

From Swept Contact to Pose: Probe-Aware Registration via Complementary-Shape Docking

Chen Chen, Yunwen Li, Yifan Xu, Xiangjie Yan, Chang Shu, Jianxia Hou, Shiji Song, and Xiang Li

Abstract—Accurate registration between a prior model and the real scene is essential for high-precision robotic manipulation, yet optical methods suffer from long calibration chains, line-of-sight constraints, and fabrication errors. We propose a calibration-free alternative that reformulates contact registration as complementary-shape docking between the object and the probe’s swept volume, explicitly accounting for probe geometry and leveraging both contact and non-contact evidence. Our solver integrates a global-to-local search via 3D FFT correlation over low-discrepancy $SO(3)$ samples, then followed by continuous $SE(3)$ refinement using Lie-algebra updates and analytic contact sensitivities. This pipeline yields efficient exploration and metric-grade convergence without fragile point correspondences. Simulation across free-form meshes achieved sub-0.04 mm and sub-0.4° accuracy and robustness to pose noise and contact loss. On a tooth-preparation robot, our method attained 0.42 mm and 3.75°, outperforming an optical tracker registration while requiring no external sensors. These results demonstrate a practical and precise registration strategy for surgical and industrial robots.

I. INTRODUCTION

Accurate registration between a prior model and the real scene is a cornerstone of robotic manipulation. It anchors task geometry to the robot frame and directly bounds the achievable precision of execution. Applications such as tooth preparation, orthopedic surgery, and precision assembly demand sub-millimeter accuracy, where even small registration errors can compromise task success.

A widely adopted solution is optical registration, which integrates hand-eye calibration, fiducial detection, and computer-aided design (CAD)-marker manufacturing. However, long kinematic chains and the accumulation of fabrication and extrinsic errors often limit accuracy even with high-end trackers. Moreover, line-of-sight constraints restrict usability in confined or occluded workspaces. These limitations motivate alternatives that bypass external tracking hardware.

Contact-based registration is one such alternative [1]. A rigid probe establishes contact with the object surface, and the robot records the swept trajectory of end-effector poses. This trajectory is then aligned to the prior model to estimate the object pose. Yet most existing methods either treat the probe

C. Chen, Y. Xu, X. Yan, S. Song, and X. Li are with the Department of Automation, Tsinghua University. Y. Li is with Tsingscribe Medical Ltd., and also with D-MAVT, ETH Zürich. C. Shu and J. Hou are with Peking University School and Hospital of Stomatology. This work was supported in part by the Brain Science and Brain-like Intelligence Technology-National Science and Technology Major Project under Grant 2021ZD0201404, in part by the National Natural Science Foundation of China under Grant 62461160307, in part by the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China under Grant JYB2025XDXM208, and in part by the Institute for Guo Qiang, Tsinghua University. Corresponding author: Xiang Li (xiangli@tsinghua.edu.cn)

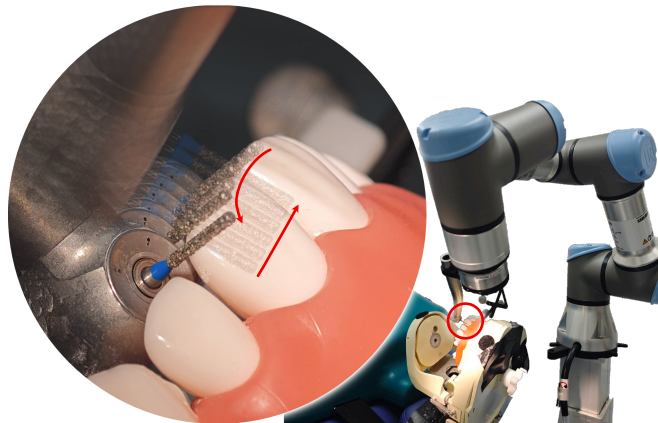


Fig. 1. Illustration of proposed registration method in a tooth preparation robot. The manipulator sweeps a cylindrical probe across the target tooth to capture its geometry, which is then registered to a preoperative 3D model for precise motion control.

as a point or approximate it with a spherical tip and radius compensation [2], [3], [4], [5]. These simplifications discard geometric constraints imposed by the true probe shape, rely on fragile point correspondences (e.g., iterative closest point (ICP) variants), and degrade when trajectories contain sparse or intermittent contact. **This leaves a gap:** a formulation that *explicitly* incorporates probe geometry, leverages both contact and non-contact evidence, and couples efficient global search with metric-grade refinement remains missing.

We address this gap by recasting contact registration as a *complementary-shape docking* problem. The admissible contact set is determined by the probe’s full swept volume rather than isolated points (Fig. 1), turning contact from sparse correspondences into dense geometric constraints. This perspective leverages all trajectory evidence, both contact and non-contact, for robust pose estimation.

Our contributions are:

- A reformulation of contact-based registration as complementary-shape docking that encodes probe geometry via swept volumes, avoiding fragile point correspondences.
- A two-stage global-to-local search over $SE(3)$ that couples FFT-based translation correlation with low-discrepancy orientation sampling, to produce robust initial hypotheses.
- A continuous refinement method on $SE(3)$ using Lie-algebraic updates and analytic contact sensitivities for precise and stable convergence.

The method requires no additional sensing hardware and

relies only on manipulator proprioception, making it practical in both medical and industrial settings. We validated our approach both in simulation and on a tooth-preparation robot using a cylindrical bur as the probe. In simulation, it achieved an average translational error of 0.031 mm and a rotational error of 0.325°. On the real robot, as shown in Fig. 1, it achieved 0.42 mm translational error and 3.75° rotational error, consistently outperforming optical pipelines while requiring no external sensors.

II. RELATED WORKS

A. Registration Strategies

Prior work has explored a range of registration strategies, from correspondence-based alignment to global inference.

Early coordinate measuring machine (CMM) pipelines register parts by probing canonical CAD features and solving least-squares with assumed correspondences [1], [6]. Grimson and Lozano-Pérez [7] established an early correspondence-driven paradigm by combining sparse points and surface normals with geometric consistency tests and pruning. Olsson *et al.* [8] showed that correspondence-based registration can be made globally optimal via branch-and-bound over points, lines, and planes, yet it still fundamentally requires reliable correspondences. Modern contact registration often follows ICP-style alignment from sparse touch points and estimated normals, which remains sensitive to outliers and initialization [2], [4].

