

RoboGrind: Intuitive and Interactive Surface Treatment with Industrial Robots

Benjamin Alt^{1,2,*}, Florian Stöckl³, Silvan Müller³, Christopher Braun^{4,5}, Julian Raible^{4,5},
Saad Alhasan³, Oliver Rettig³, Lukas Ringle¹, Darko Katic¹, Rainer Jäkel¹,
Michael Beetz², Marcus Strand³, and Marco F. Huber^{4,5}

Abstract—Surface treatment tasks such as grinding, sanding or polishing are a vital step of the value chain in many industries, but are notoriously challenging to automate. We present RoboGrind, an integrated system for the intuitive, interactive automation of surface treatment tasks with industrial robots. It combines a sophisticated 3D perception pipeline for surface scanning and automatic defect identification, an interactive voice-controlled wizard system for the AI-assisted bootstrapping and parameterization of robot programs, and an automatic planning and execution pipeline for force-controlled robotic surface treatment. RoboGrind is evaluated both under laboratory and real-world conditions in the context of refabricating fiberglass wind turbine blades.

I. INTRODUCTION

In a wide range of industries such as aerospace, consumer goods manufacturing, or the energy sector, surface treatment tasks are integral components of the value chain. With the rise of remanufacturing as a core component of the circular economy, the need for robust, cost-efficient (re-)finishing of surfaces has become even more pressing. One example is the refabrication of wind turbine rotor blades, which require considerable surface treatment after several years of use. Due to a shortage of qualified labor and high costs, economically feasible remanufacturing requires novel, robust solutions for automating surface finishing tasks with robots.

Robotic surface treatment is challenging due to the physics of contact and abrasion, which are hard to simulate and require robust force control. Moreover, in the context of remanufacturing, workpieces have been subject to different degrees of wear, requiring sophisticated perception and algorithms for automatically identifying defects. At the same time, many steps in the surface treatment workflow, from the proper parameterization of robot programs to the choice of tool, require considerable human expertise. We hypothesize that a software system that combines sophisticated perception, planning, reasoning, and control capabilities can greatly reduce the cost of automation for surface finishing tasks.




Fig. 1: RoboGrind is an intuitive, interactive system for robotic surface treatment comprising perception, program generation, planning and control.

We also propose that the human is an integral source of knowledge, guidance, and oversight, and ought to remain involved in the surface finishing process, though in a different capacity and to a lesser degree.

In this paper, we present RoboGrind, a software system for interactive, artificial intelligence (AI)-assisted surface treatment with industrial robots. It combines four distinct technical contributions:

- 1) Advanced 3D vision and automatic defect detection, enabling the automatic treatment of different surface geometries and defect locations;
- 2) Automatic planning of tool paths based on a 3D scan of the workpiece, depending on the tool and material;
- 3) AI-assisted bootstrapping of robot programs, with a deep-learning-based natural language processing (NLP) frontend for user interaction;
- 4) Automatic simulation and force-controlled execution of the generated robot programs.

RoboGrind integrates these components into a unified architecture to achieve a very high degree of automation. The perception, planning, and control systems are individually evaluated under laboratory conditions. The overall system is evaluated under real-world conditions for the refabrication of fiberglass wind turbine blades. This paper conducts one of the first investigations into force-controlled disk sanding of fiberglass with lightweight collaborative robot arms.

¹ArtiMinds Robotics, 76131 Karlsruhe, Germany

²Institute for Artificial Intelligence, University of Bremen, 28359 Bremen, Germany

³Baden-Württemberg Cooperative State University, 76133 Karlsruhe, Germany

⁴Institute of Industrial Manufacturing and Management IFF, University of Stuttgart, 70569 Stuttgart, Germany

⁵Fraunhofer Institute for Manufacturing Engineering and Automation IPA, 70569 Stuttgart, Germany

* Corresponding author: benjamin.alt@uni-bremen.de

II. RELATED WORK

A. Robotic Surface Treatment

Li et al. implemented a pipeline for automated rail grinding [1], which removes excess material from welds. They designed a system with seven main modules, including modules for measurement, motion control, and information feedback. The main difference to our system is their specialization on weld seams and lack of guided adaption of sanding parameters based on a knowledge base. Oyekan et al. researched methods to automate fan-blade reconditioning of aerospace maintenance [2], using a simulation environment based on a digital twin. Our work ties back to the point they make in their conclusion: The main bottleneck for automated grinding remains “embedding the knowledge of a skill worker into an automated cell” [2].

