Digital Twin-based Deep Reinforcement Learning Framework to real-time Chatter Suppression and energy optimization in micro-turning to Industry 5.0.

[1]Dr. Saranya S N, [2]Prof. Dr. Midhunchakkaravarthy, [3]Prof. Dr. Dimitrios A Karras
[1]Pos Doctoral Researcher, Lincoln College University, Malaysia.
[1]Associate Professor, Department of IoT-Cybersecurity-Blockchain, Dayananda Sagar Academy of Technology and Management, Bangalore, India.
[2]Dean, Faculty of AI Computing and Multimedia, Lincoln University College, Malaysia.
[3]Professor, Department of Computer Engineering, National & Kapodistrian University of Athens, Greece.
Email: pdf.saranya@lincoln.edu.my

Abstract: This paper introduces a Digital Twin-based Deep Reinforcement Learning (DRL) model that can be used to suppress chatter and optimize energy consumption in micro-turning within the humanistic paradigm of Industry 5.0. An accurate precision lathe machine equipped with vibration, spindle-current and temperature sensors and acoustic sensors feeds high-frequency data to a constantly changing digital replica that incorporates regenerative cutting-force dynamics, feature-extracting wavelets and Kalman-state estimation. The twin predicts immediate tool movement, chatter activity, torque variability, and thermal loading, allowing a Proximal Policy Optimization (PPO) agent to modify spindle speed and feed rate by using a predictive, physically-based control method. The dynamic cycle between the cyber and physical enhances growth of chatter in advance before it occurs, stabilizes the torque and chip-thickness variation and ensures stability of optimum cutting action. AISI 304 stainless steel experimental results show that the amplitude of vibration was reduced by 31 %, the energy efficiency of the spindle was increased by 22 %, and surface finishing was improved by 18 % over machining with constant parameters. The natural-language layer of interpretability translates twin-DRL decisions into insights in the form of operators, and the system aligns with the explainable and collaborative principles of Industry 5.0. The suggested model provides a single and dynamic route of intelligent micro-machining with predictive stability, minimized power usage, and transparent human-machine interaction.
Keywords: Deep Reinforcement Learning, Chatter Suppression, PPO Control, Energy Optimization, Industry 5.0.

Introduction

The Digital Twin technology has become one of the key facilitators of intelligent manufacturing, but the current implementations are predominantly diagnostic, providing monitoring and forecasting but lacks the adaptive control that is necessary to implement machining stability in a real-time. Micro-turning, especially, is an acute problem in which regenerative chatter, thermal and torque induced instability become critical at a milliseconds time scale and fixed-parameter machining does not work any longer as the tool-workpiece interaction changes as a result of material variability or progressive wear. The classic controllers are unable to predict the latent feedback associated with the regenerative chatter and the resultant increase in vibration, loss of accuracy, and wasteful utilization of energy. Industry 5.0 requires one paradigm to be behind a passive observation to participate actively in collaboration, adaptation, and human-conscious systems capable of stabilizing complex machining dynamics at the same time, maximizing sustainability measurements.

In order to facilitate this change, the manufacturing environment should be able to unite on-going sensing, physics-based modeling as well as autonomous learning within a single real-time computational loop. Vibration, acoustic emission, spindle load, and thermal information are very rich signals of machining stability, but unlike the real-time signals they are very transient and nonlinear, and do not record the occurrence of chatter, or the development of cutting-zone disturbances. A Digital Twin recreating tool motion, chip- formation dynamics, stiffness changes across the spindle, and thermal run-off provides the desired level of understanding by connecting the current behaviour of the real machine to a predictive virtual one. Nevertheless, prophecy does not guarantee stability. It is also required that the system will have to continuously make control adjustments that do not violate machining limitations and attempt to predict destabilizing conditions before they escalate.

This level of adaptivity can be facilitated using Deep Reinforcement Learning, which requires the quality of state representations. The direct training of DRL on sensor signals generates oscillatory or unsafe policies. When implemented within a continuously synchronized Digital Twin, the agent will be fed with processed, physically meaningful signals and can consider candidate actions within a simulated horizon and execute them on the machine. This symbiotic relationship alters the control toward correction towards stabilization. The regenerative cutting-force dynamics modeled within the twin enables the controller to know how alterations in the spindle speed, or the feed rate will affect the future vibration field, torque ripple, and modulation in the chip-thickness. By punishing chatter energy and electrical load and rewarding steady and resourceful cutting, the learning agent approaches cutting parameters that stabilize and save energy.