To reduce dependence on initialization and explicit correspondences, a second line of work reasons globally from contact observations. Petrovskaya and Khatib [9] formulate tactile localization as Bayesian inference that treats touch as probabilistic measurements, thereby enabling global pose reasoning. Haugaard *et al.* [5] discretize $SE(3)$ and test candidate poses against a CAD model, offering correspondence-free global search at the cost of resolution–efficiency trade-offs. Zhong *et al.* [10] output diverse pose hypotheses consistent with observed contact and free-space, emphasizing coverage rather than a single metrology-grade estimate.

In summary, correspondence-driven pipelines require accurate matches or well-chosen features and are fragile under sparse or noisy contact, whereas correspondence-free global inference leverages contact and free-space but typically trades precision for tractability.

B. Probe Geometry and Compensation

How the probe geometry is modeled or compensated has been a central factor in contact-based registration.

A common industrial practice is CAD normal offsetting, where the measured contact point is shifted inward by the stylus radius before feature fitting [1], [6]. Ristic *et al.* [11] refine this idea with CAD normal compensation coupled with NURBS surface fitting, yet the spherical assumption remains and directional effects are ignored. Xiong *et al.* [2] represent the standard metrology view by adopting spherical tips with radius correction in correspondence-driven alignment, which is adequate for ball tips but does not generalize to anisotropic probes.

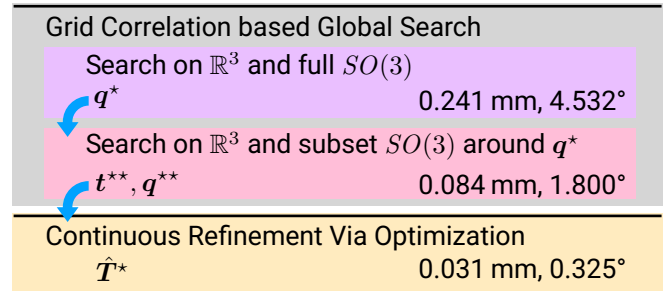


Fig. 2. An overview of the proposed registration pipeline, which consists of three stages: (a) a global search over $SE(3)$ with grid correlation and low-discrepancy orientation sampling, (b) a local search on subset $SO(3)$ around best orientation q^* to improve orientation accuracy, and (c) a continuous optimization on $SE(3)$ using Lie-algebraic updates and analytic contact sensitivities for precise convergence. Average error for each stage is shown for reference.

A second compensation strategy replaces explicit tip handling with a Minkowski sum between the CAD surface and a sphere of the probe radius, effectively inflating the object and reducing registration to point–surface alignment on the dilated model [5]. Saund *et al.* [3] implement an equivalent surrogate by subtracting the probe radius in the SDF test, which approximates a spherical Minkowski sum within a particle-filter framework.

A third line absorbs the probe radius into the measurement noise rather than modeling the geometry explicitly, which simplifies inference but caps achievable accuracy and obscures the distinction between near-contact and penetration [9]. Several recent contact-ICP systems also neglect radius or treat a sharp “pointy tool” as a point, further limiting precision under realistic tool shapes [4], [8]. Works that directly assume true surface samples or rely on deformable tactile skins bypass rigid-probe compensation but lie outside our rigid-probe metrology setting [7], [8], [10].

In summary, most pipelines simplify the probe to a sphere and then rely on offsets, Minkowski sums, or noise absorption. These heuristics ease implementation but blur geometric distinctions and limit achievable accuracy for non-spherical or anisotropic tools.

III. PRELIMINARIES

We observe a sequence of probe poses $\mathcal{P} = \{T_i\}_{i=1}^N$, with $T_i \in SE(3)$, describing a trajectory in which a rigid probe attempts to touch the surface of an object. Importantly, there is no guarantee that the probe touches the object surface along this trajectory, which makes direct alignment ambiguous.

The task is to estimate a rigid transformation $\hat{T} \in SE(3)$ that best aligns the object with the observed trajectory.

To this end, we first perform a coarse grid-based search with progressively finer orientation sampling. We then apply a continuous optimization to refine the pose estimate. The overall procedure is summarized in the three-stage diagram in Fig. 2.

IV. GRID CORRELATION BASED GLOBAL SEARCH

This section describes a two-stage search over $SE(3)$: a global exploration followed by a local search within

