that reframing size control to predict deviation rather than absolute position enables high-precision runtime adaptation under known conditions. The Volvo Trucks pilot showed that reengineering fixed dressing intervals into dynamic, quality-aware targets yields accurate predictions using boosting-based ML models. The main contributions are the deployable architecture and the demonstration that governed data linkage, rather than abundance of sensors, unlocks industrial AI value in grinding. These results suggest a practical approach to transitioning from monitoring to advisory analytics and provide opportunities to enhance grinding efficiency and sustainability.

Acknowledgements

This research was supported by EIT Manufacturing under project 24303, “AI-Driven Digital Twins in Grinding 4.0.” The authors thank Magnus Persson, Hanna Turresson and Jonas Svensson at Volvo Group, Hans-Olof Stålfelt at UVA Lidköping, and Christian Wrobel at Tyrolit – Schleifmittelwerke Swarovski for their valuable collaboration.

References

- [1] P. Krajnik, K. Wegener, T. Bergs, and A. J. Shih, “Advances in modeling of fixed-abrasive processes,” *CIRP Annals*, vol. 73, no. 2, pp. 589–614, 2024, doi: 10.1016/j.cirp.2024.05.001.
- [2] T. Bergs, D. Biermann, K. Erkorkmaz, and R. M’Saoubi, “Digital twins for cutting processes,” *CIRP Annals*, vol. 72, no. 2, pp. 541–567, 2023, doi: 10.1016/j.cirp.2023.05.006.
- [3] B. Denkena, H. Klemme, and D. Stoppel, “Condition monitoring of grinding wheels: Potential of internal control signals,” *Production Engineering*, vol. 19, pp. 65–75, 2025, doi: 10.1007/s11740-024-01295-x.
- [4] J. Jia et al., “Applications of Acoustic Emission Monitoring in Grinding: A Review,” *International Journal of Precision Engineering and Manufacturing-Green Technology*, vol. 12, no. 6, pp. 1893–1918, 2025, doi: 10.1007/s40684-025-00749-6.
- [5] M. Wahlgård, P. Krajnik, and J. Stahre, “A Scalable Edge Computing Solution for Real-Time Process Monitoring and Optimization in Industrial IoT,” *Procedia CIRP*, vol. 134, pp. 903–908, 2025, doi: 10.1016/j.procir.2025.02.219.
- [6] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, “Digital Twin in manufacturing: A categorical literature review and classification,” *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 1016–1022, 2018, doi: 10.1016/j.ifacol.2018.08.474.
- [7] J. F. Uhlenkamp, J. B. Hauge, E. Broda, M. Lutjen, M. Freitag, and K. D. Thoben, “Digital Twins: A Maturity Model for Their Classification and Evaluation,” *IEEE Access*, vol. 10, pp. 69605–69635, 2022, doi: 10.1109/ACCESS.2022.3186353.
- [8] K. Wegener et al., “Intelligent Grinding,” *International Journal of Automation Technology*, vol. 20, no. 1, pp. 57–77, 2026, doi: 10.20965/ijat.2026.p0057.
- [9] T. Kaufmann, S. Sahay, P. Niemietz, D. Trauth, W. Maaß, and T. Bergs, “AI-based Framework for Deep Learning Applications in Grinding,” in *2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMII)*, pp. 195–200, 2020, doi: 10.1109/SAMI48414.2020.9108743.
- [10] J. Zhang, C. Deng, P. Zheng, X. Xu, and Z. Ma, “Development of an edge computing-based cyber-physical machine tool,” *Robotics and Computer-Integrated Manufacturing*, vol. 67, pp. 102042, 2021, doi: <https://doi.org/10.1016/j.rcim.2020.102042>.
- [11] T. Kaufmann, J. Stanke, D. Trauth, and T. Bergs, “Conceptual Approach for Process Monitoring Based on an Assistance System for Grinding,” in *Advances in Production Research*, WGP 2018, Springer, pp. 63–71, 2019, doi: 10.1007/978-3-030-03451-1_7.
- [12] J. Mayr et al., “Thermal issues in machine tools,” *CIRP Annals*, vol. 61, no. 2, pp. 771–791, 2012, doi: 10.1016/j.cirp.2012.05.008.
- [13] L. Weng, T. Kizaki, C. Ma, W. Gao, and D. Kono, “A review of robust thermal error reduction of machine tools,” *International Journal of Machine Tools and Manufacture*, vol. 209, pp. 104298, 2025, doi: 10.1016/j.ijmachtools.2025.104298.
- [14] R. Dražumerić, R. Roininen, J. Badger, and P. Krajnik, “Temperature-based method for determination of feed increments in crankshaft grinding,” *Journal of Materials Processing Technology*, vol. 259, pp. 228–234, 2018, doi: 10.1016/j.jmatprotec.2018.04.032.
- [15] P. Krajnik, F. Hashimoto, B. Karpuschewski, E. J. da Silva, and D. Axinte, “Grinding and fine finishing of future automotive powertrain components,” *CIRP Annals*, vol. 70, no. 2, pp. 589–610, 2021, doi: 10.1016/j.cirp.2021.05.002.
- [16] P. Hoier et al., “Influence of batch-to-batch material variations on grindability of a medium-carbon steel,” *Journal of Manufacturing Processes*, vol. 73, pp. 463–470, 2022, doi: 10.1016/j.jmapro.2021.11.012.
- [17] D. Liu, Z. Liu, B. Wang, Q. Song, H. Wang, and L. Zhang, “Leveraging artificial intelligence for real-time indirect tool condition monitoring: From theoretical and technological progress to industrial applications,” *International Journal of Machine Tools and Manufacture*, vol. 202, pp. 104209, 2024, doi: 10.1016/j.ijmachtools.2024.104209.
- [18] N. Anwer, R. Stark, F. Tao, and J. A. Erkoyuncu, “Developing and leveraging digital twins in engineering design,” *CIRP Annals*, vol. 74, no. 2, pp. 843–868, 2025, doi: 10.1016/j.cirp.2025.05.002.
- [19] S. Chatterjee, N. P. Rana, Y. K. Dwivedi, and A. M. Baabdullah, “Understanding AI adoption in manufacturing and production firms using an integrated TAM-TOE model,” *Technological Forecasting and Social Change*, vol. 170, pp. 120880, 2021, doi: 10.1016/j.techfore.2021.120880.

[1] P. Krajnik, K. Wegener, T. Bergs, and A. J. Shih, “Advances in modeling of fixed-abrasive processes,” *CIRP Annals*, vol. 73, no. 2,