

Accurate 3D Reconstruction of TGVs Using Hybrid Mesh–NURBS Modeling with Z-Axis Constraints and RL-Based Optimization

Mu Fu¹, Eugene Su¹, Chao-Ching Ho^{1,*}, Liang-Chia Chen², Pradeep Kumar¹, Ming Chieh Jean¹

¹National Taipei University of Technology, Taipei, Taiwan, hochao@ntut.edu.tw

²National Taiwan University, Taipei, Taiwan, lchen@ntu.edu.tw

Abstract — Through-glass vias (TGVs) are critical for high-density semiconductor packaging; however, their internal waist geometry remains difficult to measure using non-destructive techniques. This study proposes an integrated 3D reconstruction framework combining NURBS-based parametric modeling, triangular mesh fitting, and Z-axis boundary constraints to ensure geometric accuracy. Experimentally acquired point cloud data are processed to generate high-fidelity 3D models, while a reinforcement learning (RL)-based digital twin simulation environment is developed to optimize inspection conditions and generate synthetic datasets. The proposed approach demonstrates high reconstruction accuracy and provides a robust solution for advanced metrology and inspection of TGV structures.

Keywords — TGV, 3D reconstruction, semiconductor, triangular mesh, reinforcement learning, digital twin model, industry, metrology and synthetic data.

I. INTRODUCTION

The rapid adoption of glass substrates in the semiconductor packaging industry has become increasingly important. Through-glass vias (TGVs) have emerged as one of the key enabling technologies for high-density semiconductor packaging owing to their excellent electrical and thermal properties, including low signal loss, high integration capability, high electrical resistivity, superior thermal stability, and a tunable coefficient of thermal expansion (CTE)[1]. With the rapid advancement of heterogeneous integration and advanced packaging technologies, the demand for precise and reliable characterization of TGV structures has become increasingly critical to ensure device performance and long-term reliability[1], [2]. However, the thermal stress induced by the mismatch of the CTE among multi-material encapsulation structures presents a severe threat to interfacial integrity[3]. Consequently, establishing robust filling workflows and non-destructive geometric inspection schemes has become paramount to mitigate thermo-mechanical failures [3],[4]. However, the inherent

geometry of TGVs presents significant metrological challenges. TGVs typically exhibit an hourglass-shaped or double-tapered profile, in which the diameter gradually narrows toward the waist region due to process-induced variations such as mass-transport loading effects during deep wet etching[5]. This internal geometry is difficult to characterize accurately using conventional non-destructive measurement techniques, particularly when high precision is required in industrial inspection environments. To overcome these limitations, advanced three-dimensional (3D) holographic optical metrology (HOM) based on digital holography has been investigated for full-field quantitative phase imaging [6], while photometric stereo combined with hybrid field microscopy has emerged to visualize cross-sectional textures and edge abnormalities [7]. Consequently, accurate reconstruction and quantitative analysis of the complete three-dimensional geometry remain challenging but critical for quality control.

In recent studies, TGV structures have commonly been approximated using double-cone geometric models defined by key parameters, including the top, waist, and bottom diameters along with the total depth [2]. Previous studies have demonstrated that the measurement error of the waist diameter can be effectively controlled, highlighting the importance of accurate geometric parameter estimation [2]. These geometric characteristics provide a solid foundation for developing advanced reconstruction and parametric fitting methods [6]. To facilitate three-dimensional reconstruction, several data-generation frameworks utilizing game engines, such as Unity or NVIDIA Omniverse Replicator, have been developed to generate high-fidelity synthetic data and automate instance annotations [8]. To bridge the gap between virtual simulation and real-world fabrication, recent studies have integrated advanced 3D imaging metrology [6], [7], thermo-mechanical reliability simulations [3], and novel filling workflows [4] to comprehensively evaluate structural integrity and mitigate process-induced defects [3]. Mohanty et al. [9] proposed a reinforcement learning (RL)-optimized digital twin framework to generate realistic synthetic datasets for




Fig. 1. Workflow of TGV 3D reconstruction with mesh fitting, Z-axis constraints, and RL-based simulation.

automated surface defect detection in titanium spacer rings. By integrating Blender-based three-dimensional modeling with a proximal policy optimization (PPO) algorithm, their approach automatically adjusted optical and material rendering parameters to produce high-fidelity training data. Although their experimental results demonstrated that combining synthetic and real datasets significantly improved flaw inspection compared with using limited real data alone, their methodology primarily focused on two-dimensional image generation for surface defect detection. Consequently, it remains insufficient for the three-dimensional geometric reconstruction or accurate internal profiling required for complex structures like TGVs.

