UNIVERSITY OF WEST ATTICA
eISSN: 3136-7658

Applied Manufacturing Intelligence Journal

Vol. 1, No. 1, January, 2022



Diffusion Policy Transformer Model for Contact-Dense Robotic Polishing Trajectory Generation

Saif Al-Shamsi[1,*], Omar Al-Hashimi[1], Marwan Al-Darmaki[1]

[1] *College of Engineering, Information Technology Division, University of Sharjah, Sharjah, PO Box 27272, United Arab Emirates*

Received 03 October 2021; Revised 14 December 2021; Accepted 27 December 2021; Published 04 January 2022

Abstract

Creating a Diffusion Policy Transformer Model for the Generation of Contact-Dense Robotic Polishing Trajectories. During metal and composite surface finishing, the goal is to achieve a steady polishing route while the tool position, contact force, surface curvature, abrasive wear, and vibration response all vary at the same time. Robot joint position, end-effector pose, force-torque data, spindle speed, vibration signals, local surface geometry, and roughness labels are first used to construct a polishing-state window. A Transformer encoder preserves the temporal link between contact state, tool attitude, and surface-removal reaction, while a conditional diffusion strategy learns the distribution of possible trajectory segments. Excessive force, abrupt tool angle changes, and low-speed sticking are penalised by a contact-constraint module. The final policy generates a short-horizon trajectory for robot execution and human evaluation. The simulated dataset contains 24,000 trajectory segments for each of the nine robot work stations, six material surfaces, five polishing heads, and four curvature groups. The suggested model has decreased the mean path deviation to 1.18 mm, lowered the force RMSE to 2.81 N, improved surface-roughness reduction to 27.9%, and maintained a 42 ms median inference delay when compared to dynamic movement primitives, behaviour cloning, standard Transformer policy, and diffusion policy without contact constraints. Transformer memory stabilises extended contact sequences, diffusion sampling enhances multimodal trajectory selection, and explicit contact limitations lessen force overshoot at edges and weld-seams, according to ablation analysis.

Keywords: Robotic Polishing; Contact-Dense Manipulation; Trajectory Generation; Diffusion Policy Transformer; Surface Finishing

1. Introduction

To decrease repeated manual labour and enhance process uniformity, robotic polishing is being utilised more often for high-precision metal shells of aerospace components, automobile moulds, medical implants, turbine blades, and other goods [1]. Stable physical interaction necessitates combined modulation of shape, force, friction, tool compliance, and local surface reaction, according to manipulation studies with rich contacts [2]. Because these trajectories cannot be stated by fixed rules alone when dealing with curved surfaces, weld seams, and local faults concurrently, learning-based robot control has become essential [3]. A good generative framework for generating multimodal robot behaviours from demonstrations and sensor settings is offered by diffusion-policy research [4]. In order to do work, polishing involves moving the tool smoothly over the work item. When removing defects, the robot should make excellent contact with the surface, steer clear of force spikes that might harm the workpiece, and modify its attitude when the tool comes into touch with an edge band, a concave surface, or a previously polished region.

Robotic grinding experiments [5] have demonstrated that these parameters are tightly connected to tool attitude, feed speed, and abrasive conditions. Conventional polishing automation typically employs offline path planning, manually sets contact-force thresholds, and does local tool-path correction. The techniques work well on flat surfaces,

* Corresponding author: saif.sha@sharjah.ac.ae

https://doi.org/10.68406/amij.2022.vol1iss1np1:1-13

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

but they lack the necessary precision in regions with significant surface curvature fluctuations or intensive contact. Impedance force-tracking study [6] suggests that interaction control take impact and contact stability into account because even though a trajectory is mathematically perfect, it can still fail if the abrasive tool loses touch, pushes too deeply, or changes angle too suddenly. The majority of human operators receive touch-and-sight compensation, and it is difficult to translate this information into deterministic rules. A straightforward behavior-cloning policy could average many legitimate acts and provide a trajectory that is safe in all contact states, but a data-driven policy can learn from robot demos and process measurements.

The issue with contact-dense trajectory creation is the large number of possible actions. Abolghasemi and Boloni [7] showed that end-to-end attentive visuomotor policies can learn robot manipulation actions directly from sensory observations in cluttered scenes, but the learned action should still be feasible in the presence of physical contact. Effective defect removal on a curved shell requires a somewhat slower tool path and a greater normal force; otherwise, a quick, light pass may damage the thin edge. Because of abrasive wear and prior material removal, the proper route at a weld seam may differ. Studies on imitation learning [8] have also demonstrated that there are several acceptable options for comparable observations in demonstrations. Because of variations in material hardness and contact curvature, the same measured force on a mixed batch will correspond to different surface conditions. A model that satisfies the force and smoothness conditions while representing multimodal action distributions is needed for these scenarios. Instead of treating robotic polishing as a single deterministic regression issue, the Diffusion Policy Transformer presented here treats it as a conditional trajectory generation problem.

The whole trajectory-generation route will comprise the remainder of this study. Transformer sequence modelling, diffusion-based policy learning, robotic polishing, and contact-rich manipulation are reviewed in Section 2. Contact limitations, the training target, the Diffusion Policy Transformer structure, and polishing-state building are all introduced in Section 3. The dataset distribution, trajectory accuracy, force stability, surface-quality enhancement, ablation analysis, resilience, and robot-cell deployment outcomes are shown in Section 4. The results, limitations, and future research directions are summarised in Section 5.

2. Related Work

2.1. Contact-Rich Robotic Machining

Because a tool must keep steady touch with the workpiece, robotic machining and polishing are not free-space robot movements. Research on adaptive polishing [9] demonstrates that when curved surfaces and local contact defects coexist, route correction is required. Only force feedback can verify that the route has been consistently material-removed, and force control is frequently employed in the last step of industrial robot finishing to correct for fixture mistakes and surface defects.

various feed and force settings are needed in various parts of the same surface, according to trajectory planning studies for complicated curved polishing [10]. Although surface evidence and motion production can be connected by learning-based tool-path research [11], contact stability must still be confirmed prior to execution. A trajectory model that incorporates force, position, surface evidence, and action feasibility is required in light of the findings.

2.2. Diffusion-Based Robot Policy Learning

Diffusion models may provide a wide range of outputs under the same conditions because they can learn to reverse a slow noise process. Generative action synthesis can preserve several viable behaviours rather than reducing them to a single average option, as demonstrated by Planning with Diffusion [12]. Robotic polishing requires multimodality because the robot may choose a slightly different tool angle or speed even after the contact and surface quality requirements are met.

A theoretical foundation for learning complicated distributions without requiring all outputs to be in a single mean action is provided by score-based generative models [13]. This feature works well for weld seams, curved patches, and edge areas where local movement is required. For contact-dense polishing, a diffusion policy by itself is not appropriate as the denoising procedure must be aware of the temporal contact state.

Even though a brief action segment could seem doable at one time step, it will become dangerous after a few steps of escalating normal force or cumulative vibration. Consequently, a sequence backbone and diffusion action creation are combined in this research. The policy circumstances can provide a trajectory that is both locally possible and compatible with the recent polishing history. These conditions include fresh force, position, vibration, spindle speed, surface geometry, roughness status, etc.