B. AI-Assisted Robot Programming

A wide variety of methods facilitating the creation of robot programs have been proposed. Task and Motion Planning (TAMP) approaches view robot programming as a joint task- and motion-level planning problem and combine search-based planning in task space with collision-free motion planning [3], [4]. CRAM [5] proposes a knowledge-based approach to robot programming by combining a symbolic program and knowledge representation with the KnowRob knowledge representation and reasoning (KR&R) engine [6] and a realistic simulation environment. This paradigm has been shown to support several forms of AI-assisted robot programming [7], or synthesis of robot programs from human virtual reality (VR) demonstrations [8]. Several recent approaches propose to leverage natural language as a suitable abstraction for program synthesis. Code as Policies [9] and ProgPrompt [10] use a large language model (LLM) to generate Python code that uses an API of robot primitives to solve manipulation tasks. TidyBot [11] uses natural language as an intuitive representation for user input. We propose to combine the advantages of natural language-based systems with structured KR&R and planning to achieve intuitive user interaction and precise, industrially robust execution.

III. OVERVIEW

RoboGrind is an interactive system for rapidly and intuitively automating surface finishing tasks with robots. It achieves this by combining advanced perception and planning algorithms with KR&R capabilities and natural language understanding. RoboGrind realizes a three-step workflow (see Fig. 2):

1) *3D Surface Scanning*: The workpiece is scanned with a laser scanner. Scanning and scan data processing is nearly completely automated and defects on the workpiece are automatically detected.

2) *AI-Assisted Robot Programming*: Via natural-language interaction, the user is guided through the process of generating a parameterized robot program for the given task (e.g., sanding, polishing, or deburring). A digital twin of the robot and the required robot skills to perform the task are automatically instantiated and parameterized. A path planner

generates a Cartesian tool path depending on the surface geometry.

3) *Force-Controlled Surface Treatment*: The generated robot program is simulated and displayed to the user for validation. It is then automatically compiled into executable robot code and executed on the robot. This typically involves force-controlled, abrasive motions in contact with the surface, as well as collision-free approach and depart motions. As most surface treatment tasks are iterative processes, the process is iterated until the desired surface quality is reached.

The three steps of the workflow are detailed in the following sections.

IV. 3D SURFACE SCANNING

RoboGrind performs program generation and path planning based on a 3D point cloud representation of the surface. We use a Gocator 2490¹ laser line scanner with a field of view of up to 2 m and a resolution of 6 μm . The starting point of a scan is defined manually according to the measurement range. Segmentation, outlier detection, and downsampling are performed. Based on an initial calibration step, the point cloud is transformed into the robot’s base frame. We propose an automatic defect detection pipeline that combines two outlier detection algorithms: First, curve fitting by linear regression is performed separately for the points of each scanned line and, based on several thresholds, outliers and defects are marked and excluded. Second, a StatisticalOutlierRemoval filter from the Point Cloud Library (PCL) based on k-Nearest Neighbors (k-NN) is used to detect points with few neighbors. This allows detecting scanning artefacts that cause holes in point clouds, and eliminates outliers that do not belong to a scanned part. It also helps reducing the number of false positives detected, as these have no high point density and have too few neighbors. The combined algorithms can detect all defects that deviate from an ideal surface shape, including dents, bumps, scratches, and rough surface areas. Unlike neural networks, they require no training. The code can be found in our GitLab repository².

V. AI-ASSISTED ROBOT PROGRAMMING

A core objective of RoboGrind is to solve robot-based surface treatment tasks with a high degree of automation. RoboGrind builds upon a KR&R engine to automate large parts of the surface treatment workflow. At the same time, the complexity of the task as well as safety considerations necessitate the continued involvement of human experts in the process. Recognizing this, RoboGrind realizes a robot programming paradigm that is fundamentally interactive, and integrates a natural language interface to facilitate intuitive interaction with the user.