Motivated by these challenges and recent advances, this study proposes a robust and integrated three-dimensional reconstruction framework that combines triangular mesh representation with NURBS-based parametric modeling. By integrating data-driven reconstruction with geometry-based constraints, the proposed approach enables accurate modeling of complex TGV profiles while preserving geometric consistency and boundary integrity [2], thereby advancing the automatic generation of parametric geometric digital twins (gDTs) directly from unstructured point clouds [10]. Furthermore, a deep reinforcement learning (RL)-based optimization strategy backbone by the proximal policy optimization (PPO) algorithm [11] is incorporated to refine the multi-dimensional optical and material attributes within the virtual simulation environment, enabling the generation of highly realistic synthetic inspection data through physics-based

differentiable rendering theories [12]. Unlike existing approaches, the proposed framework jointly addresses accurate TGV geometric reconstruction and physics-consistent synthetic data generation within a unified framework, thereby advancing non-destructive TGV metrology and supporting the development of next-generation semiconductor packaging inspection systems. The main contributions of this work are summarized as follows:

1. A novel 3D reconstruction framework for TGVs based on triangular mesh and NURBS modeling is proposed.
2. An RL-based digital twin framework is developed to generate realistic synthetic TGV inspection data.
3. A unified non-destructive inspection solution is established for accurate TGV metrology and intelligent semiconductor packaging inspection.

The rest of this paper is organized as follows. Section II describes the methodology and experimental setup. Finally, Section III concludes the paper and discusses future research directions.

II. METHODOLOGY AND EXPERIMENTAL RESULTS

This section describes the proposed approach for establishing a reliable geometric reference. Based on the point cloud data, we analyzed its key dimensions: top, waist, bottom diameter, and total height. Since obtaining

real data is difficult, A parametric TGV template was constructed in Rhino 8 using NURBS-based geometry to

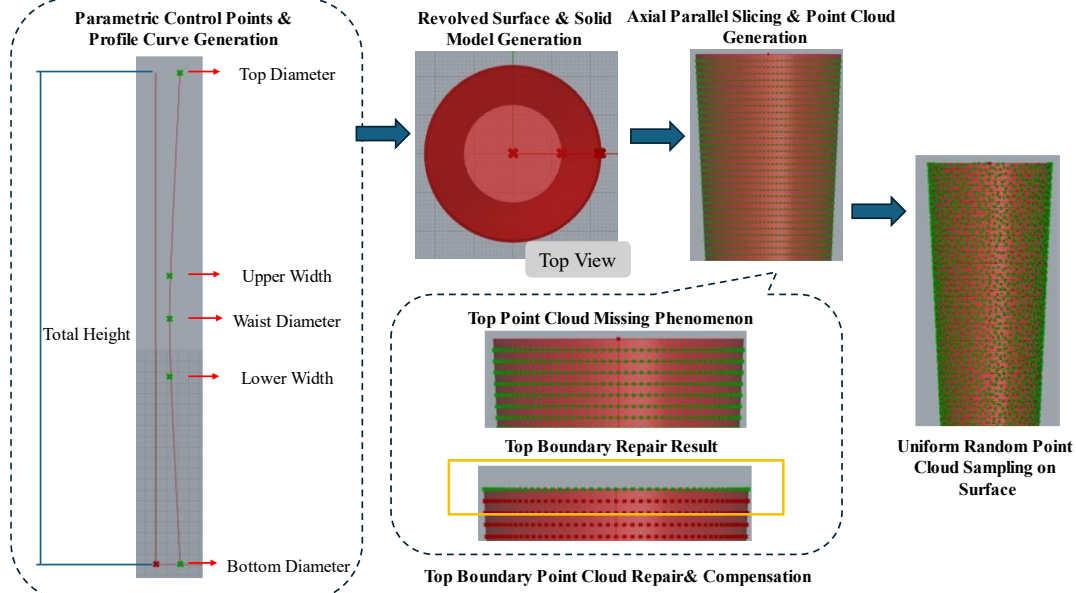


Fig. 2. Flowchart of TGV parametric virtual model generation and point cloud processing based on Rhino/Grasshopper.