2.3. Transformer Sequence Modelling for Manipulation

Transformer models are adept at handling sequence representation and using attention to link distant time steps without the need for recurrent memory. Robot and control activities can be represented as conditional sequences, according to Decision Transformer research [14]. Attention may combine observation and action history for manipulation, as demonstrated later by large-scale robot learning [15]. Polishing requires temporal memory because, while an increase in contact force may be acceptable at the start of a pass, cumulative vibration makes it harmful.

Accurate action prediction is not enough for industrial deployment. According to research on explainable artificial intelligence [16], a recommendation's supporting evidence should be presented in the decision-support output. Before applied to a safety-critical system, learned actions should be constrained by operational limitations, according to safe learning-based control [17]. As a result, the model must determine if the tool has just entered a curved region, crossed an edge, or returned to a previously polished portion.

Because the produced polishing segment must be updated often rather than being implemented as a lengthy open-loop command, model-predictive control research [18] is pertinent. If there is no force, vibration, or surface evidence, uncertainty quantification [19] is also necessary. The receding-horizon form used later in the technique section is supported by the concepts.

Maintaining the learnt policy's compatibility with controller restrictions, tool motion, and shop-floor review is supported by general robotics application studies [20]. The suggested structure is motivated by these studies: explicit penalties are included for force stability and smooth motion, Transformer attention is utilised to manage contact history, and diffusion is used for multimodal trajectory creation.

3. Methodology

3.1. Contact-Dense Polishing Data Construction

The simulated polishing dataset contains 24,000 trajectory segments from the nine robot workstations. Aluminium alloy plates, stainless steel shells, composite covers, titanium curved panels, weld-seam strips, and coated test coupons are the workpieces. Each segment consists of a 1.2-second polishing window with 60 synchronous robot-control frames captured at 50 Hz. Six-axis force-torque measurements, six robot joint locations, six joint velocities, end-effector posture, spindle speed, tool feed speed, vibration amplitude, local surface-normal direction, curvature group, starting roughness, and post-pass roughness label are among the data that are captured. The robot, surface, force, vibration, and roughness signals are transformed into a trajectory-generation state in Figure 1, which depicts the contact-dense polishing data and control workflow.

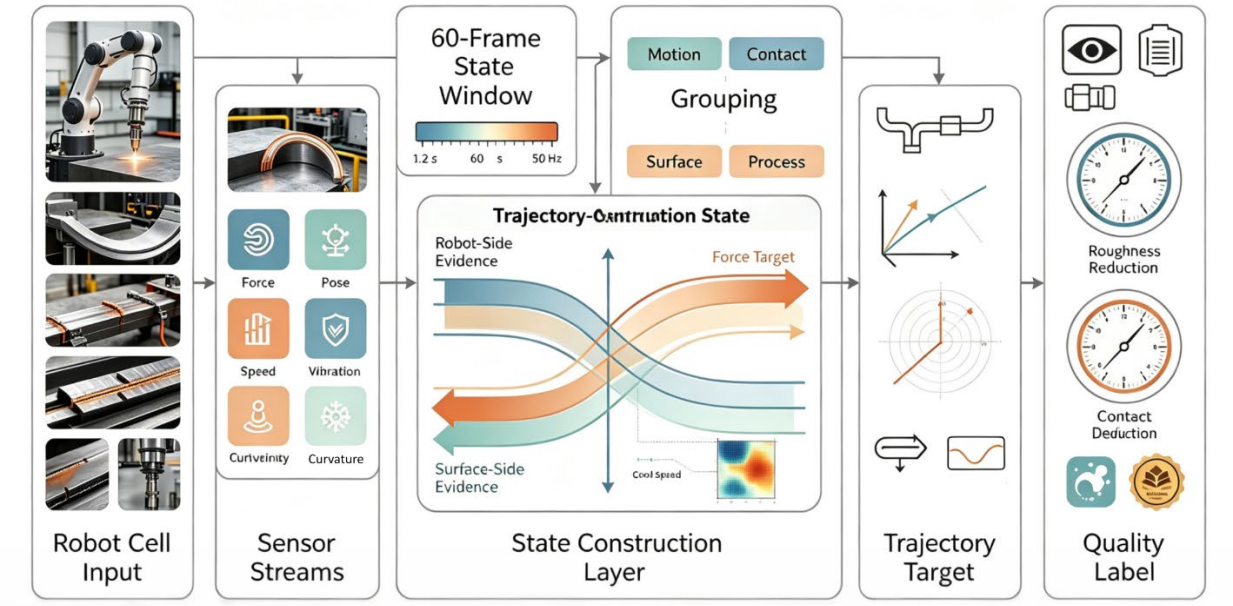


Figure 1. Contact-dense polishing data and control workflow.

A polishing state window is constructed to represent recent motion and contact behaviour. The window length is selected as 60 frames because it covers the contact response before and after the robot crosses most local defect patches in the simulated cell. The state tensor is written as

$$S_t = \{p_t, f_t, v_t, q_t\} \quad (1)$$

where p_t denotes tool pose, f_t denotes force-torque measurements, v_t denotes vibration response, and q_t denotes surface-quality cues. The representation contains both robot-side and surface-side evidence, which prevents the policy from treating a force change as meaningful without checking the local geometry. Contact deviation is measured by comparing the observed normal force with the desired force interval for the current material group:

$$c_t = \left|F_t - F_{ref}\right| \tag{2}$$

The surface-quality label is built from the roughness reduction after a polishing pass. In the simulation, the target roughness reduction is 25% for flat and curved surfaces and 20% for edge or weld-seam surfaces because these regions have stricter force-safety limits. The quality change is represented as

$$r_t = R_{\text{pre}} - R_{\text{post}} \tag{3}$$

Batches of workpieces are separated from the Data Set, not at random. Two workstations with unknown batch combinations are utilised for testing; one workstation and one material batch are used for validation; and six workstations and four material batches are used for training. The policy's suitability for a new surface shape and contact arrangement can be ascertained by split testing. There are 16,800 trajectory segments in the training set, 3,600 in the validation set, and 3,600 in the test set. All of the windows now include contact-density labels. The percentages of windows with low contact density are 34.5%, those with medium contact density are 41.2%, and those with high contact density are 24.3%. Because high-density windows are typically found around the edges, weld seams, and curved transition sections, feed rate, normal force, and tool orientation must be coordinated more precisely.

For the purpose of training the model, group the input characteristics. Features of motion include feed speed, joint velocity, and tool posture. Normal force, tangential force, torque, and force-rate change are contact characteristics. Curvature, local normal, roughness label, and defect category are examples of surface characteristics. Spindle speed, abrasive type, and predicted tool wear level are process characteristics. The action target is a brief future tool-trajectory segment, and the model uses these grouped characteristics as condition tokens. Displacement, orientation correction, feed-speed adjustment, and contact-force goal are all included in the created segment. Rather than only forecasting a scalar polishing quality score, this architecture can produce a useful robot command for the policy.

