A Framework for Combining Optimization-Based and Analytic Inverse Kinematics

Thomas Cohn*, Lihan Tang*, Alexandre Amice, and Russ Tedrake

***Abstract*— Roboticists have used both analytic and optimization methods for solving inverse kinematics (IK) problems. The two strategies have complementary strengths and weaknesses, but developing a unified approach to take advantage of both methods has proved challenging. We present a new formulation for optimization IK that uses an analytic IK solution as a change of variables, which is fundamentally easier for optimizers to solve. Extensive experimental comparisons across three popular solvers demonstrate that our new formulation achieves higher success rates than the old formulation and baseline methods across various challenging IK problems, including collision avoidance, grasp selection, and humanoid stability.**

I. Introduction

Inverse kinematics (IK) is a fundamental problem in robotics [1–5]. For a chain of rigid links connected by movable joints, IK seeks to compute a set of joint positions that achieves a desired Cartesian pose for the end-effector of the chain. Early study of IK largely focused on symbolic analysis [2]. The pose of the end-effector of can naturally be written down as a function of its joint angles, so determining joint angles for a desired target amounts to solving a highly-structured system of trigonometric polynomial equations, a process called *analytic IK*. Implementations of these solutions can usually be obtained directly from robot descriptions (e.g. URDF) via automated meta-solvers [5–8]. With redundant degrees of freedom, the system is underdetermined, requiring additional *self-motion* parameters to disambiguate solutions [9–12].

With the continued improvements to numerical computing over the past decades, numerical algorithms for solving IK have become incredibly prevalent [13]. This has lead to a broader scope, where IK problems are embedded as part of grasping, motion planning, and even humanoid robot stability [14–16]. In its most general form, IK is a broad class of kinematic optimization problems [17, §6.1].

However, optimization IK is not without its challenges. The IK constraint means the optimizer must find a solution satisfying a nonlinear equality constraint. When adding additional constraints to the problem, such as collision avoidance or humanoid stability, optimizers may get stuck in local minima, or even fail to return a feasible guess.

* denotes equal contribution. This project was supported by Amazon.com, PO No. 2D-15693043 and 2D-15694085, the National Science Foundation under Grant No. DMS-2022448, the National Science Foundation Graduate Research Fellowship Program under Grant No. 2141064, and the Robotics and AI Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the other sponsors acknowledged in this work.

We present a unifying approach for solving optimization IK problems for systems that have an analytic solution. The end-effector pose (and self-motion parameters) are the decision variables, and the joint angles are determined in terms of these variables via the analytic IK mapping, allowing us to impose the original costs and constraints. In particular, solving such an optimization problem with gradient-based solvers is enabled by modifying the implementation of an analytic IK function to support automatic differentiation, a strategy presented in [18].

The proposed formulation trades off simplicity of the IK constraint, which is *linear* in the new decision variables, for greater complexity in other costs and constraints. The relative benefit becomes apparent when inspecting the assumptions of standard nonlinear solvers, which operate in either the entire ambient space or the relative interior of the problem domain. This either precludes the use of nonlinear equality constraints altogether, as is the case for AGS [19, §8.1], CMA-ES [20, §B.5], and NFQPIM [21], or requires converting the equality to a pair of inequalities, resulting in a failure of the Linear Independence Constraint Qualification (LICQ) that many solvers rely on for numerical stability [22, §4][23, §2.2].

To evaluate our methodology, we perform a detailed comparison between formulations (and baselines) across three popular solvers representing three main paradigms for constrained nonlinear optimization (interior point, augmented Lagrangian, and sequential quadratic programming), for a variety of scenarios. The new formulation achieves a higher success rate than the old for every solver tested and for every experiment, strongly suggesting that it is fundamentally easier for optimizers to find feasible solutions for the new formulation. The results presented here are a subset of a broader work; the full preprint, containing additional details and results, is available online [24].

