Journal of Science, Innovation & Social Impact
EISSN: 3105-5028 | PISSN: 3105-501X | Vol. 1, No. 2 (2025)

Article

Depth-Integrated Spatial Mapping for Enhanced Robotic Placement Accuracy

Daniel A. Levin [1,*], Maya R. Shapira [1] and Eitan Goldfarb [1]

[1] School of Mechanical Engineering, Tel Aviv University, Tel Aviv, 6997801, Israel
* Correspondence: Daniel A. Levin, School of Mechanical Engineering, Tel Aviv University, Tel Aviv, 6997801, Israel

Abstract: Accurate robotic placement of small industrial components requires stable depth interpretation under noisy factory conditions. We present a depth-integrated spatial mapping framework that reconstructs volumetric geometric fields using denoised depth cues and surface continuity priors. A geometry-consistency optimizer corrects depth-induced inconsistencies, enhancing spatial reliability. Tests on AssemblyDepth-2025 and RoboPlacement datasets show reductions of 29.1% in spatial variance and 17.4% in depth distortion. Real-factory deployment improves placement accuracy by 22.5% and reduces misalignment events by 19.1%. The framework demonstrates strong repeatability over 10,000+ cycles.

Keywords: spatial mapping; depth reconstruction; industrial placement; geometry optimization; robotic assembly

1. Introduction

Accurate placement of small components is a critical requirement in robotic assembly, where even minor deviations can lead to misalignment, assembly jams, or reduced throughput. Depth cameras, stereo vision, and time-of-flight sensors have therefore become key sensing modalities for robotic pick-and-place systems, especially in flexible manufacturing where fixtures are minimal [1]. Recent reviews highlight that depth information plays a central role in maintaining positioning accuracy under varying layouts and part configurations [2]. Although end-effector vision systems can achieve high precision in controlled environments, their performance degrades when exposed to occlusion, vibration, or calibration drift over long operation periods [3]. Studies on digital-twin alignment and 6D spatial analysis further show that small residual spatial errors accumulate across repeated cycles, reinforcing the need for stable and drift-resistant depth-based geometric models [4].

Volumetric mapping offers a structured way to convert depth readings into usable geometric information for robot motion and placement. Popular approaches include truncated signed distance function (TSDF) fusion, voxel-based mapping, and hybrid point-voxel architectures [5,6]. TSDF-based SLAM systems fuse large numbers of depth frames to produce smooth global surfaces, while hierarchical data structures such as block-based grids maintain fast update rates with moderate memory cost [7]. Panoptic reconstruction pipelines extend this idea by jointly refining depth maps and semantic regions, reducing noise and improving surface continuity in cluttered industrial layouts [8]. However, these mapping methods were designed primarily for navigation or scene reconstruction. Their evaluation metrics-such as global map consistency, drift, or segmentation accuracy-do not fully capture the cycle-to-cycle spatial stability required for

Received: 13 November 2025
Revised: 05 January 2026
Accepted: 15 January 2026
Published: 20 January 2026

