A Comparative Study of Crossover Operators on SE(3) for Manipulator Base Placement in Telesurgery

Alec Cotton, and Carlos Rossa
Department of Systems and Computer Engineering
Carleton University, Ottawa, ON, Canada
aleccotton@cmail.carleton.ca, carlosrossa@cunet.carleton.ca

***Abstract*—In teleoperated surgery, the operator's sense of remote presence is limited by the kinematics of the patient-side manipulator. A poorly positioned robot base with inadequate task space coverage forces the operator to reposition the controller or abandon commanded motions, and a base frame position close to kinematic singularities degrades the mapping between master inputs and instrument motion.**

Correctly placing the robot's base frame relative to a remote center of motion (RCM) is an optimization over the special Euclidean group $\mathrm{SE}(3)$. Yet most evolutionary approaches represent poses as Euclidean variables, neglecting the geometry of the underlying manifold. This paper compares four crossover operators within a non-dominated sorting genetic algorithm for base frame placement, formulated as a two-objective problem that maximizes task-space reachability and log-normalized manipulability under an RCM. These operators are: single-point crossover on Euler angles, simulated binary crossover (SBX) on Euler angles, screw linear interpolation on $\mathrm{SE}(3)$ geodesics, and Lie algebra blended crossover (BLX) in $\mathfrak{se}(3)$. They are evaluated across 4 population sizes and 10 dexterity thresholds on hypervolume, generational distance, inverted generational distance, and convergence generation. For small population sizes (≤ 10), both geometry-aware operators outperform both Euler-angle operators on hypervolume, generational distance, and inverted generational distance, and Lie algebra BLX attains the highest hypervolume, significantly above SBX after correction for multiple comparisons. For larger populations (≥ 50) the operators' hypervolumes lie within 0.3% of each other. Since the base frame placement must be optimized for each patient and port placement, the operationally relevant regime is the small-population one. In this case, Lie algebra BLX recovers 98.1% of the best hypervolume observed across the study while requiring $3.2\times$ fewer inverse-kinematic solutions.

***Index Terms*—Telesurgery, teleoperation, optimization, $\mathrm{SE}(3)$, genetic algorithm, NSGA-II, crossover operators, base frame placement, manipulability, reachability, Lie algebra.**

I. Introduction

In telesurgery, the surgeon commands instrument motion from a console while a patient-side robot manipulator pivots a tool about a fixed incision point, i.e., a remote center of motion (RCM). The operator's sense of telepresence is limited by the kinematics of the follower arm: the fraction of the commanded task space admitting a valid inverse kinematic solution determines how often the operator must interrupt motion, reposition the master controller, or abandon a task, resulting in a critical loss of immersion transparency. In addition, the conditioning of the arm's Jacobian determines how faithfully the tool can reproduce commanded Cartesian velocities. Workspace coverage and tool dexterity both depend on the robot's base placement, which must be re-established for every patient, port location, and approach angle.

Placing the robot's base can be seen as an optimization over rigid-body transformations that iterates variables of the special Euclidean group $\mathrm{SE}(3)$: a smooth, nonlinear, nonconvex manifold [1]–[5]. Unlike a vector space, $\mathrm{SE}(3)$ cannot be parametrized globally without singularities, and linear interpolation between two rigid-body transformations does not in general yield a valid pose; a constraint often overlooked when evolutionary algorithms are used as the optimizer.

Genetic algorithms suit optimization problems on $\mathrm{SE}(3)$ because they require neither convexity nor differentiability [6]. They combine high-performing parent chromosomes to produce offspring inheriting structure from both parents. When genes encode transformations in $\mathrm{SE}(3)$, however, two limitations arise. First, element-wise recombination of parent homogeneous transformation matrices $\mathbf{T}_1, \mathbf{T}_2 \in \mathrm{SE}(3)$ does not produce an element of $\mathrm{SE}(3)$: a convex combination $\mathbf{T} = \lambda\mathbf{T}_1 + (1-\lambda)\mathbf{T}_2$, with $0 \leq \lambda \leq 1$, violates $\mathbf{R}^T\mathbf{R} = \mathbf{I}$ and $\det(\mathbf{R}) = 1$ for $\mathbf{R}_1 \neq \mathbf{R}_2$, where $\mathbf{R}$, $\mathbf{R}_1$, and $\mathbf{R}_2$ are the rotation matrices of $\mathbf{T}, \mathbf{T}_1$, and $\mathbf{T}_2$, yielding an offspring with no mathematical meaning. Second, configurations may be encoded as Euler-angle chromosomes $\mathbf{g} = [x, y, z, \alpha, \beta, \gamma]^T$. Although this guarantees that offspring are valid configurations, crossover then acts directly on a nonlinear, periodic, and singularity-prone parameterization of $\mathrm{SO}(3)$, so recombining two parents generally produces a rotation with no meaningful geometric relationship to either.