to increase the data points. This mathematically rigorous framework defines the key structural dimensions, including a top diameter of $44.57 \mu\text{m}$, a waist diameter of $25.22 \mu\text{m}$, a bottom diameter of $45.49 \mu\text{m}$, and a total depth of $460.00 \mu\text{m}$. Based on Rhino's Grasshopper feature to build a mathematically based, visualized workflow was built. This idealized model serves as a baseline for evaluating reconstruction accuracy and ensures consistency throughout the proposed workflow. As shown in Fig. 1, experimentally acquired point cloud data in x , y , and z -coordinates were converted into a format for compatibility with 3D modelling software. Rhino's Grasshopper workflow is illustrated in Fig. 2. The dataset consists of 46 scanning layers. Subsequently, the Model 1 employed the Poisson reconstruction method to remove the upper and lower extra boundaries, thereby achieving the NURBS surface TGV structure with an estimated volume. Another model 2 initial process cylindrical model was generated with a circumferential resolution of 64 vertices and a vertical resolution of 45 segments to capture fine geometric variations. Then, a global Z -axis boundary-locking strategy was implemented to precisely define the spatial configuration of the model by fixing the top and bottom boundaries. The geometric constraints are defined as: $Z_{\text{center}} = \frac{Z_{\text{max}} + Z_{\text{min}}}{2}$ and $H = Z_{\text{max}} - Z_{\text{min}}$. These constraints prevent geometric distortion during the fitting process and ensure boundary integrity. Subsequently, a K-Nearest Neighbors (KNN) algorithm was applied to map mesh vertices to the nearest point cloud coordinates, enabling accurate fitting of the actual TGV geometry.

Our proposed framework successfully converted

339,892 raw point cloud data into a structured triangular mesh comprising 2,946 vertices and 3,008 faces. The reconstructed TGV volume was estimated as $440,555.48 \mu\text{m}^3$, demonstrating the accuracy, robustness, and practical applicability of the proposed method. The proposed method shows improved geometric consistency compared to conventional cylindrical fitting approaches. Furthermore, a virtual opto-mechanical simulation system was developed, consisting of an 8×6 TGV array with an orthographic camera positioned at a distance of $16,500 \mu\text{m}$ from the substrate.

The data generation architecture developed in this study is illustrated in Fig. 2. The entire pipeline integrates parametric CAD modeling, regular slicing segmentation, boundary feature repair, and random surface sampling techniques. First, as shown in the section designated as "Parametric Control Points & Profile Curve Generation", a parametric control zone is established within the Rhino/Grasshopper platform. This zone allows for the dynamic input of core geometric dimensions, including Top Diameter, Waist Diameter, Bottom Diameter, and Total Height. To simulate the asymmetric deformation of TGV holes encountered in actual fabrication processes, a "Waist Height Ratio" parameter is specifically introduced to adjust the relative height along the Z -axis. Additionally, the transition curves are governed by the Upper Width and Lower Width control parameters. Based on geometric logic, the coordinates of five critical control points are computed and sequentially structured. Utilizing the NURBS curve