II. Methodology

To represent changes of coordinates in Euclidean space, we use monogram notation [17, §3.1]. ${}^{A}X^{B} \in \mathrm{SE}(3)$ is the pose of frame B relative to (and measured in) frame A (a 4×4 homogeneous transformation matrix). When this pose is a function of variables q, we write ${}^{A}X^{B}(q)$, e.g., the pose of an arm's gripper G in the world frame W with joint angles q is written ${}^{W}X^{G}(q)$. The position of frame B relative to (and expressed in) frame A is denoted ${}^{A}p^{B}$, and the orientation is denoted ${}^{A}o^{B}$. The homogeneous transformation matrix described by a position and orientation is $X({}^{A}p^{B}, {}^{A}o^{B})$.

An analytic IK function takes in three arguments:

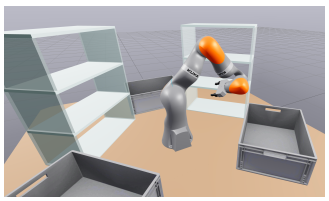

(a) Arm on a Table

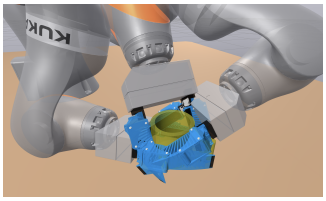

(b) Grasp Selection

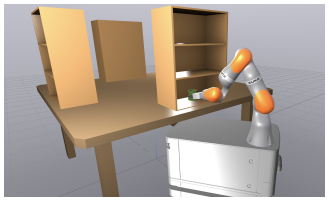

(c) Mobile Robot

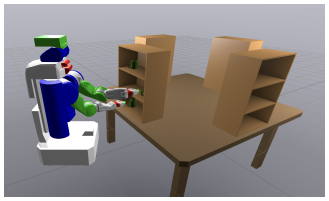

(d) Bimanual Mobile Robot

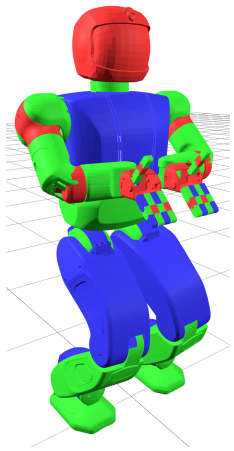

(e) Humanoid Stability

Fig. 1: Experiment setups. (b) shows 3 poses which are feasible for one instance of the grasp selection experiment, with the shelves hidden for visibility. (e) shows the nominal stable configuration used in the Hubo experiments.

1) the target end-effector pose in ${}^{W}X^{G}_{\text{des}} \in \text{SE}(3)$,
2) continuous self-motion parameters $\psi \in \Psi$, which describes the point within a self-motion manifold, and
3) discrete self-motion parameters $\kappa \in \mathscr{K}$, which specifies the self-motion manifold being considered.

The output of the function is a set of joint angles for the arm. So explicitly, an analytic IK function is denoted by

$$\text{IK} : \text{SE}(3) \times \Psi \times \mathscr{K} \to \mathbb{R}^d, \tag{1}$$

where $\mathscr{K}$ denotes the set of self-motion manifolds.

The domain of an analytic IK function typically does not match the true set of reachable configurations. Analytic IK functions generally do not consider joint limits, so that must be added as a constraint. Furthermore, a robot arm generally cannot reach all of SE(3), appearing within analytic IK as domain-restricted functions like arccos (e.g. [11, Eq. (18)]).

To handle the domain restrictions, we require that IK is constructed to clip its inputs for domain-limited functions like arccos, so for a non-reachable configuration, the joint angles output by IK will not actually achieve the specified end-effector pose. We construct *probing functions*

$$\mathcal{D}_k : \text{SE}(3) \times \Psi \times \mathscr{K} \to \mathbb{R} \tag{2}$$

that return intermediate values from within the computations of IK before a clipping operation, such that if $\mathcal{D}_k({}^{W}X^{G}, \psi, \kappa) \geq 0$ for all $k \in [K]$, then the clipping within IK has no effect. For example, if IK requires computing $\arccos(t)$ for some intermediate quantity $t = f({}^{W}X^{G}, \psi, \kappa)$, we would introduce the probing functions

$$\mathcal{D}_1 : ({}^{W}X^{G}, \psi, \kappa) \mapsto 1 - f({}^{W}X^{G}, \psi, \kappa), \tag{3a}$$
$$\mathcal{D}_2 : ({}^{W}X^{G}, \psi, \kappa) \mapsto 1 + f({}^{W}X^{G}, \psi, \kappa). \tag{3b}$$

If $\mathcal{D}_k({}^{W}X^{G}, \psi, \kappa) \geq 0$ for each $k \in [K]$, then ${}^{W}X^{G}$ is a reachable configuration.