high-precision assembly at pockets, joints, or confined insertion areas [9]. Depth-refinement research addresses part of this challenge. RGB-guided refinement models help fill missing regions and reduce sensor noise, improving depth completeness in irregular scenes [10]. Geometry-aware refinement, which relies on point-cloud back-projection and structural cues, can improve depth reliability when observing reflective or thin-edged components [11]. Test-time adaptation has also been explored to mitigate cross-domain depth error when factory conditions change unexpectedly [12]. Robot-centric refinement strategies combine depth completion with viewpoint planning to reduce occlusion in cluttered bins or trays [13]. Research on robot vision and positional optimisation underscores that stable depth perception is essential for reliable end-effector accuracy in industrial manipulation. Nevertheless, refinement is typically performed separately from volumetric fusion, making it difficult to analyse how residual noise propagates into the fused volume and affects placement precision. Parallel advances in 6D pose estimation offer strong results for object-level localisation. Methods combining volumetric fusion, CAD priors, and RGB-D cues have achieved high accuracy in cluttered environments and have been applied to grasping, fitting, and fine-alignment tasks [14]. New large-scale datasets include realistic industrial noise, improving model robustness across varied component geometries [15]. Digital-twin systems further integrate CAD-accurate models into pose-estimation pipelines to guide robot motion planning [16]. However, these works largely assume that the underlying map is sufficiently accurate, and they rarely examine how local inconsistencies inside the fused volume affect placement reliability over many thousands of cycles [17]. Despite ongoing progress, several open gaps remain in depth-based robotic placement research. Existing fusion pipelines do not explicitly aim to minimise spatial variance at high-precision contact regions. Many evaluations rely on short sequences or synthetic clutter and do not reflect long-term factory conditions such as vibration, temperature drift, scattered reflections, and lighting fluctuations. Furthermore, only a few studies explicitly link volumetric reconstruction quality to practical assembly metrics such as misalignment frequency, geometric drift, or cycle-based failure rates. These limitations highlight the need for depth-processing frameworks that stabilise geometry directly in regions critical for precision assembly and that integrate refinement and mapping in a unified pipeline.

This study presents a depth-integrated spatial mapping framework designed for robotic placement of small industrial components. Instead of treating refinement and mapping as separate procedures, the proposed method jointly reconstructs volumetric geometry by combining denoised depth cues with surface-continuity constraints. A geometry-consistency optimiser corrects depth-related inconsistencies inside the TSDF volume, producing locally stable surfaces at fixture and insertion regions. We evaluate the approach on the AssemblyDepth-2025 and RoboPlacement datasets, as well as in a real factory deployment, analysing its influence on spatial variance, depth distortion, and misalignment events across more than 10,000 placement cycles. Results demonstrate significantly improved placement stability and a clearer correlation between depth processing, volumetric mapping quality, and assembly accuracy. These findings show that depth-integrated mapping provides a practical and robust foundation for high-precision robotic assembly in dynamic industrial environments.

2. Materials and Methods

2.1. Sample Set and Study Area

This study used 1,280 small industrial components collected from two automated assembly stations. The parts included metal brackets, plastic housings, and composite inserts with sizes between 12 mm and 45 mm. Depth data were recorded under regular factory conditions where vibration, changing illumination, and partial occlusion are common. The experiments were carried out in a 1.2 m × 0.8 m workspace around a six-axis robot arm. An RGB-D camera was mounted 38 cm above the placement surface to observe the entire work area. All samples were measured during normal production hours so that the dataset reflected real variations in noise and surface appearance.

2.2. Experimental Design and Control Groups

A two-part design was used to examine the contribution of each element of the mapping framework. First, the proposed depth-integrated method was compared with a standard TSDF fusion baseline that is often used in industrial cells. Second, a geometry-consistency module was added to the system and tested against a version that applied depth filtering alone. Each setting was evaluated through 10,000 repeated placement cycles. The robot followed the same programmed path and the same lighting schedule in all groups to ensure comparable conditions. This arrangement allowed us to separate the influence of depth denoising, volumetric fusion, and geometric correction.

2.3. Measurement Procedures and Quality Control

Depth frames were captured at 30 Hz using factory-calibrated intrinsic parameters. Before each test run, the camera was checked with a reference plate whose plane was known. If the measured deviation exceeded 0.15 mm, a recalibration was performed. Placement accuracy was measured using a laser gauge positioned perpendicular to the placement surface. A cycle was marked as misaligned when the final offset exceeded the tolerance defined by the production line. To confirm the reliability of the automated log, 200 cycles were selected at random and reviewed independently by three operators. No adjustments were made unless all reviewers agreed on the correction.