A. Meta-Wizard Architecture

The cognitive subsystem for interactive robot programming, which we call *meta-wizard*, comprises a knowledge base and a novel reasoner for surface treatment tasks. Both

¹LMI Technologies GmbH, Teltow, Germany

²<https://gitlab.com/rahm-lab/robogrind-surfacetreatment>

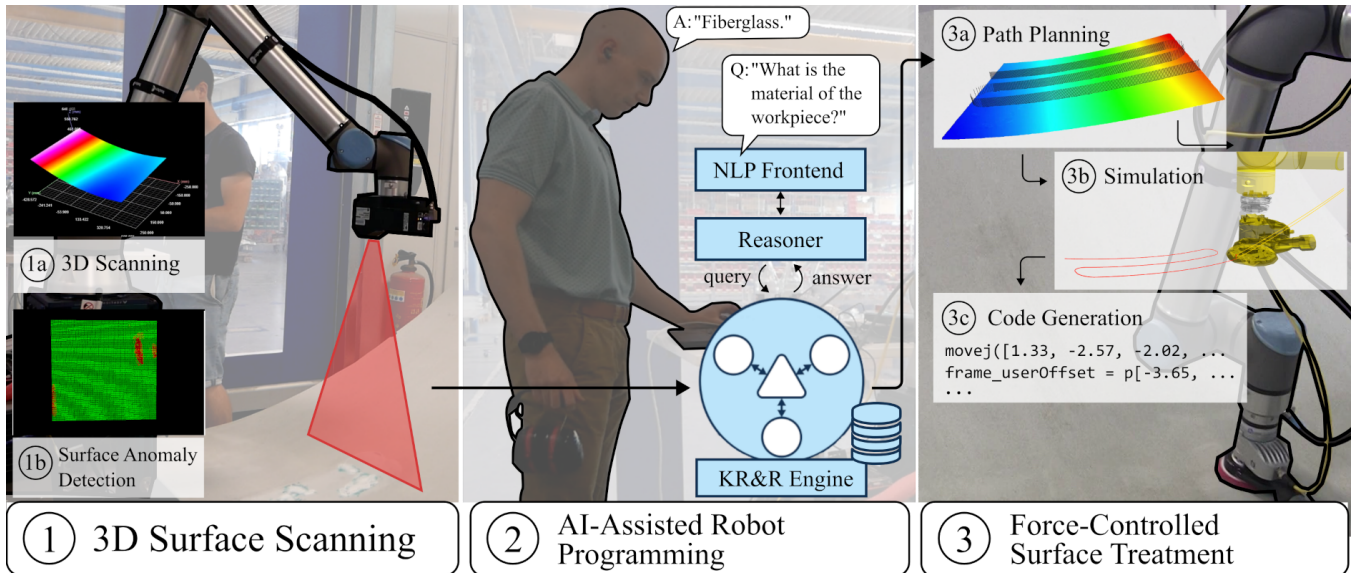


Fig. 2: RoboGrind integrates perception (1), AI-assisted programming (2), planning and force control (3) into a comprehensive assistance system for robotic surface treatment.

are integrated into the KnowRob KR&R engine [6], which provides a framework for robot cognition and metaprogramming.

1) *Knowledge Base*: The knowledge base contains commonsense, robot-specific and domain-specific knowledge and forms the basis for reasoning. In the KnowRob framework, knowledge is stored as Resource Description Format (RDF) triples in a MongoDB database. KnowRob provides utilities for the automatic conversion of Web Ontology Language (OWL) ontologies into RDF triples, which in turn allows reasoners to jointly reason over symbolic (ontological) knowledge and subsymbolic, unstructured data such as force-torque data from a robot [6]. Commonsense and robot-specific knowledge is provided via the SOMA upper-level ontology [12], while additional domain-specific knowledge is represented in a novel RoboGrind ontology. This ontology extends SOMA by surface treatment-specific concepts such as task types (e.g., Sanding), tools (e.g., OrbitalSander), or materials (e.g., Fiberglass) as well as their subsymbolic properties (e.g., rotational speeds). The knowledge base can be extended to new domains, materials, or tools by adding new classes and individuals to the ontology.