3.2. Diffusion Policy Transformer Model

Figure 2 is the denoising policy Transformer for robot polishing trajectory generation. The three modules of the model are: a condition encoder for polishing states, a diffusion policy head for action denoising, and a contact-constraint evaluator for trajectory screening. The condition encoder receives the most recent motion, contact, surface and process tokens. The diffusion head starts with a noisy trajectory segment and gradually recovers a feasible action sequence. The contact evaluator will determine if the reconstructed path exceeds the force limit, angle limit or smoothness limit, and then notify the robot-cell controller of this.

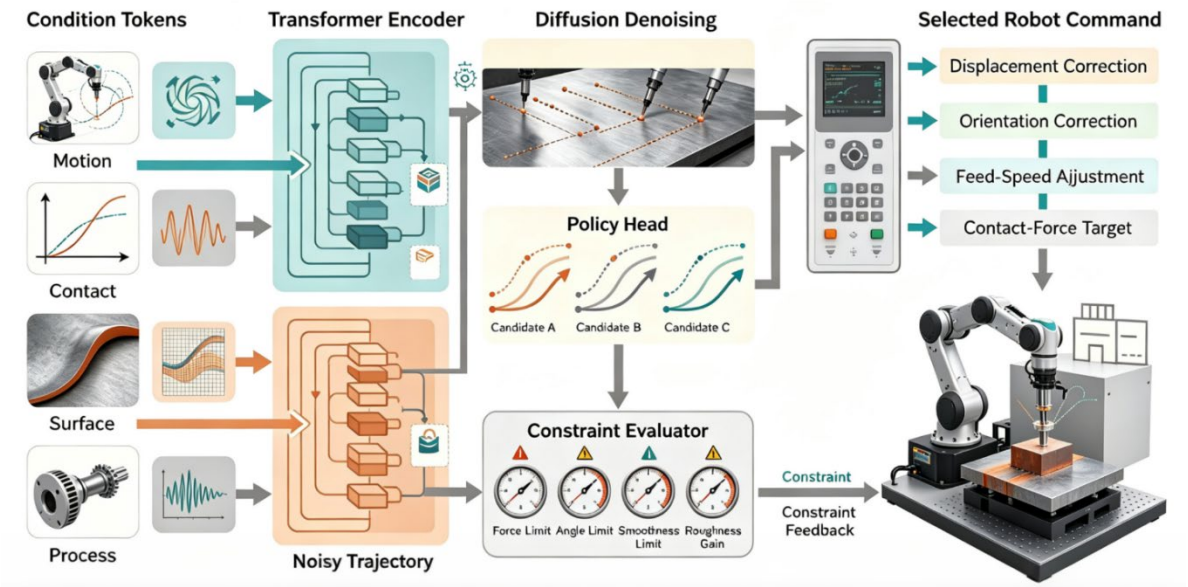


Figure 2. Denoising Policy Transformer for robot polishing trajectory generation.

The forward diffusion process adds controlled noise to the demonstration trajectory. This allows the model to learn how feasible polishing actions are distributed around a given state condition. The noisy trajectory is written as

$$a_n = a_0 + \sigma_n \epsilon \tag{4}$$

where a_0 is the demonstrated trajectory segment, σ_n is the diffusion noise level, and ϵ is sampled noise. The Transformer encoder constructs condition tokens from the recent polishing state. A compact token representation is used so that the policy can run within the robot-cell cycle time:

$$z_t = T(S_t) \tag{5}$$

The denoising head predicts a clean action from noisy action and condition tokens. The predicted action is written as

$$\hat{a}_0 = D(a_n, z_t) \tag{6}$$

The model outputs a 12-step future command, and at a 50Hz control rate, this is 0.24s of robot motion. The horizon is long enough to smooth out the polishing pass, but not too long; otherwise, the plan near the edge or where a force increases suddenly will be uneven. Sample uniformly from the 20 noise levels in the diffusion step during training. The Transformer has 4 attention layers, 8 attention heads and a 256-dimensional hidden size. The above values are chosen to keep the inference time below 50ms for the simulated robot-cell controller and retain sufficient temporal context for contact-rich operation.

The generated action is not directly accepted. The model keeps the top three denoised candidates during inference and ranks them based on contact stability, predicted roughness reduction and trajectory smoothness. If the best candidate exceeds the force limit, the policy will select the next feasible candidate or return a review flag. Add a candidate ranking step to improve accuracy; thus, the two trajectories may have similar geometric errors but different force overshoots. A generative model is only useful if the final command can be filtered by engineering constraints before execution.

3.3. Trajectory Constraint and Objective Function

The trajectory constraint module evaluates whether a generated segment is suitable for dense contact. The first constraint limits force error. It prevents the tool from pressing too hard on thin edges or losing contact on curved surfaces. The force penalty is

$$L_f = \text{mean}(c_t) \tag{7}$$

The second constraint limits sudden trajectory change. Abrupt changes may produce vibration marks, uneven removal, or controller tracking error. The smoothness term is

$$L_s = \text{mean}(|a_t - a_{t-1}|) \tag{8}$$

The third constraint links the predicted trajectory with expected surface improvement. The model should not produce a very safe but ineffective pass when the surface still contains removable roughness. This term compares measured roughness reduction with predicted roughness response in the validation window, and it is kept as part of the final objective instead of being written as an additional standalone formula.

The final objective combines diffusion denoising, force stability, smoothness, and surface-response consistency:

$$L = L_d + \alpha L_f + \beta L_s + \gamma L_r \tag{9}$$

The coefficients are set to $\alpha = 0.45, \beta = 0.20$, and $\gamma = 0.35$ after validation. These values give the largest weight to contact stability because force overshoot creates the highest engineering risk in edge and weld-seam polishing. The diffusion denoising term preserves the demonstrated action distribution, while the constraint terms prevent the policy from generating a visually smooth but physically unsafe trajectory. The model is trained for 80 epochs with a batch size of 128 trajectory windows. Early stopping is triggered if validation force RMSE does not improve for eight consecutive epochs.

For the next twelve stages, the output trajectory includes target feed speed, target normal force, displacement correction, and attitude correction. In order to create a receding-horizon operating mode, the robot controller only completes the first four steps before requesting a fresh forecast. Thus, by adjusting to fresh force and vibration data, the cumulative error is decreased. The same output, which contains the chosen candidate, rejected candidates, contact score, anticipated roughness improvement, and review flag, is also preserved as an audit record. In a challenging contact area, this record can be used to explain why the robot slowed down, decreased the typical force, or asked for manual confirmation.

4. Results and Discussion

4.1. Data Distribution and Polishing Scenario Analysis

The distribution of surface-defect class, contact-density level, and material type in the simulated dataset is displayed in Figure 3. Aluminium alloy and stainless-steel samples comprise 24.8% and 22.6% of the windows, respectively, as seen in Figure 3(a); titanium curved panels and coated coupons are less common and hence more difficult to obtain. This distribution is typical of a general-purpose polishing cell; high-value curved materials require fewer but more

safety-critical passes, while the majority of the materials generate numerous stable passes. Contact density is displayed in Figure 3(b), with a medium-density window at 41.2% being the biggest group. Although they make up 24.3%, high-density windows are responsible for 46.7% of force-overshoot labels because to the rapid changes in edge and weld-seam contact. Scratch, orange-peel texture, weld ridge, coating burr, and local pit classes are all present, as seen in Figure 3(c). According to dataset-design research [21], the evaluation should encompass not just clean examples but also challenging negative and border situations. As a result, the figure demonstrates that the model has been evaluated in both contact-sensitive sections and regular polishing passes. Additionally, the picture illustrates why a single-surface category is inappropriate for evaluation. Although aluminium alloy windows have a lot of low-vibration routes, they may nevertheless have edge bands with very little material left. Although they need a greater normal force, stainless-steel shells reduce roughness more steadily. Because the local normal varies quickly, titanium curved panels produce a bigger tool-angle correction with fewer samples. Because weld-seam strips have the largest percentage of high-density contact windows, they are utilised to assess if the policy can reject dangerous candidates instead of only fitting the tool route that has been established. To stop excessive pressure along a small-deviation route from damaging the coating, coated coupons are added. As a result, the dataset is not a trajectory-regression dataset but rather a contact-control benchmark.