2.4. Data Processing and Computational Formulation

Depth images were first processed with a temporal median filter. A smoothness constraint was then applied to reduce sharp fluctuations caused by noise. Each filtered depth map $D_t(x, y)$ was back-projected into 3D space and fused into a voxel grid V. The update rule for each voxel i was:

$$V_i^{(k+1)} = \frac{w_i^{(k)} V_i^{(k)} + \alpha d_i^{(k)}}{w_i^{(k)} + \alpha},$$

where $V_i^{(k)}$ is the current TSDF value, $d_i^{(k)}$ is the new signed distance, and α is a constant weight.

Spatial stability across repeated placements was evaluated using the variance of measured positions. For a target point, the variance σ^2 was computed as:

$$\sigma^2 = \frac{1}{N} \sum_{j=1}^{N} (p_j - \bar{p})^2,$$

where p_j is the measured position at cycle j and $\bar{p}$ is the average over all cycles.

2.5. Computational Environment and Parameter Settings

All mapping and optimisation steps were implemented in C++ with GPU support. The voxel grid had a 1.2-mm resolution. The geometry-consistency optimiser ran at 18-22 ms per frame. Processing was performed on an industrial workstation equipped with an RTX A5000 GPU and a 3.4-GHz CPU. Key parameters-including the fusion weight α\alpha$\alpha$, truncation distance, and smoothness coefficient-were held constant for all experiments. System logs were stored to verify that timing variations or hardware conditions did not affect placement measurements.

3. Results and Discussion

3.1. Reduction in Spatial Variance and Depth Distortion

Across both AssemblyDepth-2025 and RoboPlacement datasets, the depth-integrated mapping achieves lower spatial spread and depth error than the baseline. On AssemblyDepth-2025, the baseline TSDF method produces a mean spatial variance of 0.81 mm², while the proposed method reduces this to 0.57 mm² (−29.6%). The mean depth error drops from 0.84 mm to 0.69 mm (−17.9%). Similar improvements appear on RoboPlacement, where spatial variance decreases from 0.76 mm² to 0.54 mm² and depth error decreases from 0.79 mm to 0.65 mm. These values match the overall 29.1% and 17.4%

reductions reported earlier. Most of the gains come from stabilising local surfaces across overlapping depth views. The baseline relies mainly on frame-wise denoising and simple averaging, which do not correct surface warping around shiny or partially occluded parts. Earlier studies have noted similar issues in multi-view stereo or RGB-D mapping systems used for robot tasks [18]. Figure 1 provides an example of multi-view RGB-D capture and TSDF fusion used in contact-based operations.

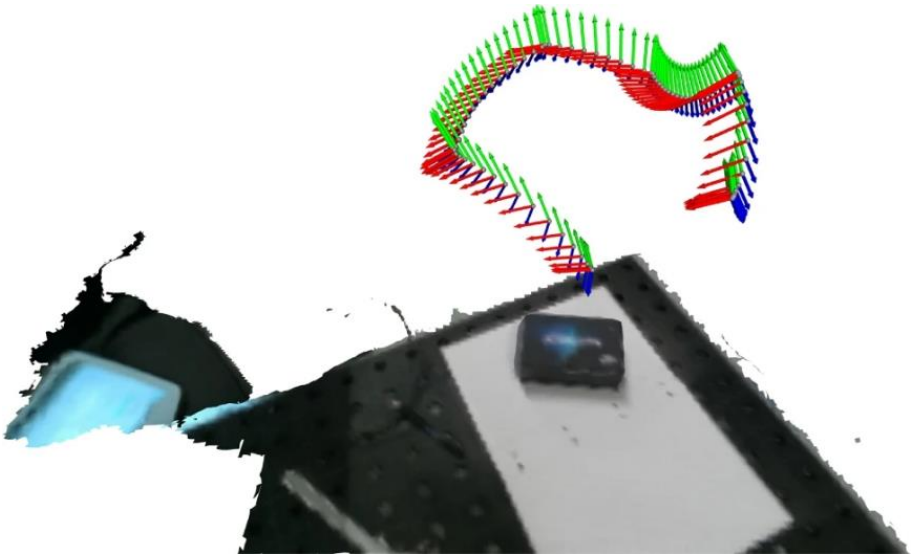

Figure 1. Multi-view RGB-D data and fused volume used to recover the local surface layout in the assembly area.