Imposing this *reachability constraint* in our optimization problems is essential for convergence. This approach differs from prior work, which directly checked that the resulting end-effector pose matched the input to IK [18]. In that formulation, the reachability constraint is always active, even when the end-effector pose is not on the boundary of the reachable set. With probing functions, the reachability constraint is inactive when the end-effector pose is in the interior of the reachable set.

Given the above machinery, we now construct the old and new formulations for an optimization IK problem. Let $q_{\text{lb}}, q_{\text{ub}} \in (\mathbb{R} \cup \{\pm\infty\})^d$ be the lower and upper joint limits, and ${}^{W}X^{G}_{\text{des}}$ the desired configuration of the end-effector. Finally, we have generic functions $f, g_i, h_j : \mathbb{R}^d \to \mathbb{R}$ for $i \in [I] = \{1, \ldots, I\}$ and $j \in [J]$ to represent the cost, inequality constraints, and equality constraints (respectively).

The standard formulation for optimization IK is

$$\min_{q} \; f(q) \tag{4a}$$
$$\text{s.t.} \; q \in \mathbb{R}^d \tag{4b}$$
$${}^{W}X^{G}(q) = {}^{W}X^{G}_{\text{des}} \tag{4c}$$
$$q_{\text{lb}} \leq q \leq q_{\text{ub}}, \tag{4d}$$
$$g_i(q) \leq 0, \qquad \forall i \in [I], \tag{4e}$$
$$h_j(q) = 0, \qquad \forall j \in [J]. \tag{4f}$$

Here, the joint angles are the decision variables, and the underlying "inverse kinematics" constraint is encoded in (4c). f is a user-defined cost function – joint-centering is a common choice. g_i and h_j are generic additional constraints, such as collision-avoidance or humanoid stability.

The variable q in (4) and constraint (4c) are readily eliminated by the analytic IK function (1). This results in an optimization over the end-effector position p and orientation o, the continuous self-motion parameters $\psi \in \Psi$, and the discrete self-motion parameters $\kappa \in \mathscr{K}$,

$$\min_{p,o,\psi,\kappa} \; f\left(\text{IK}(X(p,o), \psi, \kappa)\right) \tag{5a}$$
$$\text{s.t.} \; p \in \mathbb{R}^3, o \in \mathbb{T}^3, \psi \in \Psi, \kappa \in \mathscr{K}, \tag{5b}$$
$$p = {}^{W}p^{G}_{\text{des}}, \; o = {}^{W}o^{G}_{\text{des}}, \tag{5c}$$
$$q_{\text{lb}} \leq \text{IK}(X(p,o), \psi, \kappa) \leq q_{\text{ub}}, \tag{5d}$$
$$g_i\left(\text{IK}(X(p,o), \psi, \kappa)\right) \leq 0 \qquad \forall i \in [I], \tag{5e}$$
$$h_j\left(\text{IK}(X(p,o), \psi, \kappa)\right) = 0 \qquad \forall j \in [J], \tag{5f}$$
$$D_k(p, o, \psi) \geq 0 \qquad \forall k \in [K], \tag{5g}$$

To solve this mixed-integer nonconvex optimization problem, we fix a choice of $\kappa_0 \in \mathscr{K}$. (Since there are few choices for κ_0, one could easily solve the problem for each possible choice.) The reachability constraint (5g) requires that the joint angles output from the IK function actually achieve

the expected gripper pose. As the intent of our approach is to remove nonlinear equality constraints, we represent orientation with Euler angles. (One could also use exponential coordinates, but not unit quaternions or rotation matrices).

We draw the reader's attention to several key aspects of this optimization problem. The most important part is the elimination of the nonlinear equality constraint (4c) – the only remaining nonlinear equality constraints are the additional user-specified ones. Most solvers will be able to satisfy the linear constraints up to the accuracy of a linear system solve, ensuring feasibility of the IK constraint to a very precise tolerance. With additional nonlinear equality constraints, the performance difference between the two formulations is less clear, but empirical evidence suggests we can still match the performance of the old formulation in runtime and success rate.