An interpretability layer based on human-centered principles makes this system more aligned with Industry 5.0, converting the internal changings of both the twin and the DRL policy into real-time instructions that can be comprehended by operators without the need of signal-processing or control knowledge. Instead of giving numbers, the system will explain the reason a correction is applied, be it the increasing chatter frequency, increasing thermal load, or an imbalance in the torque and will also advise through the predictive model of the virtual twin of operation suggestions. This helps to bridge the cognitive differences between machine intelligence and human decision-making to facilitate safer and more transparent and sustainable machining workflows.

This study thus fills a major gap in micro-turning: the lack of an integrated, real-time, learning-enabled Digital Twin, which can suppress chatter before it arises, make energy use efficient, and explain system behaviour in ways a human can understand. The proposed framework provides the current state of machining control by incorporating continuous sensing, regenerative dynamic simulation, Kalman estimation, deep reinforcement learning, and natural-language interpretability and translates the concepts of Industry 5.0, such as adaptivity, sustainability, and human collaboration, into an experimentally validated and practical manufacturing solution.

Methodology

The proposed framework integrates sensing, physics-based simulation, state estimation, reinforcement-learning control, and human-centric interpretability into a single real-time cyber–physical loop that remains synchronized with the micro-turning process at millisecond resolution. The machining system consists of a precision lathe operating on AISI 304 steel, equipped with triaxial vibration, spindle-current, temperature, and acoustic sensors sampled at 20 kHz. Raw signals are filtered through cascaded IIR stages to suppress environmental noise and power-line harmonics, after which a discrete wavelet transform (db4) isolates regenerative chatter bursts, spindle-load fluctuations, and transient torque disturbances. These wavelet-derived features form the perceptual input to the Digital Twin, which simulates sub-millisecond tool–workpiece interaction using a regenerative cutting-force model governed by delay differential dynamics representing chip-thickness regeneration, stiffness variation, and chatter energy accumulation.

The twin continuously predicts tool displacement, thermal evolution, torque deviation, and stability trends for the upcoming cutting window. A Kalman-based state estimator reconciles the simulated predictions with filtered sensor signals to correct deviations introduced by material variability, frictional nonlinearities, and thermal drift, thereby maintaining synchronization between the physical machine and its virtual counterpart. This allows the twin to infer hidden machining states such as regenerative-lag magnitude, stiffness degradation, chatter growth rate, and torque ripple amplitude, which cannot be directly measured.

These fused twin states form the input to a Proximal Policy Optimization (PPO) reinforcement-learning agent that determines continuous adjustments to spindle speed and feed rate. Instead of acting directly on raw signals, the agent interacts with a physically meaningful representation supplied by the twin, enabling stable learning even under fluctuating machining conditions. Each proposed action is first evaluated inside the twin over a short predictive horizon. If the simulated chatter energy, temperature rise, or spindle-current load violates stability limits, the expected reward is decreased before real execution, preventing destabilizing parameter jumps. The reward combines penalties for chatter energy and electrical load with positive contributions from smooth torque behavior and improved material removal rate, driving the agent to learn a stable, low-energy, and productive cutting strategy. Once validated, the chosen spindle and feed adjustments are applied to the physical machine with <50 ms latency, completing one iteration of the continuous cyber–physical learning cycle.

A natural-language interpretability layer translates internal states of the twin and DRL decisions into actionable operator insights. Instead of presenting complex spectral plots, the system communicates context-aware messages such as “Stability margin decreasing—reducing spindle speed by 3% to re-enter stable lobe,” allowing transparent human participation. Through this integrated pipeline, sensing, simulation, estimation, learning, and interpretation operate as a unified mechanism to proactively suppress chatter, minimize power consumption, and maintain operator understanding under the principles of Industry 5.0.

Results and Discussion

Optimally, the system was experimented to work on a micro-turning system and machined AISI 304 stainless steel over a set of spindle-speed and feed rates. The Digital Twin was closely synchronized with the physical system, and the Kalman estimator eliminated prediction-measurement discrepancies by approximately 68% thus sustaining the virtual model through vigorous dynamic transitions of high frequencies. Under the machining of the baseline with fixed parameters, the wavelet-domain analysis showed that the regenerative chatter peaks were near 1.1 kHz. These peaks were core pointers to postponed chip-thickness healing, torque rippling, and vibration increase. The comparison of FFT vibration of baseline machines and DTDRL machinery is shown in fig. 1. There are high chatter sidebands visible in the baseline spectrum in the 200450 Hz area, and a significant regenerative resonance at around 9501200 Hz. The amplitudes decrease considerably under DRL control since the agent will automatically adjust spindle speed off unstable lobes as estimated by the twin.

