

# Globally Optimal Inverse Kinematics as a Non-Convex Quadratically Constrained Quadratic Program

Tomáš Votroubek<sup>1</sup> and Tomáš Kroupa<sup>1</sup>

**Abstract**—We show how to compute globally optimal solutions to inverse kinematics by formulating the problem as a non-convex quadratically constrained quadratic program. Our approach makes solving inverse kinematics instances of generic redundant manipulators feasible. We demonstrate the performance on randomly generated designs and real-world robots with up to ten revolute joints. The same technique can be used for manipulator design by introducing kinematic parameters as variables.

## I. INTRODUCTION

Inverse kinematics (IK) is the problem of finding control parameters that bring a robot into a desired pose within its kinematic constraints. For manipulators of six or fewer joints, this problem is practically solvable and has at most 16 different solutions [1]. Adding additional joints results in an under-constrained problem with infinitely many feasible solutions satisfying the pose constraints. Therefore, it is interesting not simply to look for feasible solutions but optimal ones. In this paper, we focus on minimizing the total movement of joints. Specifically, the task is as follows: Given the Denavit-Hartenberg parameters [2] of a kinematic chain, input joint angles, and a desired pose for its end effector, find joint angles that achieve the desired pose while minimizing the weighted sum of differences from the input (preferred) angles.

The presented problem is very hard in general. The current state-of-the-art technique is based on *sum of squares* (SOS) optimization and Gröbner bases, which are both NP-hard in the context of multivariate polynomial equations. Trutman *et al.* demonstrate [3] that for manipulators with at most seven joints, the optimal solutions admit a quadratic SOS representation, thereby being practically solvable. Such results do not extend to manipulators with more than seven joints.

A common way to avoid the complexity of IK is either to find an approximate solution or to design the robotic arm to be simple to control. Approaches from the first category include an approximation based on the Jacobian by Buss [4], which uses numerical methods to find local optima. Similarly, Dai *et al.* [5] designed a global solution method for the convex approximation of the original IK. On the other hand, the design of the KUKA LBR iiwa manipulator falls into the second category. However, having to design around

computational constraints has the downside of limiting the achievable performance of a manipulator.

The problem may be further relaxed for real-time control by dropping the guarantees on optimality or convergence to the desired pose. For example, iKinSiv finds a locally optimal solution which minimizes the distance from the desired pose using interior point optimization; the closed-loop inverse kinematics algorithm suggested in the Pinocchio library attempts to iteratively reduce the pose error, while modern approaches [6] based on convex *quadratic programming* (QP) rely on asymptotic convergence.

Global solvers recover a globally optimal solution whenever it exists or guarantee that none exists. In the context of mechanical design, it is possible to optimize a design by computing and comparing the actual limits of performance in some desired metric specified by the user. Using the optimality guarantees to test existing solvers' properties is also possible. If a solver fails to provide a solution, it is possible to answer whether the instance was solvable.

The contribution of this paper is based on the premise that while SOS optimization is, in some sense, the correct answer to the *polynomial optimization problem* (POP), incumbent implementations fall significantly short of their promised potential. We will show that highly optimized off-the-shelf global optimizers can converge faster for problems such as IK despite having much higher computational complexity. Instead of reformulating IK into a theoretically easier problem, we reformulate it into a problem for which efficient solvers exist.

Our approach finds globally optimal solutions of generic seven *degree-of-freedom* (DOF) IK instances faster than the state-of-the-art. Our technique is unique in finding globally optimal IK solutions for generic manipulators with eight or more degrees of freedom.

## II. QUADRATIC PROGRAM FORMULATION

Starting with the POP for globally solving IK recently proposed by Trutman *et al.* [3], we lift it into a *quadratically constrained quadratic program* (QCQP). Specifically, we find a new basis in lifting variables of the feasible set of the initial POP, such that the new higher-dimensional representation is at most quadratically constrained. For example, for a polynomial constraint  $x_1x_2x_3 = 0$  we introduce a variable  $y$  and rewrite the constraint as  $y = x_1x_2$  and  $yx_3 = 0$ . Solvers for non-convex QCQPs, such as SCIP [7] or Gurobi [8], can find the global solution to this new problem by the spatial branch-and-bound method.