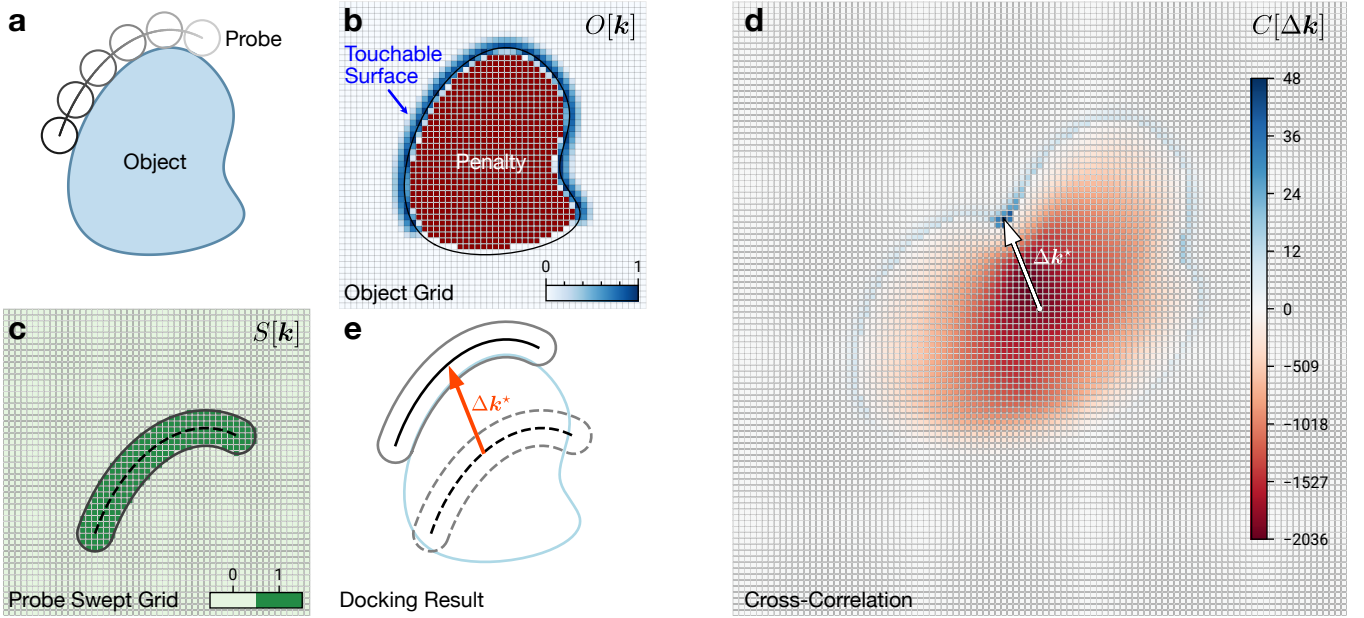


Fig. 3. Illustration of grid representations and translational search. (a) The probe sweeps along the trajectory while in contact, and the object is to be registered. (b) Object grid $O[\mathbf{k}]$ rewarding near-surface approach, tolerating shallow penetration, and penalizing deep penetration. (c) Probe swept grid $S[\mathbf{k}]$ indicating all covered voxels along probe’s trajectory. (d) Cross-correlation $C[\Delta\mathbf{k}]$ between the object $O[\mathbf{k}]$ and swept volume grids $S[\mathbf{k}]$, determining the best translation $\Delta\mathbf{k}^*$. (e) Resulting docking of the probe and object, showing complementary-shape registration respecting contact locality.

neighborhood. In the first stage, we evaluate grid-based correlation fields constructed from the object and the probe’s swept volume over uniformly sampled orientations to obtain a coarse pose hypothesis. The resulting coarse orientation is then refined by a second grid-based search constrained to a geodesic ball on $SO(3)$, yielding a more accurate candidate. Our approach is inspired by correlation-based molecular docking [12], [13], [14], which pioneered the use of grid-based correlation for efficient global search. In contrast, we adapt the scoring to contact-locality and penetration handling in physical contact and introduce a refinement stage within $SE(3)$ to achieve higher accuracy.

A. Object Grid and Probe Swept Volume Grid Representation

Both the object and the swept volume are discretized into voxel arrays at the same spacing h . Each array is defined on its own lattice, determined by a limited padding around the axis-aligned bounding box (AABB) of the corresponding geometry. A voxel index $\mathbf{k} \in \mathcal{K} \subset \mathbb{Z}^3$ maps to world coordinates $\mathbf{x} = h\mathbf{k} + \mathbf{o} \in \mathbb{R}^3$ with lattice origin \mathbf{o} .

1) *Object Grid*: Two meshes of the object serve complementary roles: a watertight mesh provides inside/outside information, and a non-watertight subset constrains feasible contacts to touchable regions.

From the watertight mesh we compute a signed distance field (SDF) $f_o(\mathbf{x})$, which is mapped into a banded template

$$O_t(\mathbf{x}) = \begin{cases} 1 - f_o(\mathbf{x})/\tau_{\text{out},1} & \text{if } 0 \leq f_o(\mathbf{x}) < \tau_{\text{out},1}, \\ 1 + f_o(\mathbf{x})/\tau_{\text{in},1} & \text{if } -\tau_{\text{in},1} < f_o(\mathbf{x}) < 0, \\ -\rho & \text{if } f_o(\mathbf{x}) \leq -\tau_{\text{in},1}, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where $\tau_{\text{in},1}, \tau_{\text{out},1}$ are the inner and outer thresholds, and ρ is a penalty term for deep penetration. This rule rewards near-surface approach, tolerates shallow penetration, and penalizes deep penetration, as shown in Fig. 3(b).

Most objects have limited surface area that can be contacted by the probe. For instance, an object placed on a table has only the upper surface reachable. Such area is represented by a subset of the object mesh, which we call the touchable surface. From the touchable surface we compute an unsigned distance $d(\mathbf{x})$, which measures the distance from the touchable surface. We define a surface mask $O_m(\mathbf{x})$ as a linearly feathered band around the zero level set as

$$O_m(\mathbf{x}) = \max\left(0, 1 - \frac{|d(\mathbf{x})|}{\tau_{\text{out},1}}\right), \quad (2)$$

which gates scoring to feasible contact regions.

The final object grid is an element-wise combination

$$O(\mathbf{x}) = \min(O_t(\mathbf{x}), O_m(\mathbf{x})), \quad (3)$$

which simultaneously encodes interior penalties and contact locality. Sampling on the object lattice yields

$$O[\mathbf{k}] = O(h\mathbf{k} + \mathbf{o}_o), \quad (4)$$

shown in Fig. 3(b), where \mathbf{o}_o is the lattice origin.

2) *Probe Swept Occupancy Grid*: We construct a binary occupancy field directly from the probe SDF and the trajectory. Let $f_p(\mathbf{x})$ denote the probe SDF in its local frame, and \mathbf{T}_i be sampled probe poses along trajectory \mathcal{P} . Define

$$S(\mathbf{x}) = \max_{i=1, \dots, N} H(f_p(\mathbf{T}_i^{-1}\mathbf{x})), \quad (5)$$

where $H: \mathbb{R} \mapsto \{0, 1\}$ is the Heaviside step function. Thus $S(\mathbf{x}) = 1$ if the swept volume covers \mathbf{x} and 0 otherwise.

Discretization on the swept-volume lattice gives

$$S[\mathbf{k}] = S(h\mathbf{k} + \mathbf{o}_s), \quad (6)$$

shown in Fig. 3(c), where \mathbf{o}_s is the lattice origin.

B. Translation Search on \mathbb{R}^3 via Correlation

We embed $S[\mathbf{k}]$ and $O[\mathbf{k}]$ on a common zero-padded canvas at spacing h without resampling. Let the embedded arrays be denoted by the same symbols for brevity. The value is zero if query is outside their supports.