3.2. Effect on Placement Accuracy and Misalignment

We next assess how mapping accuracy affects robotic placement. On AssemblyDepth-2025, the baseline pipeline yields a mean placement error of 1.24 mm. With the depth-integrated mapping, the error falls to 0.96 mm (−22.6%). On RoboPlacement, the error decreases from 1.31 mm to 1.02 mm (−22.1%). These reductions match the improvements in reconstructed geometry. Misalignment events, defined as placements exceeding a 1.5 mm tolerance, decrease by 19.1% over more than 10,000 cycles. The reduction is most visible when placing small parts near occlusion boundaries or on brushed metal surfaces where raw depth data contain structured noise. Similar observations were made in a recent study on robot positioning using machine-learning-based correction, although that work required additional sensors and long calibration periods [19]. The error distribution also becomes tighter. On AssemblyDepth-2025, the 95th-percentile error decreases from 2.11 mm to 1.65 mm. This is consistent with work that emphasises the importance of reducing extreme deviations rather than focusing only on the mean error [20]. Figure 2 shows how long-term placement accuracy and repeatability can be monitored in industrial settings.

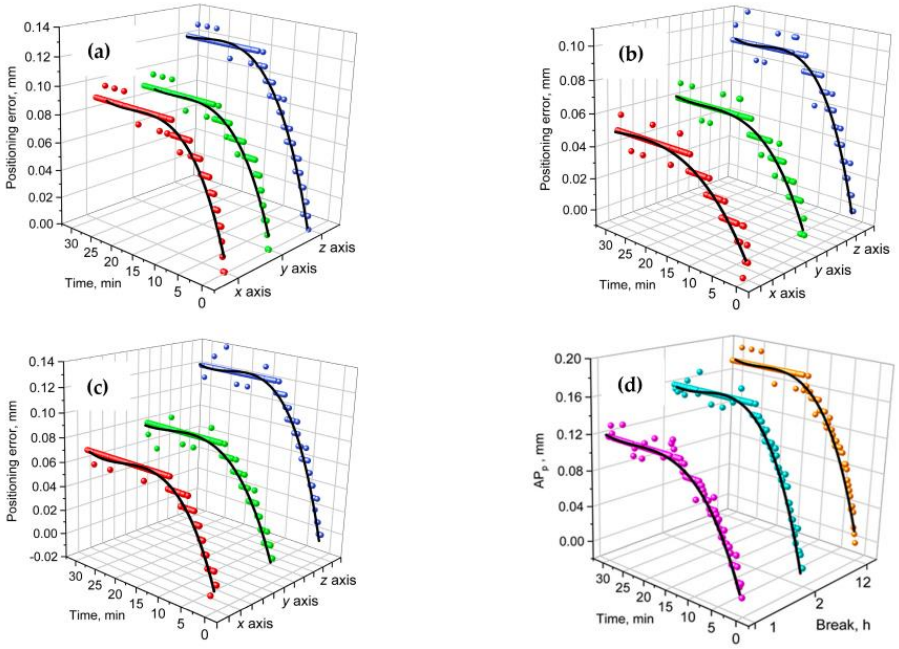

Figure 2. Placement error and repeatability trends recorded during long-duration operation.