The authors acknowledge the support of the OP VVV funded project CZ.02.1.01/0.0/0.0/16.019/0000765 "Research Center for Informatics". This research was partly funded by AFOSR award FA9550-18-1-0097.

<sup>1</sup>Tomáš Votroubek and Tomáš Kroupa are with Faculty of Electrical Engineering, Department of Computer Science, Czech Technical University in Prague, Czechia [tomas.votroubek@fel.cvut.cz](mailto:tomas.votroubek@fel.cvut.cz)

### A. Denavit–Hartenberg parameters

It is conventional to describe  $n$ -joint manipulators in terms of their *Denavit–Hartenberg* (D-H) parameters for each link  $i = 1, \dots, n$ : the offset  $d_i$ , the link length  $r_i$ , the link twist  $\alpha_i$ , and the joint angle variable  $\theta_i$ .




Fig. 1. Diagram of Denavit–Hartenberg parameters

The transformation between successive coordinate frames from link  $i$  to  $i - 1$  is defined by a D-H matrix  $T_i(\theta_i) =$

$$\begin{bmatrix} \cos \theta_i & -\cos \alpha_i \sin \theta_i & \sin \alpha_i \sin \theta_i & r_i \cos \theta_i \\ \sin \theta_i & \cos \alpha_i \cos \theta_i & -\sin \alpha_i \cos \theta_i & r_i \sin \theta_i \\ 0 & \sin \alpha_i & \cos \alpha_i & d_i \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

The end effector pose w.r.t. the base coordinate system is represented by the matrix  $P$ , which induces the *pose constraint*  $T_1(\theta_1) \cdots T_n(\theta_n) = P$ .

### B. Polynomial Formulation

In this section we describe the POP for IK initially formulated by Trutman *et al.* [3]. This formulation avoids trigonometric functions in the D-H matrices. For each link  $i = 1, \dots, n$  introduce variables  $c_i := \cos \theta_i$  and  $s_i := \sin \theta_i$ , along with the identity  $c_i^2 + s_i^2 = 1$ . After the change of variables, every D-H matrix  $T_i(\theta_i)$  depends only on  $c_i$  and  $s_i$ , and we denote it simply by  $T_i$ . We assume that each joint angle  $\theta_i$  is constrained by the maximal value  $\theta_i^{max} \in [0, \pi]$  and the minimal value  $\theta_i^{min} \in [-\pi, 0]$ . The vector of preferred joint angles is  $(\hat{\theta}_1, \dots, \hat{\theta}_n)$  and we consider some weights  $w_1, \dots, w_n \geq 0$  satisfying  $w_1 + \dots + w_n = 1$ . The goal is to minimize the weighted distance

$$\sum_{i=1}^n w_i ((c_i - \cos \hat{\theta}_i)^2 + (s_i - \sin \hat{\theta}_i)^2). \quad (1)$$

The choice of preferred joint angles  $(\hat{\theta}_1, \dots, \hat{\theta}_n)$  makes it possible to minimize the total movement of joints, or to find configurations where some links achieve desirable orientations (e.g. an upright posture in humanoid robots [9]).

Further, we split the kinematic chain with  $n$  joints into two subchains that meet in the middle<sup>1</sup>: the first anchored at the origin 1 and ending at  $\nu$  and the second starting at  $\nu + 1$  and ending at  $n$ . The pose constraint  $T_1(\theta_1) \cdots T_n(\theta_n) = P$  then becomes  $T_1 \cdots T_\nu = PT_n^{-1} \cdots T_{\nu+1}^{-1}$ . This formulation

<sup>1</sup>The placement of the split does not significantly affect the runtime as long as each subchain involves at least two joints.

enables us to lower the degree of polynomials appearing in the above constraint and, consequently, to reduce the number of lifting variables (see Section II-C).