Matching the surfaces is achieved by evaluating the cross-correlation between the two arrays:

$$C[\Delta\mathbf{k}] = \sum_{\mathbf{k} \in \mathcal{I}_o} O[\mathbf{k}]S[\mathbf{k} + \Delta\mathbf{k}], \quad (7)$$

where \mathcal{I}_o is the index set of the object lattice. We take the maximizer $\Delta\mathbf{k}^*$ as the best voxel shift on the canvas, as shown in Fig. 3(d). The corresponding physical translation between the original geometries is

$$\mathbf{t} = h\Delta\mathbf{k}^* + \mathbf{o}_s - \mathbf{o}_o. \quad (8)$$

Direct evaluation of (7) is computationally heavy, $\mathcal{O}(N^6)$. We accelerate it to $\mathcal{O}(N^3 \log N)$ with 3D fast Fourier transforms (FFTs) using

$$C = \mathcal{F}^{-1}(\overline{\mathcal{F}(S)} \odot \mathcal{F}(O)), \quad (9)$$

where \odot is element-wise multiplication and the overline denotes complex conjugation. To prevent wrap-around, each axis of the canvas is zero-padded to at least $n_o + n_s - 1$, where n_o and n_s are the side lengths of the two finite supports along that axis.

C. Orientation Sampling on $SO(3)$ and its Subspace

Exploring translation alone requires a fixed probe orientation, but full 6-DoF alignment demands systematic sampling of orientations on $SO(3)$. Rather than relying on dense grids or random sampling, we adopt a two-stage strategy: a global low-discrepancy sampler for coarse exploration and a novel geodesic ball neighborhood sampler for local refinement.

1) *Global Orientation Sampling*: For coarse global sampling, we adopt the Super-Fibonacci (SF) sequence [15], which generates unit quaternions with low discrepancy directly on the 3-sphere \mathbb{S}^3 . Compared to Hopf fibration based method [16], as used in [5], the SF method achieves lower discrepancy and more uniform coverage while remaining equally efficient.

2) *Geodesic Neighborhood Sampling*: After obtaining the best coarse orientation, we refine the search in its local neighborhood. We consider the unit quaternion space as the 3-sphere \mathbb{S}^3 , where the natural distance between two quaternions $\mathbf{q}_1, \mathbf{q}_2$ is the half relative rotation angle

$$d(\mathbf{q}_1, \mathbf{q}_2) = \arccos(|\langle \mathbf{q}_1, \mathbf{q}_2 \rangle|). \quad (10)$$

A geodesic ball of radius θ centered at quaternion \mathbf{q}_0 is thus defined as $\mathcal{B}(\mathbf{q}_0, \theta) = \{\mathbf{q} \in \mathbb{S}^3 : d(\mathbf{q}, \mathbf{q}_0) \leq \theta\}$, which in rotation space contains all orientations within an angular deviation of 2θ in the neighbor of \mathbf{q}_0 .

Uniform sampling in such a restricted region is nontrivial. The standard SF sequence construction is uniform on the full \mathbb{S}^3 , but becomes biased when restricted to a subspace. To preserve uniformity, we replace the radial law with inverse transform sampling with respect to the true radial marginal distribution. The marginal density of the radius is

$$\rho(r) \propto rL(r), \quad L(r) = 2 \arccos\left(\frac{\cos \theta}{\sqrt{1-r^2}}\right), \quad (11)$$

where $L(r)$ denotes the admissible range of the axial angle at radius r . The associated cumulative distribution is

$$F(r) = \frac{1}{Z} \int_0^r \rho(\tilde{r}) d\tilde{r}, \quad (12)$$

with Z the normalization constant. A uniform variable $u \in [0, 1]$ is then mapped to r by the inverse transform $r = F^{-1}(u)$. Since F has no closed form, we evaluate it numerically on a dense grid, store the tabulated pairs $(r, F(r))$, and recover $F^{-1}(u)$ via interpolation. The resulting GeoBall-SF sequence, summarized in algorithm 1, produces low-discrepancy samples uniformly distributed within the geodesic ball.

Algorithm 1: Generating n samples in a geodesic ball around \mathbf{q}_0 on $SO(3)$ (GeoBall-SF)

Input: $n, \mathbf{q}_0, \theta, \phi, \psi, \gamma$;

Output: Unit quaternions \mathbf{q}_i in $\mathcal{B}(\mathbf{q}_0, \theta)$

for $i \in \{0, \dots, n-1\}$ **do**

$s \leftarrow i + 0.5$;

$u_1 \leftarrow \frac{s}{\gamma} \bmod 1, u_2 \leftarrow \frac{s}{\phi} \bmod 1, u_3 \leftarrow \frac{s}{\psi} \bmod 1$;

$r \leftarrow F^{-1}(u_1), \quad R \leftarrow \sqrt{1-r^2}$;

$L \leftarrow 2 \arccos(\cos \theta / R)$;

$\alpha \leftarrow (u_2 - 0.5)L, \quad \beta \leftarrow 2\pi u_3$;

$\mathbf{q}_{\text{local}} \leftarrow (R \cos \alpha, R \sin \alpha, r \cos \beta, r \sin \beta)$;

$\mathbf{q}_i \leftarrow \mathbf{q}_0 \mathbf{q}_{\text{local}}$;

D. Combined Grid-Based Search

We combine the two components introduced above: \mathbb{R}^3 translation search by FFT correlation and orientation sampling on $SO(3)$.

We sample n_g quaternions from the global SF sequence. For each quaternion, the object grid O is rotated into the corresponding orientation and resampled onto the common canvas at spacing h via trilinear interpolation. We then compute the correlation C with the swept-volume occupancy grid S using FFT, recovering the best translation \mathbf{t} for that orientation. Maximizing the score over all orientations and translations yields a coarse pose hypothesis $(\mathbf{t}^*, \mathbf{q}^*)$.

To refine this result, we restrict the search to a geodesic ball on $SO(3)$ centered at \mathbf{q}^* . Candidate quaternions are generated with the GeoBall-SF sampler, ensuring uniform and low-discrepancy coverage within the neighborhood, with sample number n_l and radius θ . For each candidate, the object grid is again rotated, resampled, and correlated with the swept volume, updating the best hypothesis if a higher score is obtained.