III. Experiments

Throughout our experiments, we compare our new formulation to the old formulation with three solvers: the interior point solver IPOPT [23], the sequential quadratic programming solver SNOPT [22], and an augmented Lagrangian solver [25, 26] made available via the NLOPT interface [27]. As a baseline, we compare with simply drawing random samples from the self-motion manifold and picking the lowest-cost feasible one. For our first experiment, we also compare to a global mixed-integer convex relaxation, Global-IK [28]. (It cannot be used for later experiments due to scaling issues and unsupported costs/constraints.) We present a "fast" and "precise" setting, where the latter uses a finer (but more costly) relaxation, and a tighter approximation of the collision-avoidance constraint. The experimental setups are visualized in Fig. 1, and success rates, optimal costs, and runtimes for these experiments are included in Tables I–III.

Arm on a Table (Fig. 1a) is a classic collision-free IK problem: a 7DoF KUKA iiwa arm on a table, surrounded by shelves and bins. We constrain the gripper to a specific target pose and impose joint limits, collision-avoidance, and a quadratic joint centering cost. We selected 100 random targets and used 100 random initial guesses per target.

Grasp Selection (Fig. 1b) expands the previous problem to target a range of poses; we require that a specified point in between the fingers be coincident with the central axis of a mug. We selected 40 random targets and used 40 random initial guesses per target.

Mobile Manipulator (Fig. 1c) requires solving IK for an arm mounted on a mobile base that can move around a table with shelves. We return to targeting a specific gripper pose, and solve for a very similar optimization problem to *Arm on a Table*; we fix the end effector pose and impose joint limits and collision-avoidance. We selected 100 end-effector poses within the shelves and used 100 random collision-free initial guesses per target.

Bimanual Manipulator (Fig. 1d) also requires solving IK with a mobile base, but this time, two end-effector targets are specified. Each constraint corresponds to a separate kinematic chain, with the base pose and torso lift joint coupling them. For the new formulation, the decision variables are the left and right end-effector poses, the respective self-motion parameters, the base pose, and the torso lift. We selected 40 random targets and used 40 random initial guesses per target.

Humanoid Stability (Fig. 1e) considers a humanoid reaching task, where the robot must reach to a given target pose, while simultaneously placing its feet so as to guarantee stability. If both feet are flat on the ground, static stability can be ensured by keeping the center of mass above the *support polytope*: the convex hull of the points on the foot making contact with the ground. These quantities can all be computed efficiently as functions of the robot's configuration, making them amenable for optimization. In addition to stability, we fix the end-effector pose of the right hand to a randomly-selected target, and enforce joint limits and collision avoidance. For the new formulation, we added a log barrier cost on the reachability constraints, in effect forcing this constraint to be satisfied at every intermediate iteration. We do not include any joint centering, as we found it led the optimizer to sacrifice feasibility to achieve a lower objective value, and ultimately get stuck.

In the new formulation, static stability is the only nonlinear equality constraint. However, its feasible set is positive-measure in the end-effector and redundancy variables, and we can rewrite this constraint as to be inequality-only; by enumerating the $\binom{8}{3}$ possible simplices represented by the four corners of the two feet, Carathéodory's theorem [29] allows us to write static stability as a disjunctive inequality constraint. We solved IK for 100 random right-hand targets with both stability constraint formulations, each time using the known stable initial guess.

IV. Discussion

We have presented a new formulation for optimization-based inverse kinematics, that leverages a known analytic solution for the robot as a change of variables. This allows us to write IK optimization problems with the end-effector and self-motion parameters as decision variables, eliminating nonlinear equality constraints. We compare our approach with the corresponding standard formulation across three solvers, each representing a different broad class of nonlinear optimization algorithms, and we also compare to two baseline algorithms. Diverse practical experiments include shelf reaching for fixed and mobile robots, grasp selection, and humanoid stability, demonstrating relevancy to realistic robotics tasks under consideration today.