Several approaches address these two issues. Quaternion-based crossover [7], [8] avoids Euler-angle singularities in-

Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), [no. de référ. RGPIN-2025-05468]. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding ref. # RGPIN-2025-05468].

cluding gimbal lock, but linear interpolation between unit quaternions traverses the interior of S^3 rather than its surface, and therefore the result is not a valid rotation without renormalization, i.e., projecting the result back onto S^3. This projection biases the offspring distribution toward the parent midpoint [9]. Spherical linear interpolation (SLERP) [9] interpolates along the rotational geodesic instead, but still treats translation and rotation independently [10]. Screw linear interpolation (SCLerp) extends SLERP to the full group $\mathrm{SE}(3)$ via the exponential maps [11] of the parent poses, generating offspring on the geodesic screw connecting them [10], [12]. As a crossover operator, however, SCLerp restricts both offspring to the same one-dimensional geodesic, which may reduce diversity. Crossover in the Lie algebra $\mathfrak{se}(3)$ instead maps each transformation to its twist coordinates and applies a real-coded blend operator in that vector space [13], [14]. While this operation gives geometric consistency and component-wise independence, the logarithm is only locally invertible near the identity. Real-coded operators such as simulated binary crossover (SBX) [15], [16] on Euler-angle chromosomes are widely used and maintain diversity, but treat $\mathrm{SE}(3)$ as flat Euclidean space [17]. Together these illustrate the competing objectives: geometric coherence and population diversity.

This tension is evident in the telesurgical base reference frame (BRF) problem. Because the tool must pivot about a fixed point, the reachable task space is a cone-shaped submanifold of $\mathrm{SE}(3)$ extending from the RCM. Here $\mathrm{SE}(3)$ denotes the set of six-degree-of-freedom (DOF) poses a robot may take, $\mathrm{SO}(3)$ their rotational components, and $\mathfrak{se}(3)$ the associated Lie algebra, i.e., the tangent space at the identity from which poses on $\mathrm{SE}(3)$ are recovered by the exponential map [11]. The BRF determines how much of the cone the robot can reach and how close to singularities the arm operates while doing so. As a result, BRF placement must be set as a multi-objective optimization problem over $\mathrm{SE}(3)$ that maximizes reachability and manipulability simultaneously. The fitness landscape is discontinuous, nonconvex, and highly sensitive to the BRF, and is solved under a budget: setup occurs at the bedside between cases, so the optimizer must return a usable front from few costly kinematic evaluations.

This paper compares four crossover operators, i.e., naive single-point crossover, SBX, SCLerp, and Lie algebra BLX, for a two-objective BRF optimization problem under the non-dominated sorting genetic algorithm (NSGA)-II [18]. The optimization is formulated for a 6-DOF manipulator performing an RCM-constrained task. Each operator is evaluated across 4 population sizes, 10 dexterity thresholds, and 10 trials, and compared on hypervolume, generational distance, inverted generational distance, and convergence. At low population sizes, the two geometry-aware operators outperform the two Euler-angle operators on all three front-quality metrics, with Lie algebra BLX attaining the highest hypervolume, while at large population sizes the dense sampling of the decision variable space leaves all methods within 0.3% of each other in hypervolume.

II. Non-dominated Sorting Genetic Algorithm Crossover Recombination

NSGA-II is a fast, elitist multi-objective evolutionary algorithm [18]. Each instance $\mathbf{x} \in \mathcal{X}$ denotes a candidate solution in its appropriate form, and $f_1, \ldots, f_m$ denote the m real-valued objective functions defined on $\mathcal{X}$ to be minimized. Solution $\mathbf{x}_1$ Pareto-dominates $\mathbf{x}_2$ ($\mathbf{x}_1 \prec \mathbf{x}_2$) if $f_k(\mathbf{x}_1) \le f_k(\mathbf{x}_2)$ for all $k \in \{1, \ldots, m\}$ with strict inequality for at least one k. At each generation the population P of size N_p is ranked by fast non-dominated sorting into successive fronts $\mathcal{F}_1, \mathcal{F}_2, \ldots$, and individuals within a front are ordered by a crowding distance d_s, the normalized objective-space separation of the neighbors of s along each objective, with boundary individuals assigned $d_s = \infty$ so that the current extent of the explored front is preserved.

Binary tournament selection (preferring lower rank, then higher d_s) draws parent pairs, which produce offspring via crossover with probability p_c and mutation with probability p_m per gene. Mutation is identical for all operators: each gene of the Euler-angle chromosome is perturbed with probability p_m by Gaussian noise of standard deviation σ, so differences arise from crossover alone. Parent and child populations are merged into $\mathcal{M} = P \cup Q$ of size $2N_p$, re-ranked, and the next generation is filled by complete fronts in rank order, the last truncated by crowding distance, so the best solutions are never discarded.

If the above is used for element-wise recombination of transformation matrices, it violates the constraints of $\mathrm{SE}(3)$. Chart-based recombination on Euler angle chromosomes produces valid but geometrically arbitrary offspring. To address these limitations, we define four recombination operators.

A. Naive Single-Point Crossover

Given chromosomes $\mathbf{g}^{(\iota)} = [x, y, z, \alpha, \beta, \gamma]^\top$ with $\iota \in \{1, 2\}$ the parent index, a cut point $c \sim \mathcal{U}\{1, \ldots, 5\}$ is sampled once per parent pair, determining the index before and after which genes are traded, yielding two "crossover" genes $\mathbf{g}^{(c_1)}$ and $\mathbf{g}^{(c_2)}$:

$$\mathbf{g}^{(c_1)} = \left[g_1^{(1)}, \ldots, g_c^{(1)}, g_{c+1}^{(2)}, \ldots, g_6^{(2)}\right]^\top, \tag{1}$$

$$\mathbf{g}^{(c_2)} = \left[g_1^{(2)}, \ldots, g_c^{(2)}, g_{c+1}^{(1)}, \ldots, g_6^{(1)}\right]^\top. \tag{2}$$

The cut $c = 3$ exchanges translation and orientation blocks; other cuts mix within blocks. No use is made of manifold geometry.