The resulting POP (see [3, (9)]) with variables  $\mathbf{c} = (c_1, \dots, c_n)$ ,  $\mathbf{s} = (s_1, \dots, s_n)$ , the objective function (1), and the constraints formulated above is:

$$\min_{\mathbf{c}, \mathbf{s}} \sum_{i=1}^n 2w_i (1 - c_i \cos \hat{\theta}_i - s_i \sin \hat{\theta}_i) \quad (2)$$

$$\text{s.t.} \quad PT_n^{-1} \cdots T_{\nu+1}^{-1} = T_1 \cdots T_\nu \quad (3)$$

$$(c_i + 1) \tan \frac{\theta_i^{min}}{2} \leq s_i \quad i = 1, \dots, n \quad (4)$$

$$(c_i + 1) \tan \frac{\theta_i^{max}}{2} \geq s_i \quad i = 1, \dots, n \quad (5)$$

$$c_i^2 + s_i^2 = 1 \quad i = 1, \dots, n \quad (6)$$

Since (1) is equal to the linear function in (2), this POP is in fact a minimization problem with a linear objective and non-convex polynomial constraints. Linear inequalities (4)–(5) are equivalent to design constraints  $\theta_i^{min} \leq \theta_i \leq \theta_i^{max}$  by basic trigonometric identities and the convention  $0 \leq \theta_i^{max} \leq \pi$ ,  $-\pi \leq \theta_i^{min} \leq 0$ . We note that the last set of constraints from the original POP [3, (9)] is redundant.

We use  $\theta_i = \text{atan2}(s_i, c_i)$  to recover the optimal joint angles  $\theta_i$ . In the next section, we show how to make the POP above into an equivalent QCQP, that is, how to reduce the degree of polynomials in the constraints (3).

### C. Lifting Algorithms

The SOS-based method [3] to solve the non-convex POP (2)–(6) uses the hierarchy of moment relaxations [10], which are convex optimization problems and whose values converge to the optimal value of the initial POP. Thus, instead of optimizing a POP directly, the relaxations optimize its linearization using semidefinite programming. This is based on the following convexification idea. For a function  $p$  over  $X$ , the minimizer of the expected value  $\int_X p(x) d\mu(x)$  is a probability measure  $\mu$  supported by the minimizers of function  $p$ . In practice, this approach iteratively attempts to represent probability measures using truncated pseudo-moment sequences. However, to obtain a global solution rather than a relaxation, it is sufficient to constrain the first moment matrix of  $\mu$  to be rank one [11]. This constraint can be encoded in a QCQP and ensures that  $\mu$  is a Dirac measure, which guarantees the solution's global optimality. On the one hand, this approach is computationally expensive as it introduces an excessive number of variables and many redundant constraints even after picking an appropriate basis. On the other hand, the computational methods for POPs based on hierarchies of convex relaxations are certifiably globally optimal.

As an alternative, we approach the problem from the opposite direction, in a way that exploits the problem's sparsity and is also easier for a presolver to optimize. We rewrite each polynomial constraint of the POP into a set of equivalent quadratic constraints in a higher dimensional space. That is, we rewrite the monomials of each polynomial

constraint by repeatedly substituting the multiplication of two optimization variables with a new lifting variable equal to their product. For example, the monomial  $x_1x_2x_3 \cdots x_n$  could be lifted as  $y_{12}x_3 \cdots x_n$ , where  $y_{12} = x_1x_2$ . Each application of such a rewriting reduces the degree of the polynomial and can be applied repeatedly to each monomial of the original expression until only single variables remain. Since the lifting variables are quadratically constrained, we effectively transform the non-convex POP (2)–(6) into a non-convex QCQP. Such a straight-forward lifting technique was already discussed by Shor [12].