When comparing the two formulations, the new formulation consistently achieved equal or higher success rates. The new formulation is consistently able to obtain extremely feasibility tolerances (on the order of 10^{-17} or better), many orders of magnitude below what solvers can normally achieve in the presence of nonlinear equality constraints. But results were mixed as to which formulation is faster, and the poorer performance in optimal cost remains a significant limitation of the new formulation. Overall, our results strongly suggest that this new formulation is useful in practice, particularly when additional challenging constraints are present and current success rates are low.

Experiment	d	d'	Interior Point (IPOPT)		Augmented Lagrangian (NLOPT)		Sequential Quadratic Programming (SNOPT)		Global-IK		Sampling
			New	Old	New	Old	New	Old	Fast	Precise	
Arm on a Table	7	1	**0.8951**	0.5392	**0.8733**	0.049	**0.9144**	0.463	0.88	0.97	1.00 ★
Grasp Selection	7	5	**0.9893** ★	0.9875	**0.47**	**0.47**	**0.7344**	0.6625	N/A	N/A	0.525
Mobile Manipulator	10	4	**0.7182**	0.5946	**0.405**	0.0074	**0.368**	0.2884	N/A	N/A	0.95 ★
Bimanual Mobile Manipulator	18	6	**0.7819** ★	0.5243	**0.68**	N/A	**0.3289**	0.2466	N/A	N/A	0.32
Humanoid Stability (Inequality)	31	19	**0.97** ★	0.70	**0.67**	N/A	N/A	N/A	N/A	N/A	N/A
Humanoid Stability (Equality)	31	19	**0.96** ★	**0.96** ★	**0.46**	N/A	N/A	N/A	N/A	N/A	N/A

TABLE I: Success rates for each experiment, where any feasible solution is treated as a success. For each solver, the formulation with the higher success rate is bolded, and the overall highest success rate has a gold star. We also report the dimension of the configuration space (abbreviated d) and the dimension of the constraint manifold (abbreviated d'). Each general nonlinear solver achieves a equal or better success rate with the new formulation than the old in all but one case. The Global-IK baseline was only tractable for the first experiment due to runtime issues. The sampling baseline is only effective when d' is low and requires careful tuning on a per-experiment basis.

Experiment	Interior Point (IPOPT)		Augmented Lagrangian (NLOPT)		Sequential Quadratic Programming (SNOPT)		Global-IK		Sampling
	New	Old	New	Old	New	Old	Fast	Precise	
Arm on a Table	**14.1**	14.6	14.3	**11.3**	**14.2**	15.0	5.5 ★	5.8	6.1
Grasp Selection	9.22	**5.83**	10.9	**3.6** ★	**10.9**	14.0	N/A	N/A	7.44
Mobile Manipulator	8.46	**7.28**	8.62	**3.44** ★	10.8	**6.81**	N/A	N/A	14.58
Bimanual Mobile Manipulator	35.2	**25.5** ★	**30.6**	N/A	37.7	**29.8**	N/A	N/A	43.8

TABLE II: Average optimal costs for each experiment, restricted to successes. For each solver, the formulation with the better optimal cost is bolded, and the overall lowest cost has a gold star. We do not include the humanoid stability experiments, as these were pure feasibility problems. The old formulation often finds lower cost solutions, in part because it has a simpler cost landscape and is not restricted to a single branch of the IK function.

Experiment	Interior Point (IPOPT)		Augmented Lagrangian (NLOPT)		Sequential Quadratic Programming (SNOPT)		Global-IK		Sampling
	New	Old	New	Old	New	Old	Fast	Precise	
Arm on a Table	0.14 (0.11)	**0.07 (0.04)**	0.98 (0.84)	**0.24 (0.16)**	0.025 (0.025)	**0.02 (0.02)** ★	289	3362	0.101
Grasp Selection	0.45 (0.46)	**0.15 (0.15)** ★	2.68 (3.84)	**0.81 (1.13)**	1.07 (1.47)	**0.36 (0.26)**	N/A	N/A	1.69
Mobile Manipulator	0.68 (0.52)	**0.17 (0.12)** ★	1.37 (3.52)	**0.23 (0.32)**	0.90 (1.17)	**0.18 (0.17)**	N/A	N/A	0.60
Bimanual Mobile Manipulator	1.33 (0.64)	**0.40 (0.16)**	**6.01 (4.76)**	N/A	2.25 (2.30)	**0.18 (0.09)** ★	N/A	N/A	1.73
Humanoid Stability (Inequality)	**1.26 (1.12)** ★	3.45 (1.82)	**25.82 (17.82)**	N/A	N/A	N/A	N/A	N/A	N/A
Humanoid Stability (Equality)	**1.61 (1.31)** ★	1.65 (1.50)	**36.37 (24.00)**	N/A	N/A	N/A	N/A	N/A	N/A