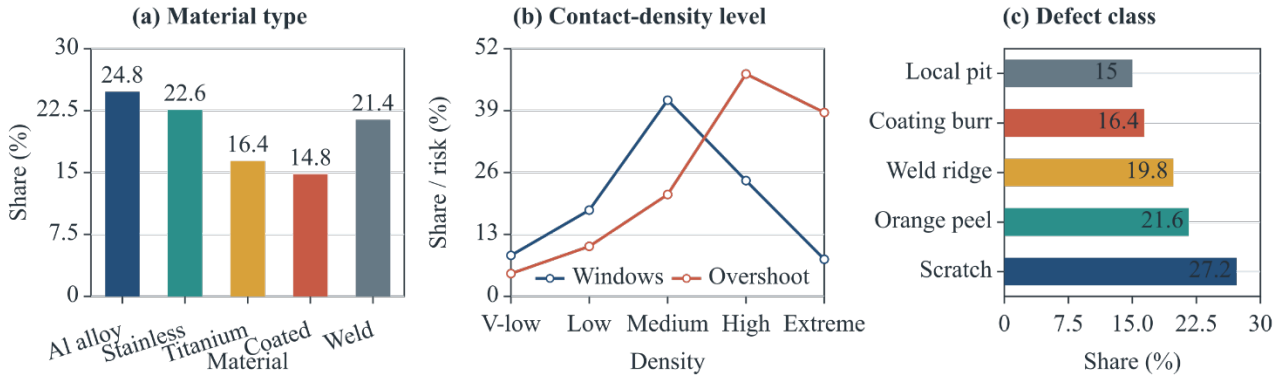


Figure 3. Dataset composition. (a) material type. (b) contact-density level. (c) defect class.

Two process-distribution perspectives are included in Figure 4. The contact-force distribution is shown in Figure 4(a); the 18–24 N range is mostly used for standard polishing, while the 30–36 N range is mostly found at weld-seam and edge windows. The distribution of trajectory lengths is shown in Figure 4(b); larger segments are more common on flat plates, and the majority of produced segments are between 12 and 28 mm long. These two distributions are taken into consideration simultaneously because robotic polishing research [22] have demonstrated that tool load and route length jointly impact surface smoothness. The hardest are short high-force windows since the controller has limited time to recover before surface damage happens. Long low-force windows are less dangerous, but the removal might not be enough if the tool loses contact. Therefore, rather than focusing just on geometric route correctness, the suggested model will be evaluated based on trajectory error, force error, surface improvement, warning output, etc. There is an uneven distribution of force. In a real-world polishing cell, safe windows are more common than harmful contact states; yet, the system's usability is determined by the risky states. There is also a relationship between trajectory-length distribution; a lengthy command segment may conceal force increase within the same execution window, whereas a short segment leads to frequent replanning. The 12-step horizon was chosen for this article out of the 8-, 12-, and 20-step possibilities. While the 20-step horizon is geometrically smoother but has greater overrun on the edges, the 8-step horizon has a lower latency but is more route fractured. The validation set will be subjected to a 12-fold cross-validation.

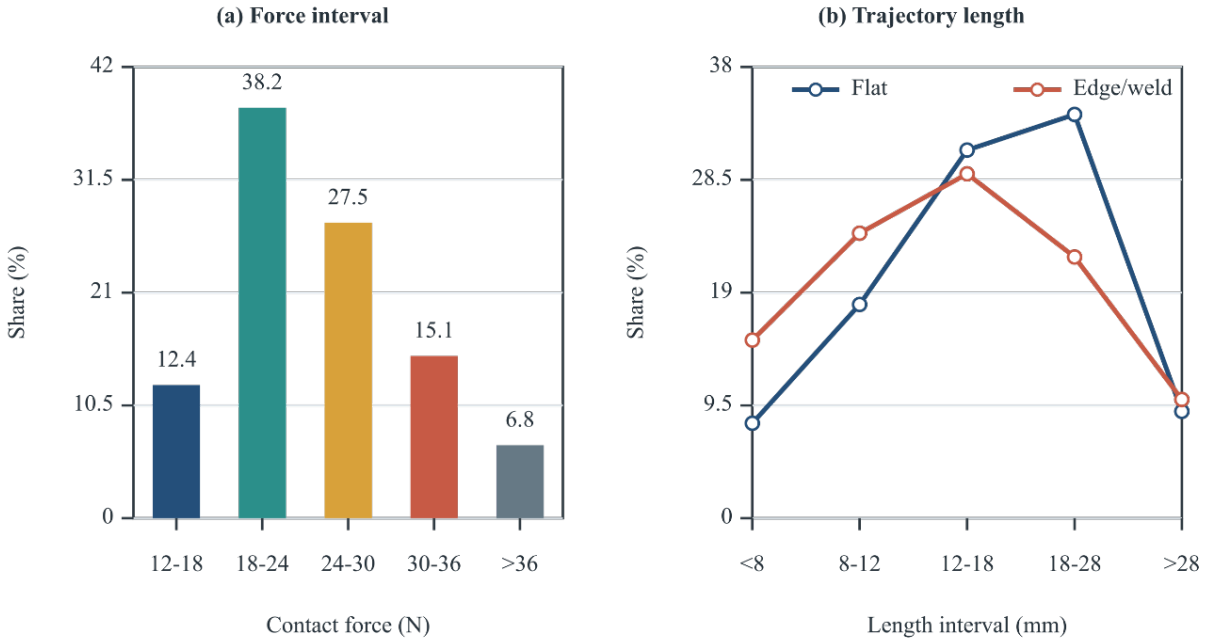


Figure 4. Contact process statistics. (a) force interval. (b) trajectory length.

Distribution also demonstrates the absence of random sampling. The model may have overestimated its capacity for generalisation if windows from the same workpiece and polishing pass show up in both training and testing. Because the unseen material batches have varying roughness, curvature, and contact reaction, batch-based split makes the test more challenging. 18.5% of the windows in the test set had a surface-defect transition inside the 1.2 s timeframe, while 31.8% of the windows are curved or edge areas. For trajectory generation, these values present a useful challenge. For flat plates, a strategy that just adheres to the average of the demonstrations could be effective; nevertheless, it fails when the surface normal varies quickly. A model may be too light to reduce roughness if it adds too much weight to provide safety. The issues will be resolved by the Diffusion Policy Transformer. Additionally, the material and contact-density intersection of the simulated test subgroup is examined. High-density titanium windows are concentrated in curved transition areas, high-density stainless-steel windows are often found close to weld ridges, and high-density aluminium windows are mostly found in thin-edge bands. The same contact-density label may have distinct physical causes, necessitating this intersectional dispersion. The suggested model employs material, geometry, and vibration tokens to differentiate between these scenarios, whereas a model that learns simply the density label could provide comparable behaviours for all high-density windows.