One important feature of the meta-wizard subsystem is that the surface treatment workflow itself is represented as classes and individuals of the ontology. Using `soma:Workflow` and related classes, surface treatment processes can be represented by modeling their respective steps (e.g., `SurfaceScanning`, `Simulation`, `Execution`, `QualityControl`) and chaining them via succedence relations. This representation of workflows as ontological knowledge permits reasoners to reason about the workflow itself, and enables the meta-wizard to change its behavior depending on its belief state—its ontological knowledge and current understanding of the task and

environment. Moreover, it permits human experts to add procedural knowledge to the wizard without changing its code.

2) *Domain-Specific Reasoner*: To realize the cognitive functions required for interactive program synthesis, we contribute a novel reasoning system consisting of a Python executable (the control module) and a KnowRob plugin (the reasoner). The control module controls the program generation process, using the Prolog-based reasoners in KnowRob and the natural language interface (see Sec. V-B) to gather information about the workflow, tools, workpieces, etc., and to make decisions about what actions to take. The reasoner is implemented as a set of Prolog predicates over classes, individuals, and relations in the knowledge base. Using these predicates, the control module can not only access explicit knowledge, but also derive new facts about existing knowledge. One example is the predicate `has_tool(Task, Tool)`, which is true for all tool types that can be used for a given task.

B. Natural Language Interface

At runtime, some information required for program synthesis, such as the exact material of the workpiece, cannot be inferred from the knowledge base and must be obtained by interacting with the user. For this purpose, we leverage an NLP pipeline combining the Google Speech Recognition API [13], based on deep neural networks, and the spaCy [14] NLP library:

- 1) Conversion of the user’s speech to text using the Google Speech Recognition API.
- 2) Grammar-based chunking of the text into semantic units using spaCy.
- 3) Semantic matching of chunks to the concept(s) to be grounded (e.g., “mm” to `Unit`, or “fiberglass” to `Material`), based on spaCy’s word vector distance.

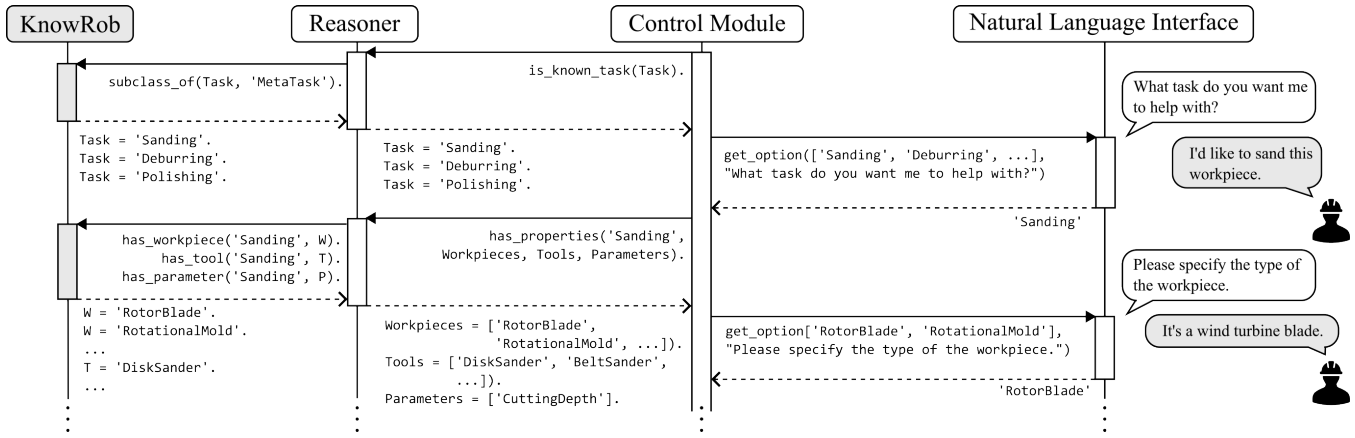


Fig. 3: Illustration of the meta-wizard’s interactive symbol grounding mechanism. The meta-wizard’s main control module retrieves symbolic knowledge about tasks, workpieces, etc., via Prolog queries to a domain-specific reasoner connected to the KnowRob KR&R engine. Missing information about the concrete task, workpiece, etc., is obtained via dialog with the user.