TABLE III: Mean runtimes for each experiment in seconds. For each solver, the formulation with the shortest runtime is bolded, and the overall shortest runtime has a gold star. For the nonlinear optimizers, the entry in parentheses is restricted to only the successes. (This is not done for Global-IK, since we do not set a timeout, and not done for sampling, since its runtime is constant.)

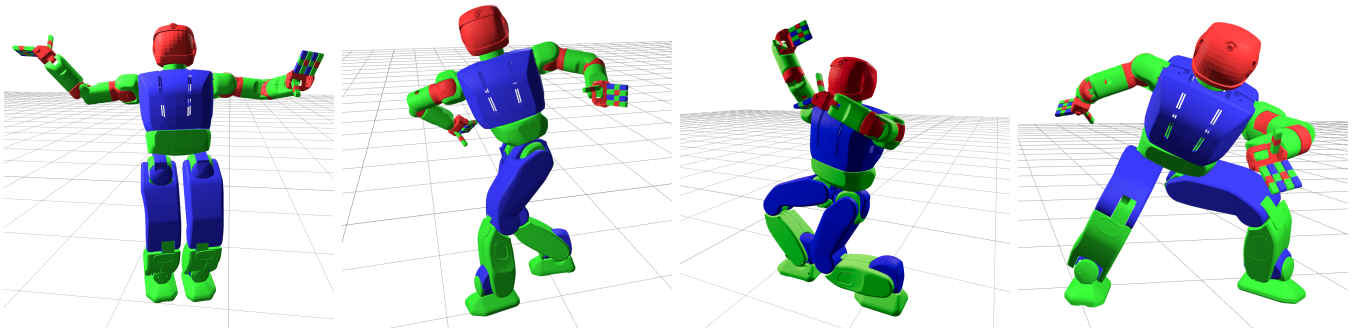

(a) Old formulation, equality. (b) Old formulation, inequality. (c) New formulation, equality. (d) New formulation, inequality.

Fig. 2: Diverse solutions from the Hubo experiments, from the different formulations and stability constraint representations. All solutions were obtained with IPOPT.

REFERENCES

[1] V. Lakshmi Narayanan, J. Narayan, H. Gritli, and S. K. Dwivedy, "A decade of inverse kinematics methods for serial manipulators: A systematic review," *Journal of Field Robotics*, vol. 43, no. 1, pp. 184–229, 2026.

[2] D. L. Pieper, "The kinematics of manipulators under computer control," Ph.D. dissertation, Stanford University, 1969.

[3] L. G. Herrera-Bendezu, E. Mu, and J. T. Cain, "Symbolic computation of robot manipulator kinematics," in *Proceedings. 1988 IEEE International Conference on Robotics and Automation*. IEEE, 1988, pp. 993–998.

[4] M. Raghavan and B. Roth, "Kinematic analysis of the 6R manipulator of general geometry," in *International symposium on robotics research*. Citeseer, 1990, pp. 314–320.

[5] R. Diankov, "Automated construction of robotic manipulation programs," Ph.D. dissertation, Carnegie Mellon University, USA, 2010.

[6] D. Zhang and B. Hannaford, "IKBT: Solving symbolic inverse kinematics with behavior tree," *Journal of Artificial Intelligence Research*, vol. 65, pp. 457–486, 2019.

[7] A. J. Elias and J. T. Wen, "IK-Geo: Unified robot inverse kinematics using subproblem decomposition," *Mechanism and Machine Theory*, vol. 209, p. 105971, 2025.

[8] D. Ostermeier, J. Külz, and M. Althoff, "Automatic geometric decomposition for analytical inverse kinematics," *IEEE Robotics and Automation Letters*, 2025.