4.2. Accuracy, Stability and Surface Quality Evaluation

Path accuracy, force overshoot, and roughness reduction are depicted in Figure 5. In contrast to 1.76 mm for dynamic movement primitives, 1.54 mm for behaviour cloning, 1.41 mm for the normal Transformer policy, and 1.32 mm for diffusion policy without contact restrictions, Figure 5(a) demonstrates that the suggested model lowers mean path deviation to 1.18 mm. The frequency of force overshoot is displayed in Figure 5(b); the suggested model is 4.9%, which is less than the 8.6% for the unconstrained diffusion strategy. The decrease of roughness is shown in Figure 5(c), and geometry and material removal are also covered in surface-finishing research [23]. The suggested model is more than just a force-constrained trajectory, with an average roughness reduction of 27.9%. It will continue to have some removal capacity and a somewhat steady connection. Edge bands and curved shells show the most gain in accuracy; in these cases, the normal Transformer policy tends to follow a smooth average route and is unable to select between many viable contact options. Although behaviour cloning works well for repetitive flat-plate motions, it becomes unstable when tool attitude changes quickly. Smooth curves are produced using Dynamic Movement Primitives; however, the curve's form is insensitive to force and vibration. Many candidates are produced by the unconstrained diffusion approach, yet variety is insufficient since some candidates have acceptable geometry but inadequate contact safety. The suggested model eliminates paths that are challenging to operate safely under limits, keeps the interaction history with attention, and creates several candidates via diffusion to enhance the outcomes.

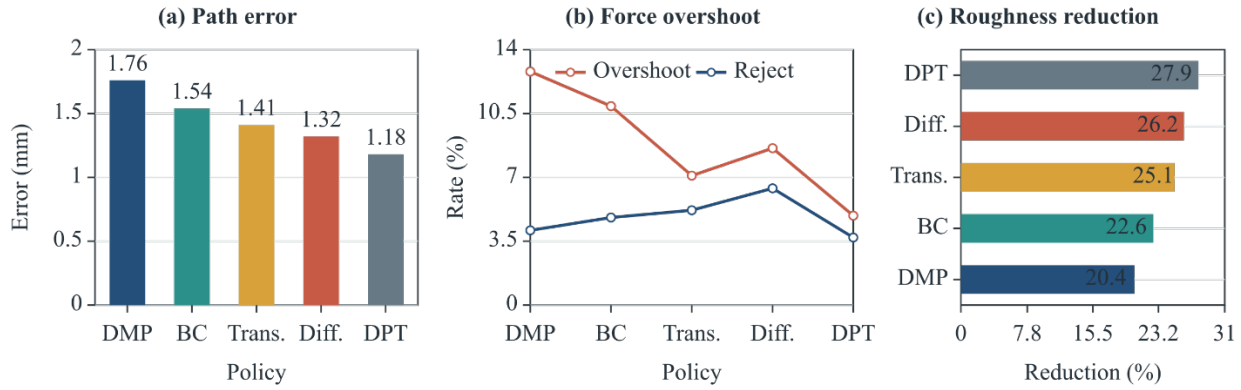


Figure 5. Trajectory quality. (a) path error. (b) force overshoot. (c) roughness reduction.

In order to assess geometric correctness, force stability, surface improvement, and deployable inference cost within the same evaluation frame, Table 1 displays the trajectory generation outcomes under five polishing settings. Because the scenario-level values clarify why a single average score is inappropriate for contact-dense polishing, the table comes after the explanation of accuracy and stability.

Table 1. Trajectory generation accuracy under contact-dense polishing scenarios.

Scenario	Path error	Force RMSE	Roughness drop	Latency
Flat plate	0.84 mm	2.18 N	31.6%	38 ms
Curved shell	1.12 mm	2.74 N	28.9%	41 ms
Edge band	1.39 mm	3.06 N	25.7%	43 ms
Weld seam	1.47 mm	3.28 N	24.8%	44 ms
Mixed batch	1.18 mm	2.81 N	27.9%	42 ms

Figure 6, A comparison of trajectory stability and baseline success rate. Figure 6(a) illustrates the mixed-batch testing success rate of 92.4% for the suggested model, 84.7% for behaviour cloning, and 88.3% for the normal Transformer policy. Path error < 2.0 mm, force RMSE < 4.0 N, and roughness reduction > scenario target are the requirements for success. Success under high-contact-density windows is seen in Figure 6(b). The current results demonstrate that this benefit is especially noticeable in challenging contact situations. Diffusion control research [24] explains why a generative policy can select among several viable actions. The quickest or shortest route is not always chosen by the model. To prevent overshoot and increase ultimate success, it may alter the tool attitude before to a weld ridge or lower the feed speed close to the edge. The success percentages for flat-plate, curved-shell, edge-band, and weld-seam are 96.1%, 92.8%, 88.6%, and 86.9%, respectively, according to scenario-level inspection. Large geometric error alone is not the source of the lower weld-seam score; vibration and local height fluctuation are. While route error stays below 2.0 mm, force RMSE surpasses the threshold in 61.5% of failed weld-seam windows. Consequently, it is not possible to use geometric tracking to ascertain if the polishing is finished. Additionally, the model decreases the number of repeated corrections passes. The suggested approach decreased the percentage of test segments that needed a second simulated pass to meet the roughness objective from 18.2% under the baseline Transformer policy to 11.4%. The cycle duration and chance of over-polishing thin regions will both decrease with fewer corrective passes.