B. Simulated Binary Crossover (SBX)

SBX acts on each gene d independently with probability 0.5. Given $x_1^{(d)} \le x_2^{(d)}$ and bounds $[lo_d, hi_d]$, spread factors $\beta_u^{(j)}$ for offspring $j \in \{1, 2\}$ are drawn independently from a bounded distribution indexed by $\eta > 0$, determining how far offspring land from their parents:

$$\mathbf{g}_{d,j}^{(c)} = \tfrac{1}{2}\left[\left(\mathbf{g}_d^{(1)} + \mathbf{g}_d^{(2)}\right) \pm \beta_u^{(j)}\left(\mathbf{g}_d^{(2)} - \mathbf{g}_d^{(1)}\right)\right]. \tag{3}$$

When $\beta_u < 1$ offspring interpolate; when $\beta_u > 1$ they extrapolate. SBX is the standard crossover in multi-objective

TABLE I
CROSSOVER OPERATOR PROPERTIES

Operator	Space	Manifold aware	Extra-polates	Independent offspring
Naive	Euler chart	No	No	Yes
SBX	Euler chart	No	Yes	Yes
SCLerp	SE(3) geodesic	Yes	Yes	No
Lie BLX	$\mathfrak{se}(3)$ (flat)	Partial	Yes	Yes

evolutionary algorithms but ignores the periodic, singular structure of the Euler chart.

C. SCLerp: Geodesic Crossover on SE(3)

SCLerp operates on the matrix representations $\mathbf{T}_1, \mathbf{T}_2 \in \mathrm{SE}(3)$ directly. The relative transform $\mathbf{T}_{\text{rel}} = \mathbf{T}_1^{-1}\mathbf{T}_2$ is mapped to $\mathfrak{se}(3)$ via the principal matrix logarithm $\hat{\boldsymbol{\xi}} = \log(\mathbf{T}_{\text{rel}})$, with rotation angle in $[0, \pi]$, and offspring are produced at interpolation parameter $t \sim \mathcal{U}(-0.5, 1.5)$:

$$\mathbf{T}_{c_1} = \mathbf{T}_1 \exp(t\,\hat{\boldsymbol{\xi}}), \quad \mathbf{T}_{c_2} = \mathbf{T}_1 \exp\big((1-t)\hat{\boldsymbol{\xi}}\big). \tag{4}$$

Both offspring lie on the geodesic screw connecting $\mathbf{T}_1$ to $\mathbf{T}_2$. Values of t outside $[0, 1]$ place one or both offspring beyond the parent transformations, enabling extrapolation and preventing population collapse.

D. Lie Algebra BLX

Each parent is mapped to a twist independently via the principal logarithm at the identity: $\boldsymbol{\xi}^{(\iota)} = [\mathbf{v}^{(\iota)\,T}, \boldsymbol{\omega}^{(\iota)\,T}]^T = \mathrm{vec}(\log(\mathbf{T}_\iota))$ for $\iota \in \{1, 2\}$, where $\|\boldsymbol{\omega}\| \in [0, \pi]$ and $\mathbf{v}$ is the linear twist component. BLX is applied per component d as:

$$\ell_d = \min\big(\xi_d^{(1)}, \xi_d^{(2)}\big) - 0.5\,\big|\xi_d^{(1)} - \xi_d^{(2)}\big|, \tag{5}$$

$$u_d = \max\big(\xi_d^{(1)}, \xi_d^{(2)}\big) + 0.5\,\big|\xi_d^{(1)} - \xi_d^{(2)}\big|. \tag{6}$$

with $\xi_{c_m,d} \sim \mathcal{U}(\ell_d, u_d)$ drawn independently for each child $m \in \{1, 2\}$ and each component d. Offspring are recovered via $\mathbf{T}_{c_m} = \exp(\hat{\boldsymbol{\xi}}_{c_m})$ after capping $\|\boldsymbol{\omega}_{c_m}\|$ at π. Unlike SCLerp, the children are not constrained to a shared geodesic. As BLX acts on each component independently, it is invariant to per-component scaling, so no balancing between $\mathbf{v}$ (mm) and $\boldsymbol{\omega}$ (rad) is required.

Table I summarizes the 4 operators, which span a spectrum from fully chart-based to fully geometry-aware.