Due to the structure of the pose constraint, we can apply the same rewriting to the expressions involving D-H matrices directly. Since the result of  $T_iT_j$  will contain at most quadratic terms, we can introduce a matrix of lifting variables for every matrix product and proceed as described above.

The number of variables and constraints in the resulting program depends on the order of application of such a rewriting rule. For example, the constraint  $x_1x_2x_3x_4 + x_2x_3x_5 = 0$  can be reformulated in at least two different ways:

$$\begin{aligned} 0 &= y_{12}y_{34} + y_{23}x_5 & 0 &= y_{123}x_4 + y_{23}x_5 \\ y_{12} &= x_1x_2 & \text{(A)} \quad y_{23} &= x_2x_3 & \text{(B)} \\ y_{34} &= x_3x_4 & y_{123} &= x_1y_{23} \\ y_{23} &= x_2x_3 \end{aligned}$$

Taking advantage of repeated terms, and even the order in which products are lifted, has an observable impact on the speed and behavior of solvers and on the cuts that they apply.

In the next section, we will focus on the performance of the method corresponding to System A and the matrix based method. We will refer to these methods as  $\text{lift}_A$  and  $\text{lift}_M$ , respectively. While  $\text{lift}_M$  often results in the fewest variables after presolve, it is not always the fastest. We found the behavior of  $\text{lift}_A$  more predictable. We refer the interested reader to our code repository for the implementation details.

#### D. Global Optimality Certificates

The global optimality certificate for the solution provided by the SOS-based method [3] is based on evaluating the rank of the matrix of floats returned by the solver.

*Branch and bound* (BB) algorithms are global optimization methods. They maintain a provable upper and lower bound on the (globally) optimal objective value, and they return a certificate proving that the solution found is globally optimal within a predefined margin of error. Our implementation involves the solver Gurobi, which uses a variant of the spatial BB method. Gurobi does not return a certificate of global optimality. The found optimal solution can be used to confirm Gurobi’s final primal bound. The log output contains the BB tree information, but it seems impossible to completely reconstruct Gurobi’s BB tree or view the specific added cuts.

Admittedly, using Gurobi means a certain trade-off between certifying the global optimality and efficiency of the computations. In our experiments, we recovered the same

global solutions as in [3] and scaled our algorithm to manipulators with more DOFs than other research papers. Recent results [13] indicate that considering random instances of some problems results in polynomial time complexity for BB methods.

#### E. Applicability of QCQP Lifting

The lifting algorithm presented in Section II-C applies to any IK problem involving polynomials since it needs no additional assumptions about the degree of polynomials or their properties. This enlarges the scope of applicability of our model significantly. For example, instead of the distance from preferred angles (1), we can optimize *manipulability*. We can also use the QCQP lifting for the *design of manipulators* by introducing new variables in place of the parameters of the manipulator. This is based on the observation that even if all the parameters in the D-H matrices are treated as optimization variables, the resulting IK problem would still be a POP. Since polynomials are closed under addition and multiplication, the resulting constraints could be lifted as described in Section II-C.

### III. RESULTS

We compare our QCQP-based procedure to the prior SOS-based method [3]. Our numerical experiments use an *Intel Xeon Gold 6146* processor limited to 4 threads. Implementing the SOS-based method involves Maple 2022 and Mosek 10, while our method uses Gurobi 10. These solvers are the latest and fastest for each problem class available to us at the time of writing. Results for the SOS-based method are based on the original implementation and use all the techniques described in [3] to improve numerical stability. We show the statistics of solve times estimated from randomly generated feasible and infeasible poses of the KUKA LBR iiwa 7-DOF manipulator, the *iCub* [14] humanoid robot’s right arm and torso (7–10 DOF), and on randomly generated designs with seven joints. Unless stated otherwise, the results in following sections pertain to the lifting approach  $\text{lift}_A$  (see Section II-C).