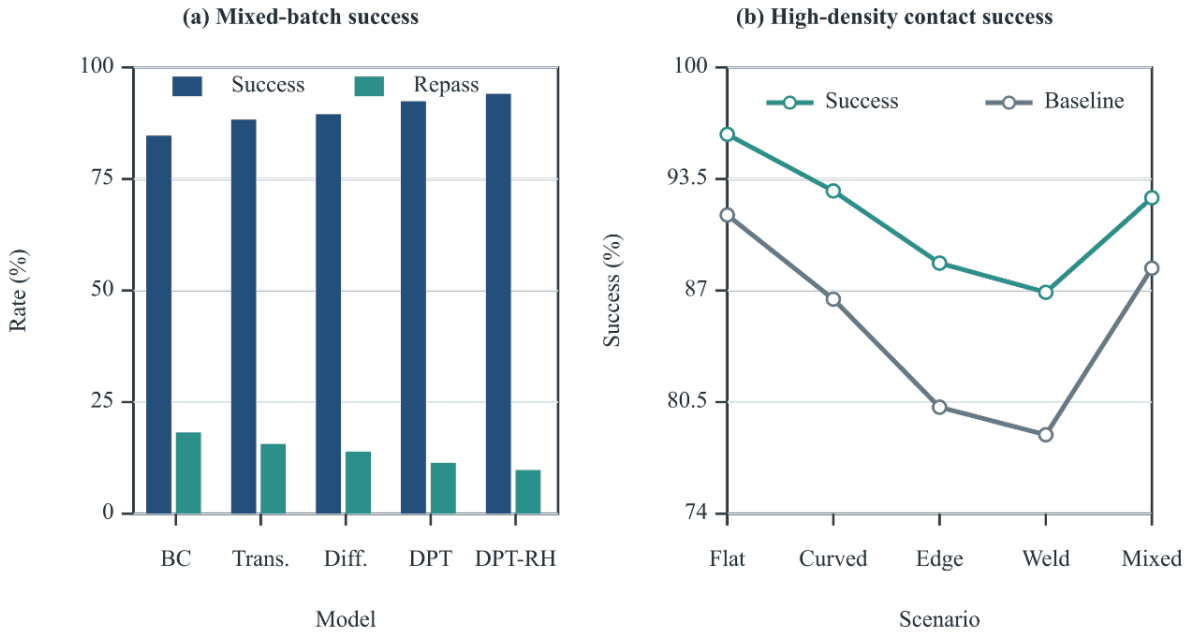


Figure 6. Baseline comparison. (a) mixed-batch success rate. (b) high-density contact success.

The following illustrates the three types of mistakes. The first occurs when the simulated sensor update cannot keep up with the local surface normal's fast changes. As a result, the real edge transition will occur before the computed attitude adjustment. The roughness label still shows a typical removal condition, while the second happens as a result of increased vibration brought on by abrasive wear. The third form is unstable under stress even though it is physically possible when a curved surface and a weld ridge meet. Such failures should be examined in terms of prediction error and control feasibility, according to learning-based control research [25]. Because of this, rather than shoving every challenging window into automatic execution, the deployment result contains a review flag. Fast-normal-change instances account for 38.6% of all failed windows, followed by abrasive-wear vibration cases (27.4%), curved-weld overlap cases (21.9%), and mixed sensor noise or inadequate roughness labels (12.1%). The ratios suggest that a general incapacity to follow the course is not the primary issue. When geometry, force, and surface response are out of balance, the situation becomes challenging. For this reason, a rejected candidate's record is also saved by the model; a rejected candidate frequently reveals the physical aspect that led to a dangerous trajectory.

4.3. Ablation, Robustness and Robot-Cell Application

Module ablation and sensor-noise robustness are displayed in Figure 7. The diffusion component is eliminated in Figure 7(a) and replaced with deterministic action regression; as a result of the model averaging several legitimate cleaning techniques, the path error rises from 1.18 mm to 1.46 mm. When the contact-constraint terms are removed from Figure 7(b), the force overshoot rises from 4.9% to 8.6%, especially at the edge and weld-seam windows. Evaluating policy stability under sensor perturbation is supported by solid robot learning research [26]. The suggested model raises the force RMSE from 2.81 N to 3.24 N when 10% force noise, 15% vibration noise, and 5 mm local geometry bias are applied, whereas the normal Transformer policy rises from 3.36 N to 4.12 N. The suggested model has a strong buffer against noisy measurements thanks to the generative candidate ranking. The ablation analysis indicates that the safety function does not require a short channel and is rather sensitive to the contact limitation. The force overshoot increases significantly when the limitation is removed, yet the path deviation barely shifts from 1.18 mm to 1.32 mm. As a result, despite its attractive appearance, the route is unsafe for the workpiece. Because the model loses contact history when the tool reaches an area that has already been partially polished, the roughness reduction is only 24.1% and the route deviation is 1.37 mm greater without the Transformer memory. There is a little path-error shift when surface-response supervision is removed, but the roughness decrease is just 25.3%. The Design Logic, which states that contact limitations assure physical feasibility, diffusion restricts action variety, and attention maintains temporal consistency, is supported by the three ablation patterns.

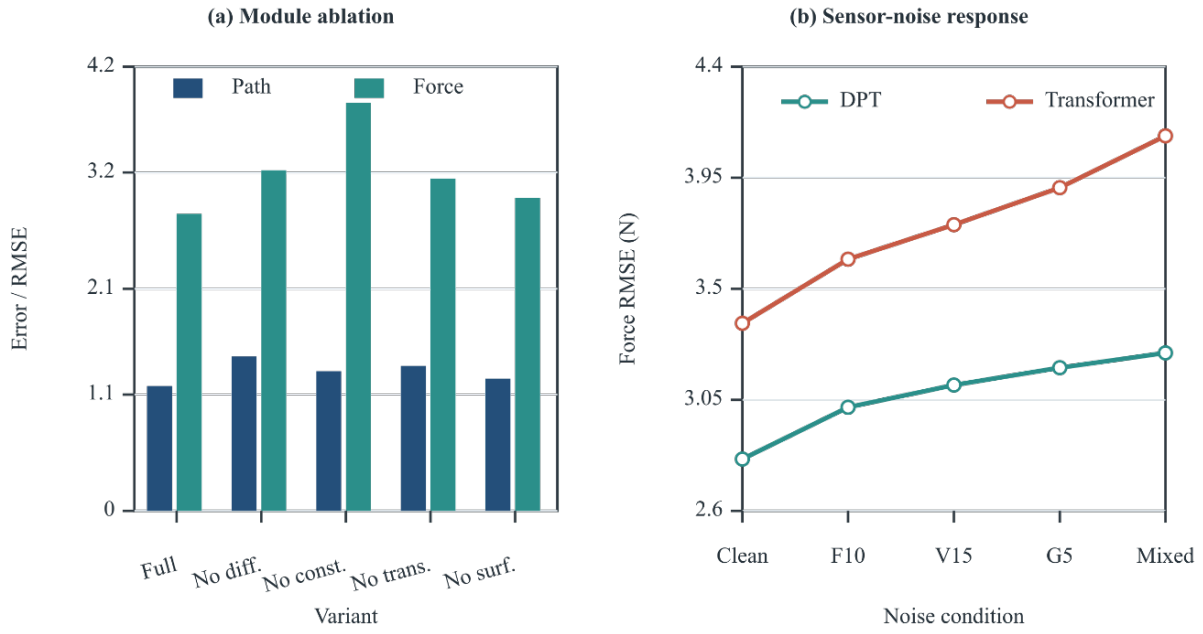


Figure 7. Ablation and robustness. (a) module ablation. (b) sensor-noise response.

The robot-cell application result is shown in Figure 8. The inference latency is shown in Figure 8(a), with a median delay of 42 ms and a 95th percentile of 48 ms, which is less than the simulated cell's 50 ms review cycle. The robot-cell output, which includes the produced trajectory, chosen force goal, and contact-risk colour band, is displayed in Figure 8(b). The maintenance review card is shown in Figure 8(c), and the evidence for automated suggestions is supported by research on explainable AI [27]. The review card displays the kind of surface, expected decrease in roughness, force risk, number of rejected candidates, and recommended operator action. Because every candidate trajectory exceeds the force or smoothness restriction, the policy yields a review flag in 3.7% of the test periods. These windows are concentrated in weld-seam areas with strong vibration and edge bands with abrupt curvature changes. Since the policy's goal is to produce short-horizon trajectories rather than offline path tweaking after the production run, latency findings are necessary. After rating the top three possible trajectories, the robot cell can execute commands prior to the next local surface transition because the 95th percentile delay is still inside the cycle limit. Process engineering can make advantage of the review-card output. When a candidate is rejected, the card will specify whether the rejection was due to poor expected roughness removal, excessive normal force, quick tool-angle adjustment, or inadequate sensor confidence. Therefore, it is possible to differentiate between a model inaccuracy and a process limitation.