In summary, the method performs a global search over $SE(3)$ followed by a local refinement in the orientation subspace. The result is an initial high-quality pose hypothesis ($\mathbf{t}^{**}, \mathbf{q}^{**}$), which serves as the starting point for the optimization-based refinement in the next section.

V. CONTINUOUS REFINEMENT VIA OPTIMIZATION

Although the two-stage global search provides a strong hypothesis for the contact registration, it remains limited by the use of discrete orientation sampling and a grid-based field representation, and thus may not fully align the contact geometries.

To address this, we introduce a continuous gradient-based refinement on $SE(3)$, optimizing a proximity objective around the best hypothesis ($\mathbf{t}^{**}, \mathbf{q}^{**}$) through Lie-algebra updates.

A. Objective and Pose Update

The object pose is represented as $\mathbf{T} = (\mathbf{R}, \mathbf{p}) \in SE(3)$ with rotation $\mathbf{R} \in SO(3)$ and translation $\mathbf{p} \in \mathbb{R}^3$. We initialize at $\mathbf{T}_{\text{init}} = (\mathbf{R}_{\text{init}}, \mathbf{t}_{\text{init}})$, which corresponds to the best hypothesis ($\mathbf{t}^{**}, \mathbf{q}^{**}$). The decision variable is a twist

$$\boldsymbol{\xi} = [\mathbf{v} \quad \boldsymbol{\omega}]^T \in \mathbb{R}^6, \quad (13)$$

where \mathbf{v} (mm) and $\boldsymbol{\omega}$ (rad) are the translational and rotational components. We map $\boldsymbol{\xi}$ to an increment $\Delta\mathbf{T} = \exp(\boldsymbol{\xi}^\wedge) = (\Delta\mathbf{R}, \Delta\mathbf{p})$ with

$$\Delta\mathbf{R} = \exp(\boldsymbol{\omega}^\wedge), \quad \Delta\mathbf{p} = \mathbf{V}(\boldsymbol{\omega})\mathbf{v}, \quad (14)$$

where $\boldsymbol{\xi}^\wedge \in \mathfrak{se}(3)$, $\boldsymbol{\omega}^\wedge \in \mathfrak{so}(3)$ denote the skew-symmetric matrix representations of the twist $\boldsymbol{\xi}$ and the rotation vector $\boldsymbol{\omega}$. The $\mathbf{V}(\boldsymbol{\omega})$ is the $SO(3)$ left Jacobian, defined as

$$\mathbf{V}(\boldsymbol{\omega}) = \mathbf{I} + \frac{1 - \cos \theta}{\theta^2} \boldsymbol{\omega}^\wedge + \frac{\theta - \sin \theta}{\theta^3} (\boldsymbol{\omega}^\wedge)^2, \quad (15)$$

where $\theta = \|\boldsymbol{\omega}\|_2$ is the rotation angle. For $\theta \rightarrow 0$, the series expansion of $\mathbf{V}(\boldsymbol{\omega})$ is used for numerical stability.

To better decouple the translational and rotational updates, an anchor point is defined in the object body frame as $\mathbf{a} \in \mathbb{R}^3$, which represents the geometry center of the mesh. We perform a left perturbation about the anchor expressed in the world frame as $\mathbf{T}(\boldsymbol{\xi}) = (\mathbf{R}(\boldsymbol{\xi}), \mathbf{p}(\boldsymbol{\xi}))$, where

$$\mathbf{R}(\boldsymbol{\xi}) = \Delta\mathbf{R}\mathbf{R}_0, \quad (16)$$

$$\mathbf{p}(\boldsymbol{\xi}) = \mathbf{p}_0 + \Delta\mathbf{p} + (\mathbf{I} - \Delta\mathbf{R})\mathbf{R}_0\mathbf{a}. \quad (17)$$

The compensation term $(\mathbf{I} - \Delta\mathbf{R})\mathbf{R}_0\mathbf{a}$ ensures that the rotation is performed about the anchor rather than the world origin.

For each probe pose \mathbf{T}_i , we compute the smallest signed distance d_i between the object and the probe geometry at that pose (see subsection V-B). We define a smooth proximity scoring function $s(d)$ that rewards near-surface proximity while softly penalizing penetration:

$$t \triangleq \frac{d}{\tau_{\text{out},2}}, \quad s(d) = \begin{cases} 0, & t > 1, \\ \cos^2(0.5\pi t), & 0 < t \leq 1, \\ At^3 + Bt^2 + 1, & -r < t \leq 0, \\ k(t+r), & t \leq -r, \end{cases} \quad (18)$$




Fig. 4. The scoring function $s(d)$ used in the continuous optimization stage.

where $\tau_{\text{out},2}, \tau_{\text{in},2}$ denote the outer and inner thresholds, $r \triangleq \tau_{\text{out},2}/\tau_{\text{in},2} = 0.15$ is their ratio, $k = 15$ is the penalty slope, and the coefficients $A = k/r^2 + 2/r^3$, $B = -k/r - 3/r^2$ ensure C^1 continuity. The resulting function encourages contact points to lie within the template bands while penalizing large deviations (Fig. 4).

The registration maximizes total proximity; equivalently, we minimize

$$J(\boldsymbol{\xi}) = - \sum_{i=1}^N s(d_i(\boldsymbol{\xi})) + \lambda(\|\mathbf{v}\|_2^2 + \|\boldsymbol{\omega}\|_2^2), \quad (19)$$

where $\lambda \geq 0$ is a Tikhonov regularizer.

Final Optimization Problem: Placing bounds per component yields the constrained problem

$$\begin{aligned} \min_{\boldsymbol{\xi} \in \mathbb{R}^6} \quad & J(\boldsymbol{\xi}) \\ \text{s.t.} \quad & \|\mathbf{v}\|_\infty \leq b_t, \\ & \|\boldsymbol{\omega}\|_\infty \leq b_r, \end{aligned} \quad (20)$$

where $J(\boldsymbol{\xi})$ is defined in (19), distances $d_i(\boldsymbol{\xi})$ are evaluated on the pose $\mathbf{T}(\boldsymbol{\xi})$ via (14)–(17), and b_t, b_r are tight physical bounds. We denote the minimizer by $\boldsymbol{\xi}^*$; it defines the refined pose $\mathbf{T}(\boldsymbol{\xi}^*)$ through (16)–(17).