Fig. 4: Path planning on point cloud data. Fig. 5: Test setup under real-world conditions.

Voice output (text-to-speech) is realized with the Windows voice APIs.

C. Interactive Programming Workflow

Fig. 3 shows part of the interactive robot programming workflow. The complete process takes the following steps:

- 1) **Meta-task grounding:** The user is prompted for the top-level task (e.g., Sanding). Suitable tools, possible workpieces and task parameters (such as the amount of material to be removed) are retrieved from the knowledge base. Missing knowledge is grounded by asking the user.
- 2) **Task selection:** To take steps towards achieving the meta-task, the next (sub-)task is identified by querying the knowledge base for the next successor in the meta-tasks workflow definition, depending on the meta-wizard’s current belief state. In some cases, this requires input from the user, e.g., when deciding whether to perform an additional pass over the surface.
- 3) **Task execution:** The selected task (e.g., Simulation) is executed by calling a corresponding handler function in the control module.
- 4) **Iteration:** Iteration of steps 2) and 3) until a terminal step in the workflow is reached.

When the information provided by the user is insufficient, e.g. due to misunderstood voice input, this information is discarded and the current step is repeated.

VI. FORCE-CONTROLLED SURFACE TREATMENT

A. Path Planning

According to the position and orientation of the scanned surface, a set of intersection planes is adaptively defined. By slicing the point cloud uniformly with multiple parallel intersection planes, a set of cross-section contours is determined. This requires identifying all points whose distance from a plane does not exceed a certain threshold. These points are then projected onto the corresponding plane and subsequently sorted, filtered, and interpolated to form the corresponding contours. As a continuous path is desired, a curve determination process enables generating sets of points that allow connecting the adjacent contours to form one meandering path (see Fig. 4). The path planner is designed to operate on vertically projectively planar surfaces. We refer to [15] for more technical details and discussion of the planner.

B. Simulation and Code Generation

The generated robot program and tool path are simulated and visualized in a 3D environment. The 3D digital twin of the robot is updated by the meta-wizard during program generation to reflect, e.g., the surface finishing tool. Simulation both permits to validate the kinematic feasibility of the planned path and the collision-freeness of approach and depart motions, but also permits the user to ensure safe execution of the program. The program is then compiled to executable robot code and executed on the robot. For simulation, compilation, and execution, the robot programming software ArtiMinds RPS³ is used.

³ArtiMinds Robotics GmbH, Karlsruhe, Germany

C. Hybrid Force-Position Control

At runtime, the planned path is executed by a hybrid force-position controller. The control law is a PID controller

$$u_{\text{wrench}}(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de(t)}{dt},$$

$$u(t) = u_{\text{pose}}(t) + u_{\text{wrench}}(t).$$

The control signal $u(t)$ is a 6-dimensional vector denoting a spatial offset (Cartesian position and orientation), which is the sum of separate wrench and pose components. The 6D wrench component $u_{\text{wrench}}(t)$ is computed by a standard PID law, where the wrench error $e(t)$ is the distance between the currently measured end-effector wrench to the allowed wrench region (a 6D hypercube spanned by the wrenches allowed for the application). For sanding of fiberglass, this region collapses to a point, where the force F_z is applied along the z -dimension, and zero along all other dimensions. The pose component $u_{\text{pose}}(t)$ is the deviation of the current end-effector pose from the point on the planned path at time t .