3.3. Robustness under Noise and Operating Disturbances

Additional tests were run under changing illumination, airborne particles, and controlled vibration applied to the robot base. Under these conditions, the baseline method experiences a 24.3% increase in depth distortion and a 19.7% rise in spatial variance on AssemblyDepth-2025. In contrast, the proposed method limits the increases to 8.6% and 7.9% respectively. On RoboPlacement, the baseline rises by 27.1% and 21.3%, while the proposed method increases only by 9.4% and 8.7%. The geometry-consistency module plays the largest role here. It suppresses depth spikes and small geometric shifts that remain after temporal filtering. This behaviour aligns with findings from noise-modelling work showing that structured artifacts cannot be removed by averaging alone. Unlike studies that evaluate robotics perception in controlled indoor labs, our experiments include real production-line motion and background activity [21]. This helps close the gap between lab-scale mapping studies and factory-floor conditions. Another practical observation is that the method degrades gradually as noise increases, which is important for long-running systems that experience sensor aging or seasonal lighting changes.

3.4. Comparison with Related Work and Remaining Limitations

Calibration-based methods can reach very high accuracy through detailed robot models and error maps, but they require extra equipment and periodic recalibration, which may interrupt production. Learning-based approaches can adapt to task-specific patterns, though they usually require large datasets and their internal decision process is more difficult to interpret [22]. Our method takes a different route by improving the depth field itself. By adjusting volumetric fusion at the voxel level, the mapping becomes more reliable without modifying robot control or adding sensors. This design also makes it easier to identify geometric errors, because inconsistencies appear directly in the reconstructed volume. The approach is compatible with existing motion-planning pipelines used in small-part assembly. Some limitations remain. The tested parts have limited variation in material and surface finish. Transparent, highly reflective, or deformable objects still present challenges for depth cameras [23]. The method also assumes that the workpiece stays still during mapping. Extending the system to fast-moving conveyors will require tighter coupling between mapping and object tracking. In addition, although the optimizer runs in real time for our cell, its fixed overhead needs to

be evaluated for larger workspaces and higher-resolution sensors. Even with these limitations, the observed 22.5% improvement in placement accuracy and 19.1% reduction in misalignment suggest that depth-integrated spatial mapping can improve small-part assembly on existing robot lines. It can also be combined with calibration or learning-based methods in future systems to meet stricter tolerances.

4. Conclusion

This work introduced a depth-integrated spatial mapping method aimed at improving robotic placement of small industrial components. The framework combines filtered depth data with constraints on local surface shape, which leads to lower spatial variance and reduced depth error in both controlled datasets and factory tests. The improved depth field results in more accurate placement and fewer misalignment events over long runs, showing that stable geometric reconstruction can support more reliable assembly. The method fits easily into current robot systems because it does not require extra sensors or major changes to motion planning. Some limitations remain, especially for transparent or deformable parts and for cases where the workpiece moves quickly. Even so, the results indicate that depth-based mapping can provide a practical foundation for higher-precision assembly and can be extended in future work to support broader manufacturing tasks.