#### A. Factors Influencing the Solution Time

As we solve IK via a non-convex QCQP, it is difficult to predict the solution time of a given instance; however, in our experience, the time increases with the number of non-zeros in the constraint matrix, the number of variables, and the volume of the space that needs to be searched. In practical terms, the runtime depends on the manipulator’s range of motion and the link twists. Specifically, designs with axes offset by  $90^\circ$ , as is the case for the *iCub* and the *KUKA LBR iiwa*, are easier to solve than ones with joints at arbitrary angles.

To demonstrate how the time required to solve IK depends on the aspects of mechanical design, we will consider three sets of randomly generated robotic designs with seven joints:

- Orth. (Orthogonal) designs with link twists sampled uniformly from  $\{-\pi/2, \pi/2\}$ , resulting in orthogonal axes and a 6-radian range of motion  $-3 \leq \theta_i \leq 3$ ,

TABLE I

QCQP RUNTIME, THE POSITIONAL AND ROTATIONAL ERROR OF THE FINAL SOLUTION. “COLD” IS A COLD-START WITHOUT A LOCAL SOLUTION. THE  $\text{lift}_M$  FORMULATION SOLVED THE ORTHOGONAL INSTANCES IN 1.23 s, 4-rad IN 1.93 s, BUT 6-rad IN 8.96 s ON AVERAGE.

Set	Time (s)	Err. ( $\mu\text{m}$ )	Err. ( $\mu\text{rad}$ )
Orth.	1.31	2.29	1.28
(cold) Orth.	2.12	2.71	1.28
4-rad	2.36	2.20	1.09
(cold) 4-rad	3.90	2.39	1.26
6-rad	4.91	2.88	1.70
(cold) 6-rad	12.20	2.51	1.15




Fig. 2. A semi-log plot of solve-time kernel-density-estimates from 1000 samples of randomly generated designs for each parameter set.

6-rad designs with arbitrary link twists  $\alpha_i$  uniformly sampled from  $[-3, 3]$  and 6-radian range of motion, 4-rad designs with arbitrary twists  $\alpha_i$  uniformly sampled from  $[-3, 3]$ , and 4-radian range  $-2 \leq \theta_i \leq 2$ .

The link offsets and lengths are generated randomly according to a uniform distribution between 10 and 100 centimeters.

Finally, warm-starting can speed up the global solution search even with low-precision local solutions. We used *IPopt* [15] and the straightforward non-linear program 7 to find an initial guess:

$$\begin{aligned}
 \min_{\theta} \quad & \sum_{i=1}^n 2w_i(1 - \cos \theta_i \cos \hat{\theta}_i - \sin \theta_i \sin \hat{\theta}_i) \\
 \text{s.t.} \quad & T_1 \cdots T_\nu = PT_n^{-1} \cdots T_{\nu+1}^{-1} \\
 & \theta_i^{\min} \leq \theta_i \leq \theta_i^{\max} \quad i = 1, \dots, n
 \end{aligned} \tag{7}$$

This warm-starting may fail to provide a solution, as local solvers may get stuck at points of local infeasibility. We limited the local-solution search to 200 iterations (approximately 200 ms). Unless indicated otherwise, all of the following benchmarks used warm-starting.

Switching to modern QP or iterative real-time techniques with higher precision and lower failure rates will likely improve solution times. Additionally, such techniques could be used as runtime branch-and-bound heuristics.

Note that the variable and constraint count alone are not a good predictor of runtime. Despite not allowing for can-

TABLE II

DOF AND THE RESULTING PROBLEM SIZE AFTER PRESOLVE FOR EACH LIFTING METHOD

Set	DOF	$\text{lift}_M$		$\text{lift}_A$	
		vars.	constrs.	vars.	constrs.
KUKA	7	71	57	76	62
Orth.	7	78	64	80	66
6-rad	7	118	104	119	105
iCub	7	78	62	78	64
iCub	8	95	79	102	86
iCub	9	136	118	146	128
iCub	10	179	159	190	170




Fig. 3. A semi-log plot comparing the normalized cumulative distributions from 100 solve times for the *KUKA LBR iiwa* 7-DOF robot using our QCQP approach (Gurobi, SCIP) and the SOS-based technique.

cellation  $\text{lift}_M$  results in the fewest variables after presolve; however, while the method is faster on some designs on average, its behavior is less predictable. In extreme cases, this resulted in runtimes almost 100 times higher than the average.