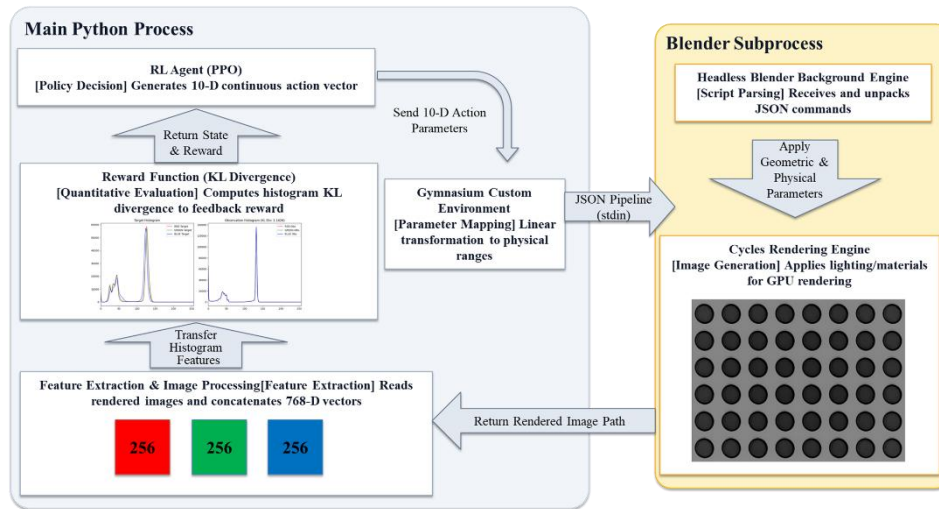


Fig. 3. Dual-process headless architecture and closed-loop reinforcement learning workflow for TGV digital twin optical alignment. Rhino/Grasshopper.

component under a degree-2 (Degree = 2) geometric definition, two continuous parabolic outlines are generated. geometric definition, two continuous parabolic outlines are generated. Subsequently, a symmetrical 3D virtual surface model of the TGV is constructed via the revolution surface (RevSrf) algorithm. The top geometric features of this model can be examined from the viewpoint depicted in "Revolved Surface & Solid Model Generation" with the "Top View" perspective.

To convert the continuous surface model into discrete point clouds that conform to physical measurement characteristics, this study implements a dual-track sampling mechanism. In the first track, the system isolates the precise circumference radius at the absolute top of the model ($Z = \text{Total Height}$) and performs uniform angular sampling to independently generate a dimensionally accurate top ring point cloud, which serves as the baseline for subsequent feature compensation. Concurrently, the second track, designated as "Axial Parallel Slicing & Point Cloud Generation", utilizes a slicing algorithm to segment the TGV virtual surface into equidistant horizontal slices moving vertically upward along the Z-axis. A fixed number of points are uniformly sampled along the perimeter of each slice to construct an organized, layered axial point cloud.

However, conventional CAD slicing operations frequently suffer from the loss of geometric features at the boundaries. As demonstrated in the magnified comparative region "Top Boundary Point Cloud Repair & Compensation", relying solely on equidistant slicing algorithms often leads to a discrepancy where the final slice layer cannot align precisely with the top edge of the model. This occurs either because the total height is not perfectly divisible by the slicing interval or due to the "inclusive start, exclusive end" boundary logic inherent

in the algorithm. Consequently, this leads to the omission of the uppermost ring point cloud (indicated by the red dot region), as shown in "Top Point Cloud Missing Phenomenon". This boundary effect severely compromises the identification accuracy of top diameters in subsequent critical dimension (CD) algorithms. To address this issue, the proposed pipeline forcibly imports the pre-computed top ring point cloud upon completion of the slicing process. Through geometric superposition, the accurate boundary points are precisely compensated back to the top of the model ($Z = \text{Total Height}$), yielding the complete feature displayed inside the yellow box of "Top Boundary Repair Result" and effectively resolving the data distortion problem.

Finally, to more realistically simulate the random scattering characteristics observed in optical measurements, the system executes a surface uniform random sampling algorithm within the region labeled "Uniform Random Point Cloud Sampling on Surface". This process scatters a specified total number of random points across the TGV surface, with a random seed (Seed) applied to guarantee experimental reproducibility. Ultimately, the system integrates the repaired regular axial slicing point clouds with the random surface points, discards the top and bottom capping faces, and outputs the non-closed TGV hole synthetic point cloud data that mirrors actual physical measurement morphology. This dataset serves as a standardized training database for subsequent high-efficiency 3D point cloud algorithms. This simulated structure connected to a digital twin RL agent to optimize rendering parameters and align simulated outputs with real inspection conditions, enabling the generation of realistic synthetic images.