III. PROBLEM FORMULATION

Let $\mathbf{q} \in \mathcal{Q} \subset \mathbb{R}^n$ denote the configuration of the patient-side serial manipulator. The forward kinematic map $\mathcal{FK}$ maps each configuration in $\mathcal{Q}$ to an end-effector frame in SE(3), and the inverse kinematic map $\mathcal{IK}$ maps each frame in SE(3) to a valid configuration in $\mathcal{Q}$, i.e., one that reproduces the full pose within the joint limits, or to $\emptyset$ if none exists; a closed-form solution for $\mathcal{IK}$ is assumed. Among valid IK branches, the one with the largest manipulability is used. Collisions are not modeled; they can be incorporated by accepting only collision-free IK branches, or as constraints under constrained domination in NSGA-II. Let cone $\mathcal{C}$ with apex at a RCM

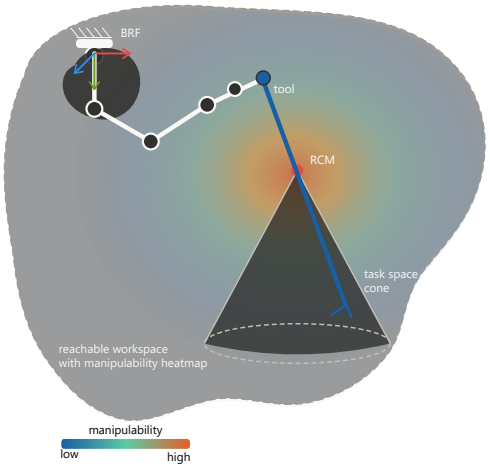


Fig. 1. Optimization problem: the robot's BRF is varied to maximize the manipulability in the conical task space on SE(3) accessible through the RCM, indicated by the heatmap.

point $\mathbf{p}_{\text{RCM}}$, axis $\hat{\mathbf{z}}_{\text{cone}}$, height h, and base radius ρ define the task space, i.e. the set of instrument poses the surgeon may command through the incision. In a frame with origin $\mathbf{p}_{\text{RCM}}$ and z-axis $\hat{\mathbf{z}}_{\text{cone}}$, tool-tip position $\mathbf{p}$ lies inside $\mathcal{C}$ if $h_0 \le \mathbf{p}^\top \hat{\mathbf{z}}_{\text{cone}} \le h$ and $\angle(\mathbf{p}, \hat{\mathbf{z}}_{\text{cone}}) \le \arctan(\rho/h)$, with a small minimum depth h_0. At each $\mathbf{p} \in \mathcal{C}$, the RCM constraint fixes the tool z-axis as $\hat{\mathbf{z}}_{\text{tool}} = \mathbf{p}/\|\mathbf{p}\|$, pointing from the RCM to the tool tip, while the roll $\phi \in [0, 2\pi)$ remains free, so the task space

$$\mathcal{W} = \big\{T(p, \phi) \in SE(3) \mid p \in \mathcal{C},\ \phi \in [0, 2\pi)\big\} \tag{7}$$

is a 4D submanifold of SE(3) (three positional DOF and roll), sampled at N_{pos} positions (depth and cosine of the polar angle drawn uniformly) and N_{roll} evenly spaced roll angles each, giving $N = N_{\text{pos}} \times N_{\text{roll}}$ frames $\{\mathbf{T}_i\}_{i=1}^N$.

A. Optimization Objectives

Given BRF $\mathbf{T}_b \in \mathrm{SE}(3)$, target frame $\mathbf{T}_i$ is expressed in the robot base frame as $\mathbf{T}_i^b = \mathbf{T}_b^{-1}\mathbf{T}_i'$, where $\mathbf{T}_i'$ accounts for the tool offset ℓ along the tool z-axis. $\mathbf{T}_i$ is reachable if $\mathcal{IK}(\mathbf{T}_i^b) \neq \emptyset$. Let $\mathcal{R}$ denote the set of frame indices for which a valid inverse kinematic solution exists, and let $N_r = |\mathcal{R}|$ denote the number of reachable frames. The two objectives are:

1) F1 - Reachability: Maximizing reachability ensures that a large proportion of the task space manifold has an inverse kinematic solution, i.e. a command issued at the console can actually be executed rather than rejected. Reachability is defined as:

$$F_1(T_b) = \frac{N_r}{N} \in [0, 1]. \tag{8}$$

2) F2 - Logarithmic-normalized manipulability: The Yoshikawa index for a non-redundant robot is $w(\mathbf{q}) = |\det(\mathbf{J})|$, where J is the geometric Jacobian. It measures how isotropically joint velocities map to tool velocities and therefore bounds the fidelity with which the master device's commanded motion is reproduced at the follower. The index

is normalized by $w_{\max}$, the largest w over the joint space (identical for all BRFs), and log-transformed to increase the spatial region over which singularities reduce manipulability within the workspace:

$$\mu_i = \log\left(\min\left(\frac{w\big(\mathrm{IK}(\mathbf{T}_b^i)\big)}{w_{\max}},\ 1\right)\right) \in (-\infty, 0]. \tag{9}$$

The function to be maximized, i.e.,

$$F_2(T_b) = \frac{1}{N_r}\sum_{i\in\mathcal{R}} \mu_i \in (-\infty, 0]. \tag{10}$$

makes the robot less likely to operate near singular poses within the task space. Averaging over reachable frames is deliberate: F_2 measures dexterity where the arm can operate and coverage is penalized by F_1, so a BRF reaching few frames with good dexterity lies at the low-F_1 end of the front. The optimization problem combining both objective functions above is:

$$\max_{\mathbf{T}_b\in SE(3)} \big(F_1(\mathbf{T}_b),\ F_2(\mathbf{T}_b)\big), \tag{11}$$

and passed to NSGA-II as $\min(-F_1, -F_2)$.