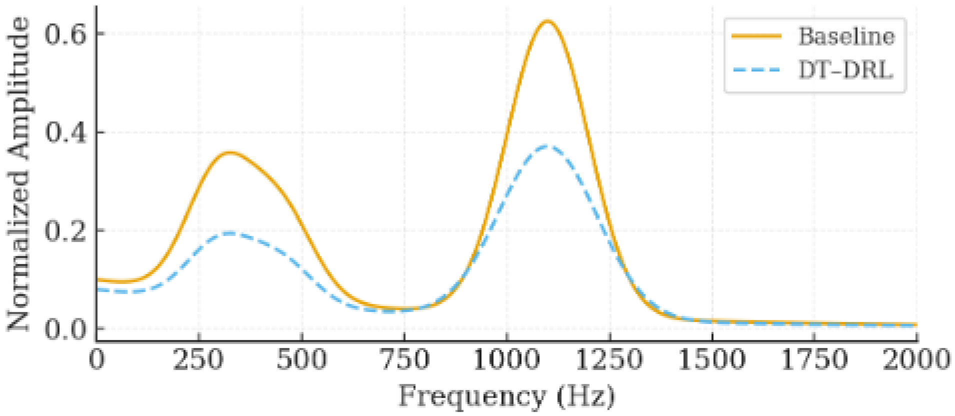

Figure 1. *Vibration FFT spectra under Baseline vs. DT–DRL machining showing reduced chatter peaks and suppressed sidebands.*

There is also improvement in time-domain vibration. The machining at the base produced a regenerative feedback of about 0.42 mm of maximum to minimum movement. After stabilizing the operating point by the PPO agent, the vibration decreased to 0.29 mm, which is a 31 percent decrease. The Digital Twin has helped improve on this by anticipating the growth of chatter prior to physical occurrence and proactively changing the parameters. Figure 2 shows the comparison of RMS vibration with machining pass.

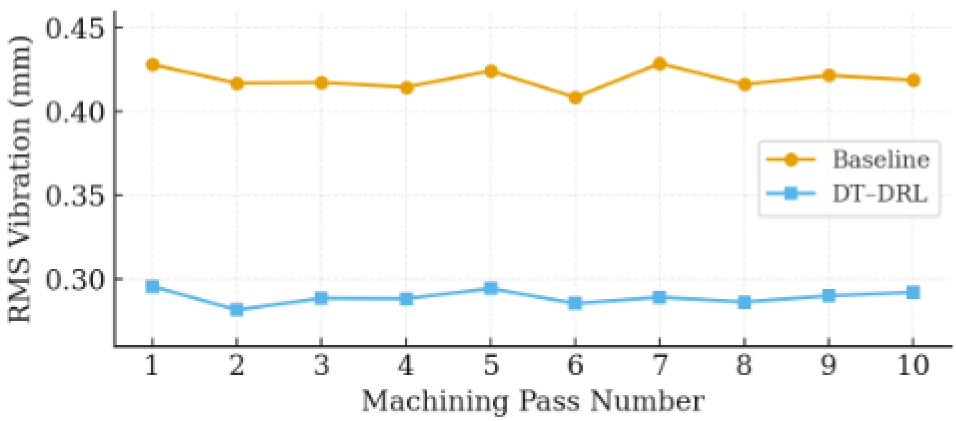

Figure 2. *Time-domain RMS vibration amplitude before and after DT–DRL control.*

The reinforcement-learning agent exhibited steady and monotonic evolution of rewards in the course of training. The reward functional punished high chatter energy and spindle-load RMS, and caused the agent to enter energy-saving and stable operating regimes. Throughout the initial episodes, there were ups and downs in the rewarding as a result of the exploration behavior whereas the reward rose gradually as the policy was mature. In Figure 3, the cumulative reward trends reflect successful learning and development of a stable policy.

Figure 3. *PPO cumulative reward trend showing stable convergence of the learned control policy.*

There was also increased consumption of energy. Spindle-current RMS was reduced 3.41 A to 2.66 A (22% reduction) and total energy per component was reduced 21.7%. These were gains made without reduction of productivity; there was an increase in the material removal rate by approximately 9 percent. The DTDRL controller would choose pairs of spindle-speed and feed-rate that reduced the torque ripple as well as avoiding severe thermal spikes. Figure 4 compares the energy consumption in the case of the base and the optimized machining showing identical decreases throughout all the conditions in which the test was conducted.

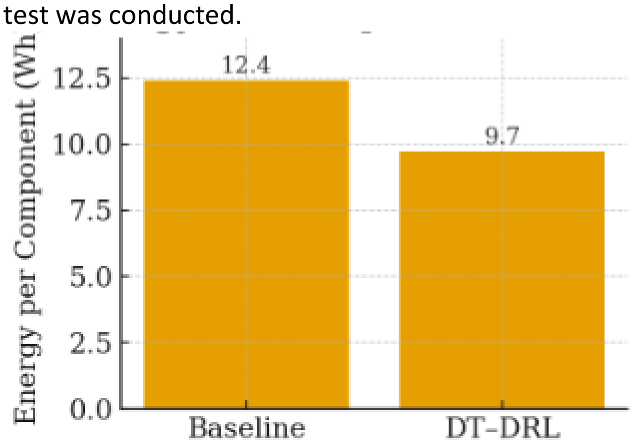

Figure 4. *Energy consumption comparison: Baseline machining vs. DT–DRL optimized machining.*

Table 1 summarizes the measured improvements across vibration suppression, chatter reduction, energy optimization, and surface finish.