### B. Comparison Across Solvers and Methods (7 DOF)

Our approach outperforms the SOS-based method [3] when tested on feasible IK instances generated by the *KUKA LBR iiwa* 7-DOF manipulator. The state-of-the-art method was explicitly optimized to take advantage of the manipulator’s structure and eliminate one of the joints. Due to Maple’s license availability, the following comparison was performed on four threads of *AMD Ryzen 5 4600G*. We did not include the modelling and Matlab overhead in the reported runtimes. The average solve-time (Figure 3) of our technique was 0.26 s, compared to 2.9 s for the optimized SOS approach. We have also tried the free and open-source solver *SCIP* instead of *Gurobi*. However, the average solve time for *SCIP* rises to 5.7 s. While a performance decrease is expected when comparing commercial solvers to free and open-source ones, in this case *SCIP* could also not effectively utilize more than one thread and did not benefit from the local solution provided by warm-starting. The prior method can take advantage of the structure of the manipulator, which significantly improves the runtime compared to generic manipulators.

To compare the methods on randomly generated designs, we rounded the lengths and offsets of the 6-rad test-set to




Fig. 4. A semi-log plot comparing the normalized cumulative distributions of solve times obtained from 100 samples of randomly generated designs with 6-radian range of motion using our approach and the SOS-based technique (rounding to two significant digits).




Fig. 5. Normalized cumulative distributions  $F(s)$  of solve times from 1000 samples of the *iCub*'s 7-DOF right arm. The full SOS-based method with symbolic reduction did not finish in time.

two significant digits, as the performance of the SOS-based method depends on the number of digits in the problem. The performance of our method is not significantly affected by the rounding. The average solve-time (Fig. 4) of our technique was 12.0s with *Gurobi* and 931s with *SCIP*. The modified SOS approach took 4328s to solve an instance on average. Finally, we note that the prior SOS-based method has roughly a 1% failure rate, a problem we did not encounter with our method.

The symbolic reduction step of the SOS-based method can be a bottleneck for some robots. For example, the 7-DOF IK instances of *iCub*'s right arm (Fig. 5) did not consistently finish within 5 hours of runtime even despite rounding to the nearest millimeter (3 significant figures). In this case however, we found that the reduction is not necessary, and *Mosek* could solve the SOS programs directly with only a 1% failure rate (this was described as the *naïve* approach in [3]). The average solve-time over 1000 samples was 0.17s for *Gurobi*, 4.5s for *SCIP*, and 3.3s for the *naïve* SOS.

Finally, we tested the SOS-based method without symbolic reduction on *Canadarm2*. While the average runtime was 6.3s, the success rate was only 31% out of 1000 feasible instances. The average runtime for QCQP with *Gurobi* was 1.8s.

### C. *iCub* Right Arm and Torso (7–10 DOF)

We also demonstrate that our technique scales to realistic instances with up to 10 degrees of freedom. Due to the low range of motion of the *iCub*'s arms, it can be beneficial to

TABLE III

THE QCQP RUNTIME STATISTICS (MEAN, FIRST AND THIRD QUARTILE) AND THE MEAN POSITIONAL AND ROTATIONAL ERRORS OF THE FINAL SOLUTION COMPUTED FROM 1000 IK INSTANCES OF *iCUB* (7–10 DOF).