To achieve physics-level optical alignment between empirical micrographs captured by the real microscope (VK3000) and synthesized images generated within a

virtual environment, a high-fidelity digital twin scene representing a through-glass via (TGV) substrate was constructed using the Blender Cycles rendering engine. However, embedding deep reinforcement learning frameworks (such as PyTorch or Stable-Baselines3) directly alongside Blender's native Python API (bpy) within a single-threaded or standard multi-threaded runtime frequently triggers fatal memory corruption and application deadlocks. To circumvent this critical computational bottleneck, this study implements a robust, closed-loop system driven by a dual-process headless communication architecture, as structurally illustrated in Fig. 3. Under this framework, the reinforcement learning agent and its customized Gymnasium environment execute within the primary Python process, whereas Blender operates strictly as an independent subprocess configured in headless mode (background execution without a graphical user interface). Inter-process communication is sustained via standard input and output streams (stdin/stdout), establishing a lightweight JSON data pipeline. Upon the generation of a continuous vector of physical action parameters by the agent, the Gymnasium environment encapsulates these variables into text-based instructions and pipes them to the Blender subprocess. Blender then applies these physical configurations within its Cycles photorealistic engine, renders the corresponding TGV image, and returns the temporary file path back to the main process for real-time optical distribution analysis. This design minimizes interface overhead and guarantees high deterministic stability throughout the automated closed-loop training cycle.

To navigate the high-dimensional and non-linear search space of optomechanical properties, the parameter alignment task is formulated as a Markov Decision Process (MDP) and optimized using Proximal Policy Optimization (PPO) proposed by Schulman et al. [11]. PPO is chosen for its superior sample complexity, implementation simplicity, and robust performance in continuous control domains compared to vanilla policy gradients and trust region methods. Traditional policy gradient estimators often suffer from destructively large updates when executing multiple optimization steps on a single trajectory. In this study, the value function is parameterized via a shared network architecture. The structural configurations of the state, action, and reward spaces within this PPO framework are defined as follows:

Observation Space: Directly employing raw pixel intensities from the rendered frames introduces prohibitive dimensionality and high sensitivity to trivial geometric noise. To compress the visual state while preserving macroscopic illumination traits, this study extracts normalized red (R), green (G), and blue (B) tri-channel color histograms. Each channel is uniformly divided into 256 bins, yielding a 768-dimensional

observation vector that robustly encapsulates exposure levels, material reflectivity, and chromatic distributions. **Action Space:** The control interface comprises a 10-dimensional continuous action vector bounded strictly within $[-1.0, 1.0]$. This vector is mapped online to physical ranges governing explicit scene parameters, including background illumination (World Strength), direct light source intensity (Sun Energy), multi-axial light incident angles, glass refractive index (IOR), surface roughness, transmittance, and specular intensity of the glass, alongside the base gray tone, roughness, and specular intensity of the supporting stage.

Reward Function: To quantify the mathematical similarity between the virtual and empirical color distributions, the reward is formulated as the negative Kullback-Leibler (KL) divergence between the rendered histogram distribution P and the real target distribution Q with

$$\text{Reward} = -D_{KL}(Q \parallel P) = -\sum_{i=1}^{768} Q(i) \ln \left(\frac{Q(i)}{P(i) + \epsilon_{num}} \right) \quad (1)$$

,where $\epsilon_{num} = 10^{-8}$ prevents numerical divergence. As the PPO agent maximizes this reward, the virtual optical distribution converges toward the target micrograph. The continuous convergence history of this optimization process is explicitly recorded and visualized in the Fig. 4.




Fig. 4. Convergence curve of the mean episode reward during the PPO-based optomechanical parameter alignment process.

During the initial baseline optimization trials, the parameter space successfully achieved numerical convergence; however, cross-examination of the final outputs revealed minor but persistent discrepancies in structural contrast and localized shadow boundaries. Diagnostic investigation indicated that the idealized TGV micro-hole pitch and spacing specified in the default Blender 3D scene deviated slightly from the physical morphology of the empirical specimen, causing the PPO agent to unnaturally inflate the environment lighting (World Strength) to compensate for geometric shadows. To rectify this spatial mismatch, the mesh primitives defining the micro-hole array within the

Blender environment were re-calibrated to mirror the precise empirical spacing of the empirical specimen. Subsequently, to avoid the computational inefficiency of re-training the agent from a random initialization under the modified geometry, a two-stage transfer learning strategy via incremental weight tuning was deployed. The PPO agent inherited the converged neural weights from the initial training, while the Adam optimizer's learning rate was annealed to a conservative 1×10^{-4} for an additional incremental tuning phase spanning 10,000 steps. This framework allowed the model to leverage its pre-acquired representations of material reflectivity and refractive dynamics, rapidly steering the policy toward a superior global optical equilibrium within the revised geometric boundaries. The shifting trend of color distributions and the reduction of systematic error during this two-stage alignment are comprehensively documented in Fig. 5, illustrating the absolute necessity of integrating precise geometric correction with incremental weight inheritance to break through the fidelity limits of optical digital twins.