B. Distance Queries on Single Probe Geometry to Mesh

To evaluate each probe pose in the objective, we require robust queries of the shortest distance between the probe and the target mesh.

The probe shape is represented by a SDF $f_p(\mathbf{x})$. The mesh is stored in a bounding volume hierarchy (BVH). For each BVH node, we evaluate f_p at the node center \mathbf{c} and keep the node if

$$f_p(\mathbf{c}) < r_{\text{node}} + \varepsilon, \quad (21)$$

where r_{node} is the circumsphere radius of the node and ε is a cutoff margin. This check ensures that triangles with clearance less than ε are preserved. Candidate triangles are further pruned using the same test with their circumsphere.

For each remaining triangle, a Frank–Wolfe optimization returns a candidate point \mathbf{x}_k and value $d_k = f_p(\mathbf{x}_k)$. The signed distance is the smallest of these,

$$d^* = \min_k d_k, \quad (22)$$

where $\mathbf{x}^* = \mathbf{x}_{k^*}$ and $k^* = \arg \min_k d_k$ is the index of the minimizing triangle. Negative values indicate penetration,

while $0 \leq d^* < \varepsilon$ indicates positive clearance. The output is a distance vector

$$\mathbf{v} = -d^* \mathbf{n}^*, \quad \mathbf{n}^* = \frac{\nabla f_p(\mathbf{x}^*)}{\|\nabla f_p(\mathbf{x}^*)\|}, \quad (23)$$

where \mathbf{n}^* is the outward SDF normal at the closest point. This vector points from the witness point to the nearest point on the SDF surface.

The entire distance query is implemented as a custom SDF–mesh collision model within MuJoCo [17].

C. Gradient and Solver

Let $\mathbf{v}_i, \mathbf{n}_i$ be the output from the SDF distance query of the i -th probe pose. $\mathbf{v}_i = \mathbf{p}_{\text{obj}}^i - \mathbf{p}_{\text{probe}}^i$, \mathbf{n}_i is the oriented unit normal consistent with the sign of d_i . Differentials of the signed distance under rigid motion follow standard contact sensitivities:

$$\frac{\partial d_i}{\partial(\Delta \mathbf{p})} = \mathbf{n}_i, \quad (24)$$

$$\left. \frac{\partial d_i}{\partial \boldsymbol{\omega}} \right|_{\text{rot}} = (\mathbf{p}_{\text{obj}}^i - \mathbf{c}) \times \mathbf{n}_i, \quad (25)$$

with pivot $\mathbf{c} = \mathbf{p}_0 + \Delta \mathbf{p} + \mathbf{R}_0 \mathbf{a}$ from (17).

Since $\Delta \mathbf{p} = \mathbf{V}(\boldsymbol{\omega}) \mathbf{v}$ by (14), the chain rule to algebra coordinates yields

$$\frac{\partial d_i}{\partial \mathbf{v}} = \mathbf{V}(\boldsymbol{\omega})^\top \mathbf{n}_i, \quad (26)$$

$$\frac{\partial d_i}{\partial \boldsymbol{\omega}} = (\mathbf{p}_{\text{obj}}^i - \mathbf{c}) \times \mathbf{n}_i + \left(\frac{\partial(\mathbf{V}(\boldsymbol{\omega}) \mathbf{v})}{\partial \boldsymbol{\omega}} \right)^\top \mathbf{n}_i. \quad (27)$$

The second term in (27) captures translation–rotation coupling and is approximated numerically by finite differences.

Let ds/dd denote the derivative of the proximity window. The full gradient of (19) is

$$\nabla J(\boldsymbol{\xi}) = - \sum_{i=1}^n \left. \frac{ds}{dd} \right|_{d_i(\boldsymbol{\xi})} \begin{bmatrix} \partial d_i / \partial \mathbf{v} \\ \partial d_i / \partial \boldsymbol{\omega} \end{bmatrix} + 2\lambda \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\omega} \end{bmatrix}. \quad (28)$$

We solve (19) with L-BFGS-B optimizer [18] on $\boldsymbol{\xi}$, using initial value $\boldsymbol{\xi}_0 = \mathbf{0}$, bounds $\|\mathbf{v}\|_\infty \leq b_t$, $\|\boldsymbol{\omega}\|_\infty \leq b_r$. The analytic gradient (28) is supplied to the optimizer. We terminate when the projected gradient norm and the step size fall below prescribed thresholds.

VI. RESULTS

We evaluate the proposed method in both simulation and real-world experiments. The parameter settings used in all experiments are summarized in Table I. The average nearest neighbor distance of the global and local orientation samples are 8.0° and 0.7° , respectively.

A. Simulation

We generated contact trajectories in simulation to evaluate the proposed method, using a cylindrical probe with a diameter of 1.4 mm and a length of 20 mm.

Five object meshes were used in the simulation, as shown in Fig. 5: a workpiece (simple geometric shape suitable for traditional CMM registration), the Matlab logo (mathematical free-form surface), a toy car (CAD-derived object with

TABLE I
PARAMETER SETTINGS USED IN ALL EXPERIMENTS.

Parameter	Value	Parameter	Value
ρ	-50	h	0.2 mm
n_g	20000	$\tau_{\text{out},1}$	0.5 mm
n_l	5000	$\tau_{\text{in},1}$	0.2 mm
θ	5°	$\tau_{\text{out},2}$	0.2 mm
λ	10×10^{-4}	$\tau_{\text{in},2}$	0.03 mm

complex geometry), and front and molar teeth (meshes scanned by an intraoral scanner). All meshes were scaled so that their longest side measured approximately 5 mm to 8 mm.

For each object, we generated a contact trajectory via OpenCAMLlib [19], simulating the probe contacting the object surface. Each trajectory consisted of three segments with different contact directions.

To better mimic real-world execution, we perturbed each trajectory by adding Gaussian noise with a standard deviation of 0.02 mm independently to each pose sample. The trajectory was sampled to be as realistic as possible for a robot to execute using a force–position hybrid controller. The trajectories are visualized in Fig. 5.