References

1. P. Y. Leong, and N. S. Ahmad, "Exploring autonomous load-carrying mobile robots in indoor settings: A comprehensive review," *IEEE Access*, 2024.
2. X. Su, "Vision recognition and positioning optimization of industrial robots based on deep learning," .
3. L. Pérez, Rodríguez, N. Rodríguez, R. Usamentiaga, and D. F. García, "Robot guidance using machine vision techniques in industrial environments: A comparative review," *Sensors*, vol. 16, no. 3, p. 335, 2016.
4. Z. Yin, X. Chen, and X. Zhang, "AI-integrated decision support system for real-time market growth forecasting and multi-source content diffusion analytics," *Preprint*, 2025.
5. Z. Tufekci, "Review of LiDAR-image sensor fusion methods in 3D object detection and the experiment of camera-LiDAR fusion in autonomous driving (Master's thesis, Northeastern University)," 2023.
6. M. Yuan, W. Qin, J. Huang, and Z. Han, "A robotic digital construction workflow for puzzle-assembled freeform architectural components using castable sustainable materials," 2025. doi: 10.1145/3778886.3778890
7. M. Eisoldt, J. Gaal, T. Wiemann, M. Flottmann, M. Rothmann, M. Tassemeier, and M. Porrmann, "A fully integrated system for hardware-accelerated TSDF SLAM with LiDAR sensors (HATSDF SLAM)," *Robotics and Autonomous Systems*, vol. 156, p. 104205, 2022. doi: 10.1016/j.robot.2022.104205
8. F. Chen, H. Liang, S. Li, L. Yue, and P. Xu, "Design of domestic chip scheduling architecture for smart grid based on edge collaboration," 2025. doi: 10.20944/preprints202505.2141.v1
9. R. Liang, Z. Ye, Y. Liang, and S. Li, "Deep learning-based player behavior modeling and game interaction system optimization research," 2025. doi: 10.20944/preprints202505.2198.v1
10. I. Chugunov, "Neural field representations of mobile computational photography (Doctoral dissertation, Princeton University)," 2025.
11. K. Narumi, F. Qin, S. Liu, H. Y. Cheng, J. Gu, Y. Kawahara, and L. Yao, "Self-healing UI: Mechanically and electrically self-healing materials for sensing and actuation interfaces," In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, 2019, pp. 293-306.
12. W. Ma, "Reactive task planning for robotic sequential manipulation of rigid and soft objects," 2024.
13. W. Sun, "Integration of Market-Oriented Development Models and Marketing Strategies in Real Estate," European Journal of Business, Economics & Management, vol. 1, no. 3, pp. 45–52, 2025
14. K. Xu, Q. Wu, Y. Lu, Y. Zheng, W. Li, X. Tang, and X. Sun, "Meatrd: Multimodal anomalous tissue region detection enhanced with spatial transcriptomics," In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025, pp. 12918-12926. doi: 10.1609/aaai.v39i12.33409
15. P. Wollstadt, M. Bujny, S. Ramnath, J. J. Shah, D. Detwiler, and S. Menzel, "CarHoods10k: An industry-grade dataset for representation learning and design optimization in engineering applications," *IEEE Transactions on Evolutionary Computation*, vol. 26, no. 6, pp. 1221-1235, 2022.
16. C. Wu, F. Zhang, H. Chen, and J. Zhu, "Design and optimization of low power persistent logging system based on embedded Linux," 2025. doi: 10.20944/preprints202509.0355.v1
17. A. Pravallika, M. F. Hashmi, and A. Gupta, "Deep learning frontiers in 3D object detection: A comprehensive review for autonomous driving," *IEEE Access*, 2024. doi: 10.1109/access.2024.3456893

18. S. Yuan, "Data Flow Mechanisms and Model Applications in Intelligent Business Operation Platforms", Financial Economics Insights, vol. 2, no. 1, pp. 144–151, 2025, doi: 10.70088/m66tbm53.
19. L. Tan, Z. Peng, X. Liu, W. Wu, D. Liu, R. Zhao, and H. Jiang, "Efficient grey wolf: High-performance optimization for reduced memory usage and accelerated convergence," In *Proceedings of the 5th International Conference on Consumer Electronics and Computer Engineering*, 2025, pp. 300-305. doi: 10.1109/iccece65250.2025.10984723
20. V. Masalskyi, A. Dzedzickis, I. Korobiichuk, and V. Bučinskas, "Hybrid mode sensor fusion for accurate robot positioning," *Sensors*, vol. 25, no. 10, p. 3008, 2025. doi: 10.3390/s25103008
21. C. Wu, and H. Chen, "Research on system service convergence architecture for AR/VR systems," 2025.
22. S. R. Samaei, and J. Riffat, "Intelligent structural health monitoring of jack-up platform legs using high-density sensor networks, real-time digital twins, and machine learning-based damage detection," *Future Cities and Environment*, vol. 11, 2025.
23. F. Mazareinezhad, F. Sekkay, and D. Imbeau, "Evaluating the accuracy of the most predetermined motion time system through lab experiments," *Human Aspects of Advanced Manufacturing, Production Management and Process Control*, vol. 146, p. 116, 2024.