B. Dexterity Threshold Sweep

The log-normalized manipulability metric in (9) tends to $-\infty$ only when the robot is at a singularity. The logarithmic scale widens the regions around singularities, so F_2 varies smoothly rather than approximating a count of *singular* and *non-singular* configurations as on the linear scale. We therefore define a thresholding hyperparameter τ below which the manipulability is deemed quasi-singular, i.e., below the minimum dexterity level acceptable for teleoperated motion. With $w_i = w(\mathrm{IK}(\mathbf{T}_b^i))/w_{\max}$ the normalized Yoshikawa index for frame i, quasi-singular frames are floored at τ:

$$\tilde{w}_i = \max\big(\min(w_i,\ 1),\ \tau\big), \qquad \mu_i = \log \tilde{w}_i \in [\log\tau,\ 0], \tag{12}$$

so that every quasi-singular frame receives the finite penalty $\log\tau$. Accordingly $F_2 \in [\log\tau, 0]$, and $F_2 = \log\tau$ when $N_r = 0$. We sweep τ across ten values evenly spaced in log-space, $\log(\tau) \in \{-4.00, -3.56, -3.12, -2.68, -2.24, -1.81, -1.37, -0.93, -0.49, -0.05\}$, and run NSGA-II across all methods and population sizes. The sweep tests whether the comparison between operators depends on the threshold, and reveals how the hypervolume achieved by each method varies with the dexterity threshold.

IV. Validation Setup

Evaluation considers the Mecademic Meca500 6-axis serial manipulator with a reach of approximately 300 mm, standing in for a patient-side arm at bench scale, with a rigid 200 mm tool attached to its last joint along its z-axis. Its closed-form IK has up to eight branches, each checked against the joint limits ($\pm175°$, $-70°/90°$, $-135°/70°$, $\pm170°$, $\pm115°$, last joint unbounded). $w_{\max} = 3.41 \times 10^6$ mm^3 is the largest $|\det(\mathbf{J})|$ over 5×10^4 configurations sampled uniformly within these limits. The workspace cone $\mathcal{C}$ has height $h = 90$ mm, base radius $\rho = 50$ mm, and $h_0 = 0.05h$, sampled at $N_{\text{pos}} = 50$ positions and $N_{\text{roll}} = 6$ roll angles, yielding $N = 300$ points per evaluation. Roll changes only the unbounded last joint, on which neither reachability nor w depends, so each position is solved once ($N_{\text{eff}} = 50$ IK solutions per evaluation) with identical objectives.

Each operator is evaluated at population sizes $N_p \in \{5, 10, 50, 100\}$ with $R = 10$ independent random-seed replicates per τ, giving 160 runs per τ and 1600 in total; the four operators share initial populations. NSGA-II parameters are $p_c = 0.9$, $p_m = 0.3$ per gene, mutation strength $\sigma = 0.05 \times$ range, $G_{\max} = 50$, SBX distribution index $\eta = 10$, and BLX extension 0.5. The chromosome is bounded by $x, y \in [-200, 200]$ mm, $z \in [-400, 0]$ mm, $\alpha, \gamma \in [-\pi, \pi]$, and $\beta \in [-\pi/2, \pi/2]$ in the RCM frame, with $\mathbf{R} = \mathbf{R}_z(\alpha)\mathbf{R}_y(\beta)\mathbf{R}_x(\gamma)$. Offspring of every operator, including extrapolated ones, are converted to Euler angles, wrapped to $[-\pi, \pi)$, and clipped to the bounds.

Pareto fronts are evaluated using hypervolume (HV), generational distance (GD), inverted generational distance (IGD), and convergence generation. The hypervolume is computed exactly in the negated objective space $(-F_1, -F_2)$ with reference point $(0, -\log\tau)$, the worst value of each objective, and divided by $|\log\tau|$ so that HV $\in [0, 1]$. Metrics are reported at $\log(\tau) = -2.24$ ($\tau \approx 0.106$), the fifth sampled value. For GD and IGD, let $\mathcal{P} \subset \mathbb{R}^2$ denote the non-dominated solutions obtained at termination and $\mathcal{P}^*$ an independent reference front pooled from three runs per operator at $N_p = G_{\max} = 200$ with separate seeds; $-F_2$ is divided by $|\log\tau|$. Convergence error and reference-front coverage are

$$GD = \frac{1}{|\mathcal{P}|}\sum_{\mathbf{u}\in\mathcal{P}} \min_{\mathbf{u}^*\in\mathcal{P}^*} \|\mathbf{u} - \mathbf{u}^*\|_2, \tag{13}$$

$$IGD = \frac{1}{|\mathcal{P}^*|}\sum_{\mathbf{u}^*\in\mathcal{P}^*} \min_{\mathbf{u}\in\mathcal{P}} \|\mathbf{u}^* - \mathbf{u}\|_2. \tag{14}$$

Lower GD and lower IGD indicate, respectively, that solutions lie close to the reference front, and that all regions are covered. Convergence generation G_{95} is the earliest generation at which the Pareto front hypervolume reaches 95% of its final value.