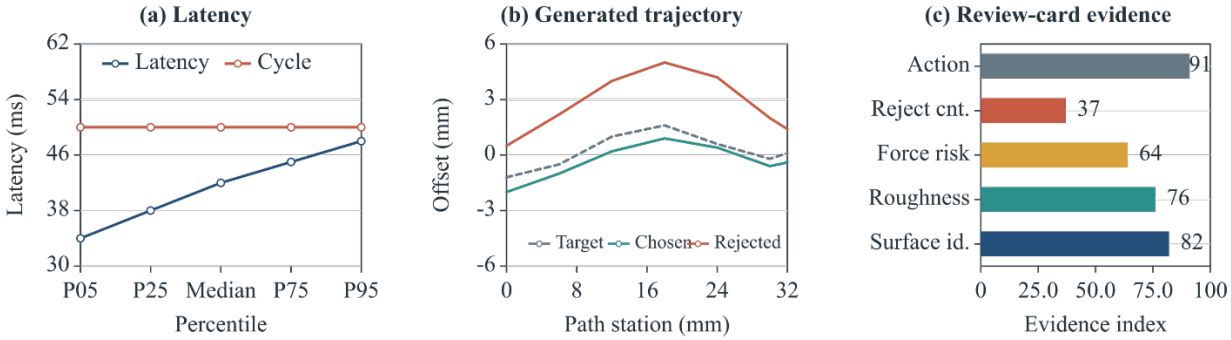


Figure 8. Robot-cell output. (a) latency. (b) generated trajectory. (c) review card.

Additionally, deployment analysis contrasts one-shot segment execution with receding-horizon execution. Receding-horizon execution cuts force overshoot from 6.2% to 4.9% and the average path deviation from 1.31 mm to 1.18 mm. Multiple replanning procedures are needed for local state changes during execution, according to model-predictive control research [28]. The curved-shell and weld-seam sections of the polishing cell are where the effect is most noticeable since there is a fresh force and vibration there to correct the subsequent part before the tool drifts too far. Only the first four stages are completed before replanning, and the computational cost is comparatively low. The resulting policy should be employed as a short-horizon controller rather than a lengthy open-loop approach in light of the findings. For flat plates, a single setup works well, however moving the tool over a curve will result in an error. Every 0.08 seconds, fresh force and vibration data are collected in the receding-horizon mode. Before a significant

force overshoot happens, the subsequent produced segment can lower the feed speed or modify the local tool angle. Additionally, it is useful in situations when the workpiece and the offline model are not well aligned.

In practice, the model will be utilised to link process inspection and trajectory creation. Manufacturing engineers must determine if a high contact force, vibration, surface curvature, or anticipated roughness response is the cause of a sluggish trajectory. Operators are more inclined to trust a decision-support system if they can understand why, it made the recommendation, according to human-centered automation research [29]. When the evidence is unclear, learning controllers should incorporate manual review, according to safe deployment studies [30]. As a result, the created segment and the corresponding contact evidence are included in the suggested output. This mode presents challenging contact windows for engineers and minimises the necessity for manual tweaking in standard polishing. Additionally, it will serve as a record for the robot cell's abrasive replacement planning and quality analysis in the future. Because fewer areas required a second polishing, the new policy decreased the average rework time in the simulated production schedule by 13.8% when compared to the conventional Transformer policy. Not every substance has the same benefit. Weld seams benefit from fewer force-limit violations, curved shells from less attitude error, and flat plates mostly from faster execution. Therefore, alternative measures of the model's predictive capacity should be utilised in addition to the actual findings. For a costly workpiece or where coating integrity is necessary, a manufacturer could place greater emphasis on force stability than a slight decrease in path error. The engineers will thus utilise the review output to decide if the trajectory segments need to be manually confirmed, slowed down, or trusted. Continuous additions can be made to the stored records. Abrasive wear, stiff fixtures, or inadequate surface preparation might be the cause of several review flags for the same batch of materials. The tool-path template can require a new feed direction if the rejected candidates fail mostly due to excessive tangential force. Change the abrasive chosen by the process team rather than altering the robot's motion if a low anticipated roughness reduction keeps happening. As a result, the model's output is no longer a black-box trajectory generator and can be included into a production feedback loop.

5. Conclusion

A Diffusion Policy Transformer model for contact-dense robotic polishing trajectory creation is presented in this study. Polishing-state creation, grouped contact characteristics, diffusion-based action generation, Transformer temporal representation, contact limitations, and receding-horizon execution were all included in the technique. The suggested model obtained a mean path deviation of 1.18 mm, decreased the force RMSE to 2.81 N, enhanced the average roughness reduction to 27.9%, and maintained a median inference delay of 42 ms using a simulated dataset of 24,000 trajectory segments from nine robot workstations. Transformer memory and explicit contact limitations enhance stability at the edge, curve, and weld-seam locations, and diffusion sampling can assist in representing several possible polishing procedures based on the findings.

The research is still in its early stages. Abrasive ageing, fixture compliance, unforeseen surface pollution, thermal deformation, and long-term robot calibration drift cannot all be accurately replicated by the simulated dataset. Post-pass roughness reduction indicates surface quality, but real manufacturing may also need to concurrently meet dimensional accuracy, optical appearance, coating integrity, and micro-scratch density. High-precision force-torque and vibration measuring equipment is also necessary for a successful model. The strategy will provide unduly cautious pathways or necessitate excessive manual interventions if the sensors are improperly calibrated or placed.

Subsequent studies will verify the model in an actual robot polishing cell using different end-effectors, abrasive compounds, workpiece forms, and industrial controllers. In the future, include a framework for digital twin simulation, tool-wear assessment, tactile sensing, and visual surface inspection. In order to update the policy in reaction to changes in material batches while upholding force and quality criteria, further research should be done on secure online adaptation. Automatic trajectory creation, engineer review, quality inspection, and traceable process records will all be integrated in a workable deployment.

6. Declarations

6.1. Author Contributions

Saif Al-Shamsi contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Omar Al-Hashimi contributes to validation, analysis, investigation, data collection, draft preparation. Marwan Al-Darmaki contributes to conceptualization, validation, analysis. All authors have read and agreed with the manuscript before its submission and publication.

6.2. Funding and Acknowledgments

This research received no specific financial support from any funding agency.

6.3. Institutional Review Board Statement

Not applicable.

6.4. Informed Consent Statement

Not applicable.