VII. EXPERIMENTS

The following comprehensive experimental evaluation of RoboGrind assesses its capabilities regarding robot-assisted surface treatment for remanufacturing, exemplified by wind turbine blades. The perception, planning, and control subsystems are individually assessed under laboratory conditions. Moreover, the overall system is evaluated holistically under real conditions. In all experiments, we utilize a UR10e robot⁴ with a disk sander⁵ for safe human-robot collaboration. As robots by ABB, KUKA, FANUC or MOTOMAN are commonly used for polishing tasks [16], we additionally provide an assessment of the suitability of UR robots with respect to surface treatment.

A. Laboratory Experiments

1) *Perception*: Three identical 500 mm x 750 mm pieces of the same uncoated fiberglass rotor blade with the same curvature are partitioned into four segments, resulting in a total of twelve concave segments with various surface defects (indentations and rough areas). The segments are scanned, followed by filling of the defects with fiberglass putty, scanned again, and finally all the segments are scanned a third time after sanding. The result is 36 point clouds, 12 without defects and 24 with 96 detectable defects. The detected, undetected and falsely detected defects are marked on the 24 point clouds with detectable defects.

2) *Planning*: Precise path planning is characterized by how well the path aligns with the surface it approximates. To measure this alignment, we determine the distance from each point on the path to the closest point within the surface point cloud obtained after being primed for subsequent sanding. Using these distances, we estimate the accuracy of the approximation by computing the metrics Root Mean Square

⁴Universal Robots A/S, Odense, Denmark

⁵OnRobot GmbH, Soest, Germany

TABLE I: Evaluation of the path planning algorithm w.r.t. to its ability of approximating the surface to be sanded. Metrics are computed for each of the 12 point clouds independently. Averages and standard deviations are reported.

RMSE / mm	MAE / mm	MAX / mm
0.372 ($\sigma = 0.030$)	0.348 (0.034)	1.046 (0.330)

Error (RMSE), Mean Absolute Error (MAE), and Maximum Absolute Error (MAX). It is worth mentioning that the scores computed are directly correlated with the resolution of the surface point cloud, i.e., higher resolution leads to lower scores in general and vice versa, requiring normalization for ideal comparability. We report the raw scores without normalization as all point clouds have nearly the same amount of total points ($\mu = 83,501.75$, $\sigma = 2,288.78$), which introduces a small but defensible bias in favor of reporting values with an associated dimension (mm).

3) *Control*: For each of the 36 scans, the surface is sanded at least once using the proposed simulation, planning, and control subsystems. Robot end-effector trajectories, comprising end-effector forces, are measured and analyzed. Each surface segment is sanded with different parameters, varying the contact pressure F_z , angle of attack α (in $^\circ$) and number of passes. The contact pressures used in the z direction were 5 N and 10 N, the angles of attack between the disc and the workpiece were 2° and 5° , and the process was repeated for 5 or 10 passes. A total of 120 executions have been performed. To quantify the precision of the force controller, we compute the MAE and MAX error of the measured end-effector forces along the z -axis with respect to the contact force setpoint F_z . In addition, we compute the rise time of the controller, defined as the time it takes for the measured end-effector force to reach 90 % of F_z .

B. Real-World Experiments

To evaluate RoboGrind under realistic conditions, we deploy it on-site at a company specialized in robot-based surface treatment⁶. In contrast to the laboratory experiments, the experiments are conducted on the complete blade and not on a cut-out section. First, a scan of the surface and two sanding passes are performed (rough and fine sanding). Then fiberglass putty is applied, the surface is scanned again and two additional passes are performed. The main aim of the experiment is to determine to what extent RoboGrind can provide user assistance for robotic surface treatment. The experiment setup is shown in Fig. 5.

VIII. RESULTS

A. Laboratory Experiments

1) *Perception*: Evaluation of laboratory experiments regarding the perception and damage detection demonstrated that perception is robust in over 71 % of the 24 samples and 96 included damages. 69 defects were correctly detected,

⁶SHL AG, Böttingen, Germany

TABLE II: Measured force control metrics for the three tested contact force threshold values F_z . MAX is computed after F_z has first been reached.

Parameters		Metrics		
F_z / N	MAE / N	MAX / N	Rise time	Trials
5	1.47	9.74	0.96	40
10	1.84	7.17	1.75	26
20	2.10	5.52	1.35	17




Fig. 6: Four randomly selected, exemplary executions for each of the three tested contact forces. One execution per group has been highlighted. Smaller force setpoints lead to better controller behavior.