Fig. 5. Comparison of the Kullback-Leibler (KL) divergence between empirical micrographs and aligned virtual rendered images.

III. CONCLUSIONS

This study presents a robust workflow for the 3D reconstruction of through-glass vias (TGVs) by integrating NURBS-based modeling, triangular mesh fitting, and Z-axis boundary constraints. The proposed approach ensures high geometric accuracy and surface integrity while efficiently converting point cloud data into reliable 3D models for non-destructive metrology. Furthermore, reinforcement learning (RL)-based simulation digital twin model enables the generation of high-quality synthetic datasets with defect characteristics, which will support future development of deep learning models for automated defect detection and intelligent inspection of TGV structures.

IV. ACKNOWLEDGMENTS

This work was supported by the National Science and Technology Council (NSTC), Taiwan, under Grant No. NSTC 114-2218-E-027-001.

REFERENCES

- [1] Lai, Y., Pan, K., Park, S.: Thermo-mechanical reliability of glass substrate and Through Glass Vias (TGV): A comprehensive review, *Microelectronic Reliability*, Vol. 161, 2024, 115477.
- [2] Lu, K., Lim, Y., Chun, H., Park, J., Lee, C.: Shadow-graphic approach to through glass vias (TGV) metrology and inspection, *Optics and Laser Technology*, Vol. 193, 2026, 114261.
- [3] Zhao, J., Chen, Z., Qin, F., Yu, D.: Thermo-Mechanical Reliability Study of Through Glass Vias in 3D Interconnection, *Micromachines*, Vol. 13, No. 10, 2022, 1799.
- [4] Liu, X., Yu, C., Zhang, Z., Yu, D.: Through Glass via (TGV) Filling Process With Conductive Silver Paste Based on 3D Printing Technology, *Proceeding of the International Conference on Electronic Packaging Technology (ICEPT)*, August 2024, pp. 1-3.
- [5] Xie, Q., et al.: Loading effects mitigation for wet etching of high-aspect-ratio through-glass vias in thick glass, *Social Science Research Network*, Rochester, NY, 2026, 6446207.
- [6] Huang, C.H., Cheng, C.J.: Through glass via detection and inspection by holographic optical metrology, *Proceeding of Metrology, Inspection, and Process Control XXXIX*, San Jose, United States, April 2025, p. 56.
- [7] Sung, G.: Advanced 3D Imaging Approach to TSV/TGV Metrology and Inspection Using Only Optical Microscopy, *arXiv preprint*, arXiv:2505.04913, 2025.
- [8] Zhang, J., Fukuda, T., Yabuki, N.: Automatic generation of synthetic datasets from a city digital twin for use in the instance segmentation of building facades, *Journal of Computational Design and Engineering*, Vol. 9, No. 5, 2022, pp. 1737-1755.
- [9] Mohanty, S., Su, E., Ho, C.C.: Enhancing titanium spacer defect detection through reinforcement learning-optimized digital twin and synthetic data generation, *Journal of Electronic Imaging*, Vol. 33, No. 1, 2024, 013021.
- [10] Mirzaei, K., Arashpour, M., Asadi, E., Masoumi, H., Li, H.: Automatic generation of structural geometric digital twins from point clouds, *Scientific Reports*, Vol. 12, No. 1, 2022, 22321.
- [11] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal Policy Optimization Algorithms, *arXiv preprint*, arXiv:1707.06347, 2017.
- [12] Zhao, S., Jakob, W., Li, T.M.: Physics-based differentiable rendering: from theory to implementation, *Proceeding of ACM SIGGRAPH 2020 Courses*, New York, NY, USA, August 2020, pp. 1-30.