Operators are compared at each N_p with Welch's t-test over the $R = 10$ replicates, Holm–Bonferroni-corrected over the six pairs (quoted p-values), with an exact Mann–Whitney U test as a check. We also report $N_{\text{IK}} = N_p(G_{95} + 1)\, N_{\text{eff}}$, the number of inverse-kinematic solutions consumed up to convergence, as a hardware-independent proxy for setup cost.

V. Results and Discussion

Fig. 2 shows the final hypervolume against τ (a), the convergence (b), and the final Pareto fronts (c), with scalar metrics in Table II. Hypervolume decreases as τ increases, as a stricter floor removes dexterity margin. At $N_p = 10$ the ranking of all four operators is identical at every τ, and at $N_p = 5$ Lie algebra BLX attains the highest mean hypervolume at nine of the ten thresholds; at $N_p \geq 50$ the order varies because the differences are within replicate noise.

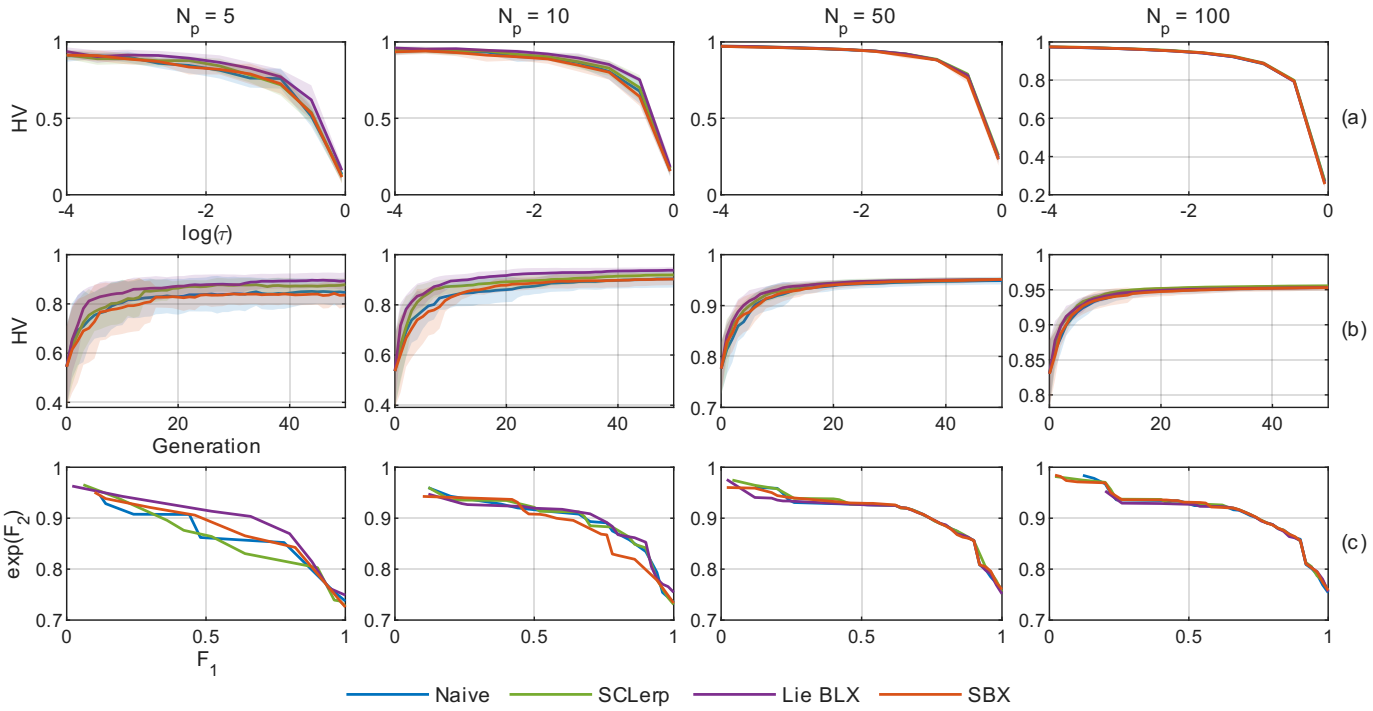


Fig. 2. Comparison between crossover methods for different population sizes N_p: (a) normalized converged hypervolume *versus* the threshold τ (mean $\pm$ std, 10 replicates); (b) hypervolume convergence over generations; (c) final Pareto fronts pooled over replicates. (b) and (c) are at $\log(\tau) = -2.24$.

The small-population comparison is thus not an artifact of the threshold; at $\log(\tau) = -2.24$, under 0.1% of reachable frames on final fronts are at the floor.

A. Small Populations

At $N_p = 5$ and 10 the operators separate according to how they treat the geometry of $\mathrm{SE}(3)$: both geometry-aware operators outperform both Euler-chart operators in mean HV, GD, and IGD. Lie algebra BLX attains the highest hypervolume at both sizes, 0.893 ± 0.034 and 0.938 ± 0.012, against 0.836 ± 0.049 and 0.903 ± 0.024 for SBX, a difference significant after correction ($p = 0.045$ and 0.007). SCLerp ranks second in hypervolume and attains the lowest IGD at both sizes, while Lie algebra BLX attains the lowest GD. In IGD, SCLerp and Lie algebra BLX improve significantly on naive crossover at $N_p = 5$ ($p = 0.042$), and SCLerp on SBX at $N_p = 10$ ($p = 0.021$); other differences are not significant. No operator produced a failed replicate (lowest single-run HV 0.73), and dispersion is similar across operators.