The calibration results were evaluated by the translational error of the mesh geometry center and the rotation error. Table II reports the results of three stages. The global search stage achieved an average translation error of 0.241 mm and rotation error of 4.532° . The local search stage achieved an average translation error of 0.084 mm and rotation error of 1.800° . The refinement stage further improved the accuracy to 0.031 mm and 0.325° , demonstrating a significant benefit. The average runtime was 115 s.

To evaluate robustness, we added independent Gaussian noise to each pose in the trajectory. The noise variance was varied across experiments, and the results are shown in Fig. 6. The proposed method remained robust up to 0.1 mm noise: below this threshold, both translational and rotational errors remained small and the simulation results were stable. When the noise standard deviation exceeded 0.1 mm, accuracy dropped sharply and the simulation results became unstable. This failure occurs when a large portion of the swept volume intrude into the mesh, causing heavy penetration penalties and negative scores during local refinement.

In addition to the Gaussian noise experiments, we also evaluated the method under sparse contact conditions, where only a portion of the trajectory points were in actual contact with the mesh. This behavior reflects realistic scenarios, where the object may “hop” or intermittently touch the surface rather than maintaining continuous contact. Remarkably, the registration remained accurate even with up to 98% of the trajectory not touching the surface, as shown in Fig. 7. The results indicate that the method can achieve reliable registration with sparse data, demonstrating robustness in practical scenarios.

B. Real-World Experiments

Since acquiring ground truth for real-world registration is difficult, the accuracy is indirectly assessed through the




Fig. 5. Simulation overview. Contact trajectories were generated with a cylindrical probe on five object meshes. Translucent swept volumes show the simulated probe motion used for registration. Gaussian noise was added during experiments but omitted here for clarity.

TABLE II
SIMULATION RESULTS

Mesh Name	Trajectory Length	Global Search Results		Local Search Results		Refinement Stage Results	
		Trans. Err(mm)	Rot. Err(°)	Trans. Err(mm)	Rot. Err(°)	Trans. Err(mm)	Rot. Err(°)
Workpiece	395	0.089 ± 0.007	3.117 ± 0.000	0.102 ± 0.007	2.750 ± 0.000	0.021 ± 0.003	0.214 ± 0.042
Matlab Logo	525	0.368 ± 0.011	3.117 ± 0.000	0.073 ± 0.039	2.163 ± 0.513	0.040 ± 0.004	0.512 ± 0.112
Toy Car	542	0.177 ± 0.033	3.117 ± 0.000	0.068 ± 0.006	1.678 ± 0.979	0.034 ± 0.003	0.322 ± 0.099
Front Teeth	242	0.107 ± 0.075	5.290 ± 0.000	0.101 ± 0.031	1.110 ± 0.254	0.030 ± 0.003	0.403 ± 0.235
Molar Teeth	339	0.463 ± 0.009	8.018 ± 0.000	0.078 ± 0.012	1.299 ± 0.511	0.032 ± 0.001	0.174 ± 0.041
Mean Across Meshes		0.241 ± 0.166	4.532 ± 2.164	0.084 ± 0.016	1.800 ± 0.666	0.031 ± 0.007	0.325 ± 0.138




Fig. 6. Noise robustness evaluation. The plot shows the registration error under different levels of Gaussian noise added to the poses.




Fig. 7. Sparse contact evaluation. The registration remained accurate with up to 98% of the trajectory not touching the surface.

outcomes of tooth preparation.

We evaluated the proposed method on a tooth preparation robot system, consisting of a UR3e manipulator equipped with a dental handpiece with a cylindrical bur. Its task is to perform tooth preparation surgery on a phantom teeth, shaping the target tooth according to a predesigned model. Such surgery requires an accurate model registration between the preoperative 3D model and the real tooth.

The accurate tooth mesh was obtained by a high-accuracy intraoral scanner (3Shape TRIOS). The cylindrical bur acted

as a probe, with its diameter of 1 mm and length precisely defined. The tool center point (TCP) was calibrated using the four-point method provided by the UR controller.

We used a force-position hybrid controller to maintain a contact force of 0.3 N during sweep contact. The contact path was planned similar to the simulated trajectories, with three segments of different contact directions. The trajectory execution time was 80 s. After the registration, the tooth preparation was executed using the obtained results.

For comparison, we also evaluated an optical-tracker-based calibration chain. Specifically, we used an NDI Polaris Lyra optical tracking system, which is widely adopted for high-precision pose measurement, with a reported error of 0.40 mm at the 95% confidence level and a volumetric RMS accuracy of 0.20 mm [20]. In our implementation, the NDI setup used a standard hand-eye calibration ($AX = XB$) to determine the robot-NDI transformation. The tooth-NDI relationship is obtained using a custom guide plate with attached markers. The guide is CAD-designed based on the intraoral scan and fabricated by 25 μ m-accuracy 3D printing, but it still introduced noticeable printing and mounting errors. Errors in this chain led to trajectory shifts: during experiments, the robot damaged both the neighboring teeth and the target tooth, and excessive cutting even broke the bur (see Fig. 8).

In contrast, our method requires no such calibration chain, and the preparation closely matched the planned tooth shape without damaging adjacent teeth or overloading the bur.

To evaluate its accuracy, we used an indirect method: the unprepared teeth served as stable references, since their geometry remains unchanged during preparation. Using these neighbors, we aligned the designed tooth model (from the initial intraoral scan) with the postoperative oral scan using ICP, treating this as an approximate ground truth between the

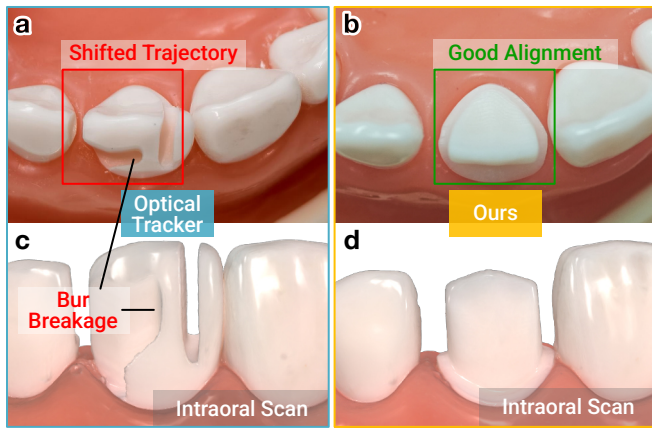


Fig. 8. Comparison of the prepared tooth using the optical tracker and the proposed registration method. (a) With the optical tracker, the trajectory was shifted leftward, damaging adjacent teeth, and the bur broke mid-procedure when forced to remove excess material caused by the misalignment. (b) Using our method, the prepared tooth closely matched the designed shape, demonstrating substantially improved registration accuracy.