Set	Time			Loc. Err. Avg. ( $\mu\text{m}$ )	Rot. Err. Avg. ( $\mu\text{rad}$ )
	Avg. (s)	Q1 (s)	Q3 (s)		
7 DOF	0.2	0.1	0.2	0.2	1.0
8 DOF	0.8	0.6	0.9	0.3	1.4
9 DOF	6.5	2.8	9.8	0.5	2.1
10 DOF	77.6	28.0	90.4	0.8	2.9

include the torso in the kinematic chain. We start with the 7-DOF right arm and progressively add torso joints up to the full 10-DOF chain. We are aware of no suitable baseline for this class of problems.

We also checked the scalability by running the 10-DOF instances on 126 threads of *AMD EPYC 7543* processor, which resulted in a mean runtime of 10.1s over 100 samples.

### D. Quality of Local Solutions

In case of the *iCub* test sets, the objective value of the local solutions found by *IPopt* was often either globally optimal already, regardless of the DOF. For poses from the 10-DOF set, *Gurobi* was able to use the local solutions or find nearby feasible solution in 93% of cases.

The median relative suboptimality of the local solutions was  $-0.07\%$  for 10 000 feasible poses of the 8-DOF *iCub* kinematic chain. Compare this to the *KUKA iiwa*, where the median relative distance was  $-14.8\%$ .

### E. Infeasible Poses

Deciding that a kinematic chain is incapable of achieving a given pose may be faster or slower on average than computing the optimal angles for a feasible pose. We show this behavior on poses generated by computing the forward kinematics of randomly sampled joint angles without considering joint limits.

Out of a set of 1000 poses generated for the *KUKA* manipulator, 345 were infeasible, and the average time to detect infeasibility was 1.1s compared to 0.44s to solve a feasible pose. The infeasibility of poses for *iCub*'s arm (7 DOF) was detected in 0.02s on average over 980 samples.

## IV. CONCLUSIONS

We have demonstrated that our QCQP-based method finds globally optimal solutions to IK problems without symbolic reduction in exact arithmetic. Our results agree with those in [3] for cases where IK instances can be set to solve the same objective. To the best of our knowledge, our technique is unique in the ability to globally solve IK instances of generic manipulators with eight or more degrees of freedom.

Although our method does not apply to real-time control, additional pre-processing and appropriate heuristics may make such a use case possible. For example, we could use modern real-time IK solvers based on convex QP as runtime




Fig. 6. Difference between the objective achieved by *IPOpt* from the actual solution over 10 000 feasible poses showing that local solutions are often nearly optimal for the *iCub* 8 DOF set, but not for *KUKA iiwa*.

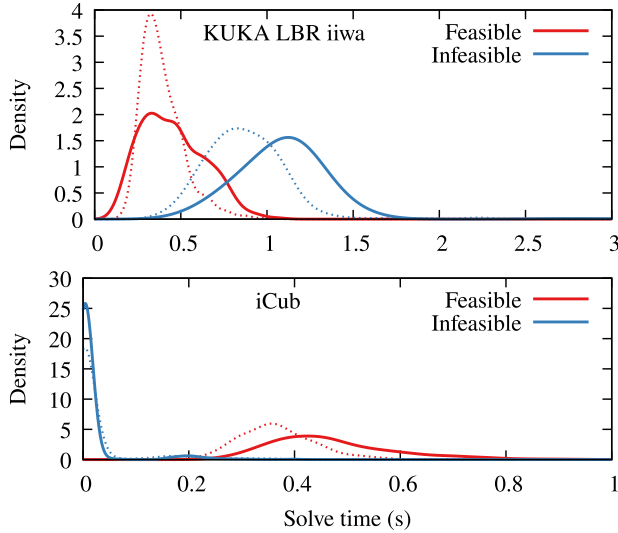


Fig. 7. Kernel density estimation of solve times of feasible/infeasible poses using  $\text{lift}_A$ . The performance of  $\text{lift}_M$  is overlaid as a dotted line.

heuristics or generate initial feasible solutions to warm-start the search. While running the experiments, we observed that a significant fraction of the solution time for *Gurobi* was spent proving that the solution was optimal. Including outer approximations in original variables as additional constraints or cuts will likely result in faster solve times. Similarly, the performance of *SCIP* will likely improve with higher-quality initial solutions. It may also be possible to accelerate the detection of infeasible poses using the outer approximation technique described by Dai *et al.* [5].