27 were missed and 7 marked spots were false positives. Remaining problems are holes in point clouds caused by failed measurements and false positive defect detections. Failed measurements result from darker colored or dirty surfaces that are not detectable by the laser scanner. False positives occur in sections where the depth of the defects is low because of the dynamic thresholds of the algorithms.

2) *Planning*: The planner performs robustly on all 12 point clouds (see Table I and Fig. 4). Scores ranging from 0.331 mm to 0.425 mm for RMSE and 0.307 mm to 0.420 mm for MAE indicate a consistent performance, although some point clouds do contain void areas, leading to higher MAX scores.

3) *Control*: Out of the 120 sanding attempts, 69 % were successful. The remainder failed due to exceeding force-torque limits of the UR10e, which occurred because of the natural vibration of the workpiece. For the 83 successful executions, the computed metrics are shown in Table II. The MAE increases with the force setpoint, but is limited to 2.1 N, indicating sufficiently precise control for fiberglass sanding. The rise time remained below 1 second for $F_z = 5$. Overall, we found that a low force setpoint of 5 N achieves the best quantitative controller behavior and resulting surface quality. Exemplary force trajectories for each force setpoint are plotted in Fig. 6.

In the laboratory experiments, we found that the sanding parameters shown in Table III deliver the best results. Additionally to these parameters the rotational speed and the traverse speed were set to 6,000 rpm and 20 m/s, respectively. The surface roughness of the rotor blade before the process is about $90 \mu\text{m}$. We compare the roughness after filling R_{a1}

TABLE III: Evaluation of the surface quality in laboratory experiments. Measured surface roughness for four parameter sets.

Parameters			Metrics	
F_z / N	$\alpha / ^\circ$	Passes	$R_{a1} / \mu\text{m}$	$R_{a2} / \mu\text{m}$
10	5	5	199.57	122.44
10	2	10	169.45	104.25
5	5	5	123.46	162.02
5	2	10	225.01	119.82

to roughness after sanding R_{a2} , which is close to the initial value.

B. Real-World Experiments

In the field experiment, we found that all components of RoboGrind performed as well in an industrial setting as in the laboratory. The 3D scanning and path planning algorithms performed robustly under realistic, uncontrolled light conditions, both before and after putty has been applied. For the path planner, we report errors close to the values observed in the laboratory (0.378 mm, 0.357 mm and 0.994 mm for RMSE, MAE and MAX). After putty has been applied, small regions with strong noise regarding the scanned surface led to noisy point normals and therefore a slight, but sudden change in the orientation of the sander. This can be avoided by applying a global low-pass filter on the point cloud before path planning. The interactive meta-wizard understood user inputs well despite loud background noise. The control behavior is influenced by the much larger dimensions of the workpiece (see Fig. 5) and the resulting strongly differing natural frequencies. A sufficient surface quality for the application was achieved. We leave the optimization of controller parameters for future work. Overall, the UR10e robot proved suitable for sanding tasks.

IX. CONCLUSIONS

In this work, we introduced RoboGrind, an AI-assisted robotic system for automating surface treatment tasks.

We contribute the first quantitative evaluation of robotic sanding of fiberglass with a collaborative robot. Moreover, we provide a qualitative evaluation on a real-world use case.

Our findings indicate that RoboGrind is able to largely automate the process from perception to program execution for surface treatment. The system components achieve a level of reliability and competence suitable for real-world use.

However, some limitations remain. The meta-wizard is currently limited to tasks within its existing knowledge base—future work will explore methods for learning new tasks from demonstration or observation. Applications beyond surface treatment, such as peg-in-hole assembly, also remain open directions for future work.

ACKNOWLEDGMENT

This work was supported by the state of Baden-Württemberg in the project RoboGrind under grant BW1.0079/01, the DFG CRC EASE (CRC #1320) and the EU project euROBIN (grant 101070596).