Fig. 9. Comparison of designed tooth and prepared tooth using the proposed registration method. (a) Mesh visualization. (b) Per-point distance error between designed mesh and prepared tooth mesh.

two oral scans. We also performed ICP using the prepared tooth itself as the reference, assuming the prepared shape ideally matched the design except for a rigid offset.

By comparing the transformation from the approximate ground truth with the transformation estimated by our registration method, we obtained a pose error of 3.75° in rotation and 0.42 mm in translation. Furthermore, we computed the per-point distance error between the designed mesh and the prepared tooth mesh (Fig. 9). The distance error ranged from -0.59 mm (over-cut) to 0.71 mm (under-cut).

These results confirm that our registration method produces more reliable and accurate preparations than the optical-tracker-based approach, achieving better adherence to the planned tooth geometry.

VII. CONCLUSIONS

We proposed a calibration-free, probe-aware registration method that reformulates alignment as complementary-shape docking between the object and the probe's swept volume. The method achieved average 0.03 mm and 0.32° accuracy in simulation and 0.42 mm and 3.75° in real world experiments, outperforming an optical calibration chain while requiring only manipulator proprioception. In practice, the method delivers high precision on free-form surfaces while requiring neither external sensors nor careful initialization, making it a practical registration solution for surgical and industrial robots operating in constrained or occluded workspaces. A limitation is the reliance on an accurate object model, and

future work will explore confidence evaluation and the use of contact force cues to further strengthen robustness.

REFERENCES

- [1] K. T. Gunnarsson and F. B. Prinz, "CAD model-based localization of parts in manufacturing," *Computer*, vol. 20, no. 8, pp. 66–74, Aug. 1987.
- [2] Z. Xiong, M. Y. Wang, and Z. Li, "A near-optimal probing strategy for workpiece localization," *IEEE Trans. Robot.*, vol. 20, no. 4, pp. 668–676, Aug. 2004.
- [3] B. Saund, S. Chen, and R. Simmons, "Touch based localization of parts for high precision manufacturing," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. Singapore, Singapore: IEEE, May 2017, pp. 378–385.
- [4] Y. Kim, A. Kramberger, A. G. Buch, and C. Sloth, "Contact-based pose estimation of workpieces for robotic setups," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. London, United Kingdom: IEEE, May 2023, pp. 12 324–12 330.
- [5] R. L. Haugaard, Y. Kim, and T. M. Iversen, "Fixture calibration with guaranteed bounds from a few correspondence-free surface points," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. Yokohama, Japan: IEEE, May 2024, pp. 6718–6724.
- [6] Z. Li, J. Gou, and Y. Chu, "Geometric algorithms for workpiece localization," *IEEE Trans. Robot. Autom.*, vol. 14, no. 6, pp. 864–878, Dec. 1998.
- [7] W. E. L. Grimson and T. Lozano-Pérez, "Model-based recognition and localization from sparse range or tactile data," *Int. J. Robot. Res.*, vol. 3, no. 3, pp. 3–35, Sep. 1984.
- [8] C. Olsson, F. Kahl, and M. Oskarsson, "The registration problem revisited: Optimal solutions from points, lines and planes," in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, vol. 1. New York, NY, USA: IEEE, 2006, pp. 1206–1213.
- [9] A. Petrovskaya and O. Khatib, "Global localization of objects via touch," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 569–585, Jun. 2011.
- [10] S. Zhong, D. Berenson, and N. Fazeli, "CHSEL: Producing diverse plausible pose estimates from contact and free space data," in *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, Jul. 2023.
- [11] M. Ristic, I. Ainsworth, and D. Brujic, "Contact probe radius compensation using computer aided design models," *Proc. Inst. Mech. Eng. B: J. Eng. Manuf.*, vol. 215, no. 6, pp. 819–834, Jun. 2001.
- [12] E. Katchalski-Katzir, I. Shariv, M. Eisenstein, A. A. Friesem, C. Aflalo, and I. A. Vakser, "Molecular surface recognition: Determination of geometric fit between proteins and their ligands by correlation techniques," *Proc. Natl. Acad. Sci.*, vol. 89, no. 6, pp. 2195–2199, Mar. 1992.
- [13] J. M. Paggi, A. Pandit, and R. O. Dror, "The art and science of molecular docking," *Annu. Rev. Biochem.*, vol. 93, no. 1, pp. 389–410, Aug. 2024.
- [14] N. S. Pagadala, K. Syed, and J. Tuszynski, "Software for molecular docking: A review," *Biophys. Rev.*, vol. 9, no. 2, pp. 91–102, Apr. 2017.
- [15] M. Alexa, "Super-fibonacci spirals: Fast, low-discrepancy sampling of $SO(3)$," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. New Orleans, LA, USA: IEEE, Jun. 2022, pp. 8281–8290.
- [16] A. Yershova, S. Jain, S. M. LaValle, and J. C. Mitchell, "Generating uniform incremental grids on $SO(3)$ using the Hopf fibration," *Int. J. Robot. Res.*, vol. 29, no. 7, pp. 801–812, Jun. 2010.
- [17] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Vilamoura-Algarve, Portugal: IEEE, Oct. 2012, pp. 5026–5033.
- [18] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM J. Sci. Comput.*, vol. 16, no. 5, pp. 1190–1208, Sep. 1995.
- [19] A. Wallin, "OpenCAMLlib." [Online]. Available: <https://github.com/aewallin/opencamlib>
- [20] Northern Digital Inc, "Polaris Lyra," Mar. 2023. [Online]. Available: <https://www.ndigital.com/optical-navigation-technology/polaris-lyra/>