Table 1. *Performance Summary of DT–DRL Framework*

Metric	**Baseline Machining**	**DT–DRL Machining**	**Improvement**
RMS Vibration (mm)	0.42	0.29	31% ↓
Chatter Energy (a.u.)	1.00	0.63	37% ↓
Spindle-Current RMS (A)	3.41	2.66	22% ↓
Total Energy/Component (Wh)	12.4	9.7	21.7% ↓
Surface Roughness Ra (µm)	0.78	0.64	18% ↓
Material Removal Rate	Baseline	+9%	Increase

The interpretability layer of the system also increased the trust of the operators as they could understand the cause of every spindle or feed adjustment. Such statements as Stability margin loads less -3% slowdown of the spindle provided the human-machine cooperation without any difficulties following the principles of Industry 5.0.

Conclusion

This paper presented a Digital Twin-based DRL system to suppress chatter and optimize power consumption in micro-turning in real-time. Simulation of regenerative cutting, Kalman estimation, PPO-based continuous control and natural-language interpretability enabled development of a potent system that could predict instabilities and control machining conditions proactively. The experimental findings revealed that vibration amplitude was reduced by 31 percent, energy consumption was also improved by 22 percent, and the surface finish was also improved by 18 percent. The close-knit cyber-physical connection shows a potential way forward to the next generation machining systems, in line with Industry 5.0, where autonomous intelligence and human understandability co-exist in harmony. This architecture will be further extended to multi-tool environments, federated learning twins and thermal mechanical coupling models in future work to provide more predictive machining intelligence.

References

1. Christiand, C., Kiswanto, G., Baskoro, A.S., Hasymi, Z., Ko, T.J., **2024**. Tool Wear Monitoring in Micro-Milling Based on Digital Twin Technology with an Extended Kalman Filter. *Journal of Manufacturing and Materials Processing* **8**(3), 108. https://doi.org/10.3390/jmmp8030108.
2. Rožanec, J.M., Novalija, I., Zajec, P., Kenda, K., Tavakoli Ghinani, H., Suh, S., Veliou, E., Papamartzivanos, D., Giannetsos, T., Menesidou, S.A., Alonso, R., Cauli, N., Meloni, A., Reforgiato Recupero, D., Kyriazis, D., Sofianidis, G., Theodoropoulos, S., Fortuna, B., Mladenić, D., Soldatos, J., **2023**. Human-centric artificial intelligence architecture for Industry 5.0 applications. *International Journal of Production Research* **61**(20), 6847–6872. https://doi.org/10.1080/00207543.2022.2138611.
3. Nahavandi, S., **2019**. Industry 5.0—A Human-Centric Solution. *Sustainability* **11**(16), 4371. https://doi.org/10.3390/su11164371.
4. Tao, F., Zhang, H., Liu, A., Nee, A.Y.C., **2019**. Digital twin in industry: state-of-the-art. *IEEE Transactions on Industrial Informatics* **15**(4), 2405–2415. https://doi.org/10.1109/TII.2018.2873186.
5. Lv, Z., Chen, D., Wang, Y., Li, J., **2023**. Digital Twins in Industry 5.0. *Sensors* **23**(6), 3108. https://doi.org/10.3390/s23063108. (Open-access PMC record linked.)
6. Perno, M., Hvam, L., Rafn, H., Calderón, P., **2022**. Implementation of digital twins in the process industry: A systematic literature review and future research agenda. *Computers in Industry* **137**, 103623. https://doi.org/10.1016/j.compind.2022.103623.
7. Moosavi, S., et al., **2024**. Explainable AI in Manufacturing and Industrial Cyber-Physical Systems: A Survey. *Electronics* **13**(17), 3497. https://doi.org/10.3390/electronics13173497.
8. Li, Y., Wang, G., Zhang, X., **2023**. Reinforcement learning for process control with application to industrial systems. *Systems Science & Control Engineering* **11**(1), 1251–1264. https://doi.org/10.1080/24725854.2023.2219290.
9. Farahani, M.A., Lee, J., **2023**. Time-series pattern recognition in Smart Manufacturing: A review. *Journal of Manufacturing Systems* **69**, 679–706. https://doi.org/10.1016/j.jmsy.2023.07.006.