REFERENCES

- [1] J. Li, T. Yuan, W. Wu, H. Zhu, C. Zhang, and J. Xie, "Automatic programming system for grinding robot of chsr rail," *Proceedings of the 2018 IEEE International Conference on Robotics and Biomimetics*, 2018.
- [2] J. Oyekan, M. Farnsworth, W. Hutabarat, D. Miller, and A. Tiwari, "Applying a 6 DoF robotic arm and digital twin to automate fan-blade reconditioning for aerospace maintenance, repair, and overhaul," *Sensors*, vol. 20, no. 16, p. 4637, 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/16/4637>
- [3] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated Task and Motion Planning," *arXiv:2010.01083 [cs]*, Oct. 2020, arXiv: 2010.01083. [Online]. Available: <http://arxiv.org/abs/2010.01083>
- [4] L. P. Kaelbling and T. Lozano-Perez, "Hierarchical task and motion planning in the now," in *ICRA*, May 2011, pp. 1470–1477. [Online]. Available: <http://ieeexplore.ieee.org/document/5980391/>
- [5] M. Beetz, L. Mösenlechner, and M. Tenorth, "CRAM — A Cognitive Robot Abstract Machine for everyday manipulation in human environments," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2010, pp. 1012–1017, iSSN: 2153-0866.
- [6] M. Beetz, D. Bessler, A. Haidu, M. Pomarlan, A. K. Bozcuoglu, and G. Bartels, "Know Rob 2.0 — A 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents," in *ICRA*, May 2018, pp. 512–519. [Online]. Available: <https://ieeexplore.ieee.org/document/8460964/>
- [7] S. Koralewski, G. Kazhoyan, and M. Beetz, "Self-Specialization of General Robot Plans Based on Experience," *IEEE Robotics and Automation Letters*, 2019.
- [8] B. Alt, F. K. Kenfack, A. Haidu, D. Katic, R. Jäkel, and M. Beetz, "Knowledge-Driven Robot Program Synthesis from Human VR Demonstrations," Rhodes, Greece, Sep. 2023, arXiv:2306.02739 [cs]. [Online]. Available: <http://arxiv.org/abs/2306.02739>
- [9] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, "Code as Policies: Language Model Programs for Embodied Control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 9493–9500.
- [10] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, "ProgPrompt: Generating Situated Robot Task Plans using Large Language Models," Sep. 2022, arXiv:2209.11302 [cs]. [Online]. Available: <http://arxiv.org/abs/2209.11302>
- [11] J. Wu, R. Antonova, A. Kan, M. Lepert, A. Zeng, S. Song, J. Bohg, S. Rusinkiewicz, and T. Funkhouser, "TidyBot: Personalized Robot Assistance with Large Language Models," May 2023, arXiv:2305.05658 [cs]. [Online]. Available: <http://arxiv.org/abs/2305.05658>
- [12] D. Beßler, R. Porzel, M. Pomarlan, A. Vyas, S. Höffner, M. Beetz, R. Malaka, and J. Bateman, "Foundations of the Socio-physical Model of Activities (SOMA) for Autonomous Robotic Agents," *arXiv:2011.11972 [cs]*, Nov. 2020, arXiv: 2011.11972. [Online]. Available: <http://arxiv.org/abs/2011.11972>
- [13] "Speech-to-Text: Automatic Speech Recognition." [Online]. Available: <https://cloud.google.com/speech-to-text>
- [14] M. Honnibal and I. Montani, "spaCy: Industrial-strength Natural Language Processing (NLP) in Python." [Online]. Available: <https://github.com/explosion/spaCy>
- [15] J. Raible, C. Braun, and M. Huber, "Automatic Path Planning for Robotic Grinding and Polishing Tasks based on Point Cloud Slicing," in *ISR Europe 2023; 56th International Symposium on Robotics*, 2023, pp. 382–389.
- [16] X. Zeng, G. Zhu, Z. Gao, R. Ji, J. Ansari, and C. Lu, "Surface polishing by industrial robots: a review," *The International Journal of Advanced Manufacturing Technology*, vol. 125, no. 9-10, pp. 3981–4012, 2023.