Both chart-based operators recombine Euler angles coordinate-wise, so the offspring of two good placements need not lie near either parent in $\mathrm{SE}(3)$, and a small population cannot compensate by volume. The geometry-aware operators place offspring relative to their parents in $\mathrm{SE}(3)$: on the connecting screw (SCLerp) or in an extended twist-space neighborhood (Lie algebra BLX). Lie algebra BLX, with independently placed children, converges closest to the reference front, while SCLerp covers it slightly more evenly, so sharing one screw did not limit diversity at these budgets. As a placement is optimized once per procedure, a single run matters: at $N_p = 10$ even the worst Lie algebra BLX replicate (0.907) exceeds the mean of both chart-based operators.

B. Large Populations

From $N_p = 50$ onward, all methods reliably cover the Pareto front, as a dense population samples $\mathrm{SE}(3)$ broadly enough to compensate for geometric naivety. Final hypervolumes span 0.949–0.952 at $N_p = 50$ and 0.953–0.955 at $N_p = 100$, under 0.3%; the only significant difference, SCLerp over Lie algebra BLX at $N_p = 100$ ($p = 0.027$), is 0.0025. The lowest IGD is attained by SCLerp at $N_p = 50$ and SBX at $N_p = 100$, neither significantly. Convergence generation does not differ significantly at any N_p; SBX's low value at $N_p = 5$ reflects early stagnation at the lowest final hypervolume. Naive crossover converges slowest at every N_p. The pooled front of all study runs lies within GD $=$ IGD $= 0.005$ of the independent reference, so these comparisons do not depend on the reference.

C. Implications for Telesurgery

Setup time in the operating room scales with the IK solutions consumed (N_{IK}, Table II). The best mean hypervolume in this study, 0.955 (SCLerp, $N_p = 100$), requires 1.2×10^5 IK solutions; Lie algebra BLX reaches 98.1% of it at $N_p = 10$ for 3.75×10^4, $3.2\times$ fewer. Under a bedside evaluation budget, the geometry-aware operators deliver the better front per kinematic evaluation.

The Pareto fronts in row (c) show a consistent negative slope in the F_1 versus $\exp(F_2)$ plane: placing the base nearer the

TABLE II

PERFORMANCE METRICS AT $\log(\tau) = -2.24$. MEAN $\pm$ STD, 10 REPEATS; BEST IN BOLD. GD AND IGD AGAINST THE INDEPENDENT REFERENCE FRONT, SCALED BY 10^2. N_{IK} IS THE NUMBER OF INVERSE-KINEMATIC SOLUTIONS TO CONVERGENCE, $N_p(G_{95}+1)N_{\text{EFF}}$, IN UNITS OF 10^3.

N_p	Metric	Naive	SCLerp	LieAlg.	SBX
5	HV	0.846±0.051	0.878±0.026	**0.893±0.034**	0.836±0.049
	GD ($\times 10^2$)	7.1±2.7	5.3±1.8	**5.2±2.5**	8.2±3.5
	IGD ($\times 10^2$)	14.08±2.87	**10.63±2.07**	10.75±1.86	13.04±4.53
	Conv.	12.2±9.7	11.8±6.4	11.2±8.8	**7.0±4.4**
	N_{IK} ($\times 10^3$)	19.8	19.2	18.3	**12.0**
10	HV	0.903±0.034	0.919±0.020	**0.938±0.012**	0.903±0.024
	GD ($\times 10^2$)	5.6±2.7	4.0±1.3	**2.6±1.9**	5.0±1.8
	IGD ($\times 10^2$)	7.21±1.69	**5.59±1.09**	6.30±2.37	8.80±2.59
	Conv.	16.6±7.8	11.9±6.9	**11.5±8.6**	12.7±4.2
	N_{IK} ($\times 10^3$)	52.8	38.7	**37.5**	41.1
50	HV	0.949±0.008	**0.952±0.004**	0.951±0.001	0.951±0.003
	GD ($\times 10^2$)	**1.0±0.6**	1.1±0.6	1.1±0.4	1.0±0.4
	IGD ($\times 10^2$)	4.92±2.82	**3.23±1.44**	4.06±3.55	4.09±2.93
	Conv.	7.1±2.7	6.4±3.1	**4.5±1.6**	6.8±4.0
	N_{IK} ($\times 10^3$)	121.5	111.0	**82.5**	117.0
100	HV	0.953±0.002	**0.955±0.002**	0.953±0.001	0.953±0.004
	GD ($\times 10^2$)	0.9±0.5	**0.6±0.3**	0.7±0.2	1.1±0.7
	IGD ($\times 10^2$)	2.92±0.86	3.08±3.08	3.03±1.53	**2.23±0.69**
	Conv.	4.1±1.7	**3.0±1.2**	3.1±1.5	3.8±1.7
	N_{IK} ($\times 10^3$)	153.0	**120.0**	123.0	144.0

cone apex increases task-space coverage but leaves less margin from singular postures. The two objectives correspond to two failure modes in teleoperation: low F_1 means the follower cannot execute operator commands, requiring repositioning of the master device; low F_2 means motion is executed with degraded velocity and force transmission, so the master–follower mapping becomes anisotropic, with loss of transparency.