[9] D. H. Gottlieb, "Topology and the robot arm," *Acta Applicandae Mathematica*, vol. 11, no. 2, pp. 117–121, 1988.

[10] C. L. Luck and S. Lee, "Self-motion topology for redundant manipulators with joint limits," in *[1993] Proceedings IEEE International Conference on Robotics and Automation*. IEEE, 1993, pp. 626–631.

[11] C. Faria, F. Ferreira, W. Erlhagen, S. Monteiro, and E. Bicho, "Position-based kinematics for 7-DoF serial manipulators with global configuration control, joint limit and singularity avoidance," *Mechanism and Machine Theory*, vol. 121, pp. 317–334, 2018.

[12] Y. He and S. Liu, "Analytical inverse kinematics for Franka Emika Panda–a geometrical solver for 7-dof manipulators with unconventional design," in *2021 9th International Conference on Control, Mechatronics and Automation (ICCMA)*. IEEE, 2021, pp. 194–199.

[13] S. R. Buss, "Introduction to inverse kinematics with Jacobian transpose, pseudoinverse and damped least squares methods," *IEEE Journal of Robotics and Automation*, vol. 17, no. 1-19, p. 16, 2004.

[14] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 295–302.

[15] M. Fallon, S. Kuindersma, S. Karumanchi, M. Antone, T. Schneider, H. Dai, C. P. D'Arpino, R. Deits, M. DiCicco, D. Fourie, *et al.*, "An architecture for online affordance-based perception and whole-body planning," *Journal of Field Robotics*, vol. 32, no. 2, pp. 229–254, 2015.

[16] P. Marion, M. Fallon, R. Deits, A. Valenzuela, C. P. D'Arpino, G. Izatt, L. Manuelli, M. Antone, H. Dai, T. Koolen, *et al.*, "Director: A user interface designed for robot operation with shared autonomy," in *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*. Springer, 2018, pp. 237–270.

[17] R. Tedrake, *Robotic Manipulation*, 2024. [Online]. Available: http://manipulation.mit.edu

[18] T. Cohn, S. Shaw, M. Simchowitz, and R. Tedrake, "Constrained bimanual planning with analytic inverse kinematics," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6935–6942.

[19] R. G. Strongin and Y. D. Sergeyev, *Global optimization with non-convex constraints: Sequential and parallel algorithms*. Springer Science & Business Media, 2013, vol. 45.

[20] N. Hansen, "The CMA evolution strategy: A tutorial," *arXiv preprint arXiv:1604.00772*, 2016.

[21] Y. Shang, Z.-F. Jin, and D. Pu, "A new filter QP-free method for the nonlinear inequality constrained optimization problem," *Journal of Inequalities and Applications*, vol. 2018, no. 1, p. 278, 2018.

[22] P. E. Gill, W. Murray, and M. A. Saunders, "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM review*, vol. 47, no. 1, pp. 99–131, 2005.

[23] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[24] T. Cohn, L. Tang, A. Amice, and R. Tedrake, "A framework for combining optimization-based and analytic inverse kinematics," *arXiv preprint arXiv:2602.05092*, 2026.

[25] A. R. Conn, N. I. M. Gould, and P. Toint, "A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds," *SIAM Journal on Numerical Analysis*, vol. 28, pp. 545–572, 1991.

[26] E. Birgin and J. Martínez, "Improving ultimate convergence of an augmented Lagrangian method," *Optimization Methods and Software*, vol. 23, pp. 177–195, 2008.

[27] S. G. Johnson, "The NLopt nonlinear-optimization package," https://github.com/stevengj/nlopt, 2007.

[28] H. Dai, G. Izatt, and R. Tedrake, "Global inverse kinematics via mixed-integer convex optimization," *The International Journal of Robotics Research*, vol. 38, no. 12-13, pp. 1420–1441, 2019.

[29] C. Carathéodory, "Über den variabilitätsbereich der Fourier'schen Konstanten von positiven harmonischen funktionen," *Rendiconti Del Circolo Matematico di Palermo (1884-1940)*, vol. 32, no. 1, pp. 193–217, 1911.