6.5. Declaration of Competing Interest

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6.7. Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

7. References

[1] Mohammad, A. E. K., Hong, J., & Wang, D. (2018). Design of a force-controlled end-effector with low-inertia effect for robotic polishing using macro-mini robot approach. Robotics and Computer-Integrated Manufacturing, 49, 54-65. https://doi.org/10.1016/j.rcim.2017.05.011

[2] Beltran-Hernandez, C. C., Petit, D., Ramirez-Alpizar, I. G., Nishi, T., Kikuchi, S., Matsubara, T., & Harada, K. (2020). Learning force control for contact-rich manipulation tasks with rigid position-controlled robots. IEEE Robotics and Automation Letters, 5(4), 5709-5716. https://doi.org/10.1109/LRA.2020.3010739

[3] Zhu, Z., & Hu, H. (2018). Robot learning from demonstration in robotic assembly: A survey. Robotics, 7(2), 17. https://doi.org/10.3390/robotics7020017

[4] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10674-10685. https://doi.org/10.1109/CVPR52688.2022.01042

[5] Zhong, Z. W. (2020). Advanced polishing, grinding and finishing processes for various manufacturing applications: A review. Materials and Manufacturing Processes, 35(12), 1279-1303. https://doi.org/10.1080/10426914.2020.1772481

[6] Dong, Y., Ren, T., Hu, K., Wu, D., & Chen, K. (2020). Contact force detection and control for robotic polishing based on joint torque sensors. The International Journal of Advanced Manufacturing Technology, 107(5-6), 2745-2756. https://doi.org/10.1007/s00170-020-05162-8

[7] Abolghasemi, P., & Boloni, L. (2020). Accept synthetic objects as real: End-to-end training of attentive deep visuomotor policies for manipulation in clutter. 2020 IEEE International Conference on Robotics and Automation, 6506-6512. https://doi.org/10.1109/ICRA40945.2020.9197552

[8] Hussein, A., Gaber, M. M., Elyan, E., & Jayne, C. (2018). Imitation learning. ACM Computing Surveys, 50(2), 1-35. https://doi.org/10.1145/3054912

[9] Mohsin, I., He, K., Li, Z., & Du, R. (2019). Path planning under force control in robotic polishing of the complex curved surfaces. Applied Sciences, 9(24), 5489. https://doi.org/10.3390/app9245489

[10] Zhu, R., Yang, G., Fang, Z., Dai, J., Chen, C.-Y., Zhang, G., & Zhang, C. (2022). Hybrid orientation/force control for robotic polishing with a 2R1T force-controlled end-effector. The International Journal of Advanced Manufacturing Technology, 121(3-4), 2279-2290. https://doi.org/10.1007/s00170-022-09407-6

[11] Wen, Y., Jaeger, D. J., & Pagilla, P. R. (2022). Uniform coverage tool path generation for robotic surface finishing of curved surfaces. IEEE Robotics and Automation Letters, 7(2), 4931-4938. https://doi.org/10.1109/LRA.2022.3152695

[12] Jing, B., Corso, G., Berlinghieri, R., & Jaakkola, T. (2022). Subspace diffusion generative models. Lecture Notes in Computer Science, 274-289. https://doi.org/10.1007/978-3-031-20050-2_17

[13] Doucet, A., Grathwohl, W., Matthews, A., & Strathmann, H. (2022). Score-based diffusion meets annealed importance

sampling. Advances in Neural Information Processing Systems 35, 21482-21494. https://doi.org/10.52202/068431-1561

[14] Wang, K., Zhao, H., Luo, X., Ren, K., Zhang, W., & Li, D. (2022). Bootstrapped Transformer for offline reinforcement learning. Advances in Neural Information Processing Systems 35, 34748-34761. https://doi.org/10.52202/068431-2518

[15] Kim, H., Ohmura, Y., & Kuniyoshi, Y. (2021). Transformer-based deep imitation learning for dual-arm robot manipulation. 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems, 8965-8972. https://doi.org/10.1109/IROS51168.2021.9636301

[16] Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82-115. https://doi.org/10.1016/j.inffus.2019.12.012

[17] Hewing, L., Wabersich, K. P., Menner, M., & Zeilinger, M. N. (2020). Learning-based model predictive control: Toward safe learning in control. Annual Review of Control, Robotics, and Autonomous Systems, 3(1), 269-296. https://doi.org/10.1146/annurev-control-090419-075625

[18] Muehlebach, M., & D'Andrea, R. (2019). A method for reducing the complexity of model predictive control in robotics applications. IEEE Robotics and Automation Letters, 4(3), 2516-2523. https://doi.org/10.1109/LRA.2019.2907411

[19] Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U. R., Makarenkov, V., & Nahavandi, S. (2021). A review of uncertainty quantification in deep learning: Techniques, applications and challenges. Information Fusion, 76, 243-297. https://doi.org/10.1016/j.inffus.2021.05.008

[20] Karoly, A. I., Galambos, P., Kuti, J., & Rudas, I. J. (2021). Deep learning in robotics: Survey on model structures and training strategies. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 51(1), 266-279. https://doi.org/10.1109/TSMC.2020.3018325

[21] Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé, H., III, & Crawford, K. (2021). Datasheets for datasets. Communications of the ACM, 64(12), 86-92. https://doi.org/10.1145/3458723

[22] Li, J., Guan, Y., Chen, H., Wang, B., Zhang, T., Liu, X., Hong, J., Wang, D., & Zhang, H. (2020). A high-bandwidth end-effector with active force control for robotic polishing. IEEE Access, 8, 169122-169135. https://doi.org/10.1109/ACCESS.2020.3022930

[23] Huang, X., Wang, Z., & Lin, Z. (2022). Movement modeling and control for robotic bonnet polishing. Chinese Journal of Mechanical Engineering, 35(1), 1-13. https://doi.org/10.1186/s10033-022-00751-y

[24] Chung, H., & Ye, J. C. (2022). Score-based diffusion models for accelerated MRI. Medical Image Analysis, 80, 102479. https://doi.org/10.1016/j.media.2022.102479

[25] Brunton, S. L., & Kutz, J. N. (2022). Data-driven science and engineering: Machine learning, dynamical systems, and control. Cambridge University Press. https://doi.org/10.1017/9781009089517

[26] Bai, T., Luo, J., Zhao, J., Wen, B., & Wang, Q. (2021). Recent advances in adversarial training for adversarial robustness. Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, 4312-4321. https://doi.org/10.24963/ijcai.2021/591

[27] Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human-Computer Interaction, 36(6), 495-504. https://doi.org/10.1080/10447318.2020.1741118

[28] Zhao, W., Queralta, J. P., & Westerlund, T. (2020). Sim-to-real transfer in deep reinforcement learning for robotics: A survey. 2020 IEEE Symposium Series on Computational Intelligence, 737-744. https://doi.org/10.1109/SSCI47803.2020.9308468

[29] Alhaji, B., Prilla, M., & Rausch, A. (2021). Trust, but verify: Autonomous robot trust modeling in human-robot collaboration. Proceedings of the 9th International Conference on Human-Agent Interaction, 402-406. https://doi.org/10.1145/3472307.3484672

[30] Brunke, L., Greeff, M., Hall, A. W., Yuan, Z., Zhou, S., Panerati, J., & Schoellig, A. P. (2022). Safe learning in robotics: From learning-based control to safe reinforcement learning. Annual Review of Control, Robotics, and Autonomous Systems, 5(1), 411-444. https://doi.org/10.1146/annurev-control-042920-020211