There is no universally optimal BRF: the choice depends on whether task completion or dexterity reserve is the priority, and the front provides the full set of trade-offs. By $N_p = 100$ the fronts of all methods overlap, so this surface is a property of the robot and task rather than of the optimizer.

VI. CONCLUSION

This paper compares 4 crossover operators for placing a teleoperated manipulator base under an RCM constraint. The choice of operator matters most under small evaluation budgets, as in telesurgery. At $N_p \leq 10$, ignoring the geometry of $\mathrm{SE}(3)$ carries a measurable cost: both geometry-aware operators produce better fronts than both Euler-chart operators in hypervolume, GD, and IGD, consistently across dexterity thresholds. Lie algebra BLX, drawing two independently placed offspring in $\mathfrak{se}(3)$, is recommended in this regime, with the highest hypervolume, significantly above SBX; SCLerp is a close second that covers the front most evenly. At $N_p \geq 50$, population density compensates for geometric naivety and all operators lie within 0.3% in hypervolume, at roughly three times the kinematic evaluations of Lie algebra BLX at $N_p = 10$ for under 2% more hypervolume. Naive crossover converges slowest throughout and is not recommended.

Three limitations should be considered. First, with $R = 10$ replicates, several pairwise differences at small N_p do not survive correction, although the ranking is stable across thresholds. Second, SCLerp and Lie algebra BLX use the principal logarithm, discontinuous at rotations of π, which the bounds admit (about 10% of final BRFs lie within 10° of π), so nearby parents can map to distant twists. Third, manipulability is a proxy for teleoperation transparency rather than a measurement of it. Future work can include adaptive crossover selection, a parent-centered logarithm, fixed-budget comparisons, collision constraints, and experiments with a human operator in the loop, relating the Pareto position of a BRF to measured task performance and subjective presence.

REFERENCES

[1] M. Posa *et al.*, "A direct method for trajectory optimization of rigid bodies through contact," *The Inter. Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014.

[2] M. Posa and R. Tedrake, "Direct trajectory optimization of rigid body dynamical systems through contact," in *Tenth Workshop on the Algorithmic Foundations of Robotics*, 2013, pp. 527–542.

[3] A. Hereid and A. D. Ames, "Frost: Fast robot optimization and simulation toolkit," in *Inter. Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 719–726.

[4] J. L. Blanco-Claraco, "A tutorial on $\mathrm{SE}(3)$ transformation parameterizations and on-manifold optimization," *arXiv preprint arXiv:2103.15980*, 2021.

[5] J. Urain *et al.*, "Se (3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion," in *Inter. Conference on Robotics and Automation*, 2023, pp. 5923–5930.

[6] L. Zhang and L. Fang, "Kinematic parameter identification of robotic manipulators based on se (3) lie group error modeling and improved genetic algorithm," in *Journal of Physics: Conference Series*, vol. 3108, no. 1, 2025, p. 012012.

[7] Y. Oshman and A. Carmi, "Attitude estimation from vector observations using a genetic-algorithm-embedded quaternion particle filter," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 4, pp. 879–891, 2006.

[8] T. T. Khuat and M. H. Le, "A genetic algorithm with multi-parent crossover using quaternion representation for numerical function optimization," *Applied Intelligence*, vol. 46, no. 4, pp. 810–826, 2017.

[9] K. Shoemake, "Animating rotation with quaternion curves," in *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, 1985, pp. 245–254.

[10] A. Sarker, A. Sinha, and N. Chakraborty, "On screw linear interpolation for point-to-point path planning," in *Inter. Conference on Intelligent Robots and Systems*, 2020, pp. 9480–9487.

[11] X. Huo, S. Yang, B. Lian, T. Sun, and Y. Song, "A survey of mathematical tools in topology and performance integrated modeling and design of robotic mechanism," *Chinese Journal of Mechanical Engineering*, vol. 33, no. 1, p. 62, 2020.

[12] R. Laha *et al.*, "Point-to-point path planning based on user guidance and screw linear interpolation," in *Inter. Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 85451, 2021, p. V08BT08A010.

[13] A. Marthinsen, "Interpolation in lie groups," *SIAM Journal on Numerical Analysis*, vol. 37, no. 1, pp. 269–285, 1999.

[14] A. Mota *et al.*, "Lie-group interpolation and variational recovery for internal variables," *Computational Mechanics*, vol. 52, no. 6, pp. 1281–1299, 2013.

[15] K. Deb and H.-G. Beyer, "Self-adaptive genetic algorithms with simulated binary crossover," *Evolutionary computation*, vol. 9, no. 2, pp. 197–221, 2001.

[16] K. Deb, R. B. Agrawal *et al.*, "Simulated binary crossover for continuous search space," *Complex systems*, vol. 9, no. 2, pp. 115–148, 1995.

[17] K. Deb, "Multi-objective optimisation using evolutionary algorithms: an introduction," in *Multi-objective evolutionary optimisation for product design and manufacturing*. Springer, 2011, pp. 3–34.

[18] K. Deb *et al.*, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.