Our proof-of-concept implementation is available on GitHub at <https://github.com/votroto/IK.jl>.

#### ACKNOWLEDGMENT

We would like to thank Didier Henrion (LAAS-CNRS, University of Toulouse) and Jakub Mareček (Czech Technical University in Prague) for their valuable comments.

#### REFERENCES

- [1] D. Manocha and J. Canny, "Efficient inverse kinematics for general 6R manipulators," *IEEE Transactions on Robotics and Automation*, vol. 10, no. 5, pp. 648–657, 1994.
- [2] M. W. Spong, S. Hutchinson, and M. Vidyasagar, *Robot Modeling and Control*, 2nd ed. Standards Information Network, Feb. 2020.
- [3] P. Trutman, M. S. El Din, D. Henrion, and T. Pajdla, "Globally optimal solution to inverse kinematics of 7dof serial manipulator," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6012–6019, 2022.
- [4] S. Buss, "Introduction to inverse kinematics with Jacobian transpose, pseudoinverse and damped least squares methods," *IEEE Transactions in Robotics and Automation*, vol. 17, 05 2004.
- [5] H. Dai, G. Izatt, and R. Tedrake, "Global inverse kinematics via mixed-integer convex optimization," *Int. J. Rob. Res.*, vol. 38, no. 12-13, pp. 1420–1441, Oct. 2019.
- [6] S. S. M. Salehian, N. Figueroa, and A. Billard, "A unified framework for coordinated multi-arm motion planning," *The International Journal of Robotics Research*, vol. 37, no. 10, pp. 1205–1232, 2018. [Online]. Available: <https://doi.org/10.1177/0278364918765952>
- [7] K. Bestuzheva, M. Besançon, W.-K. Chen, A. Chmiela, T. Donkiewicz *et al.*, "The SCIP Optimization Suite 8.0," Optimization Online, Technical Report, December 2021. [Online]. Available: [http://www.optimization-online.org/DB\\_HTML/2021/12/8728.html](http://www.optimization-online.org/DB_HTML/2021/12/8728.html)
- [8] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2022. [Online]. Available: <https://www.gurobi.com>
- [9] U. Pattacini, F. Nori, L. Natale, G. Metta, and G. Sandini, "An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots," 10 2010, pp. 1668–1674.
- [10] J. B. Lasserre, "Global optimization with polynomials and the problem of moments," *SIAM Journal on Optimization*, vol. 11, no. 3, pp. 796–817, 2001. [Online]. Available: <https://doi.org/10.1137/S1052623400366802>
- [11] M. Laurent, *Sums of Squares, Moment Matrices and Optimization Over Polynomials*. New York, NY: Springer New York, 2009, pp. 157–270. [Online]. Available: [https://doi.org/10.1007/978-0-387-09686-5\\_7](https://doi.org/10.1007/978-0-387-09686-5_7)
- [12] N. Z. Shor, "An approach to obtaining global extremums in polynomial mathematical programming problems," *Cybernetics*, vol. 23, no. 5, pp. 695–700, Sep 1987. [Online]. Available: <https://doi.org/10.1007/BF01074929>
- [13] S. S. Dey, Y. Dubey, and M. Molinaro, "Branch-and-bound solves random binary IPs in polytime," in *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2021, pp. 579–591.
- [14] A. Parmiggiani, M. Maggiali, L. Natale, F. Nori, A. Schmitz, N. Tsagarakis, J. Santos-Victor, F. Becchi, G. Sandini, and G. Metta, "The design of the iCub humanoid robot," *International Journal of Humanoid Robotics*, vol. 9, 12 2012.
- [15] 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, pp. 25–57, 2006.