

Self-Contained and Automatic Calibration of a Multi-Fingered Hand Using Only Pairwise Contact Measurements

Johannes Tenhumberg^{*1,2,3} Leon Sievers^{*1,2,3} Berthold Bäuml^{1,2}

Abstract—A self-contained calibration procedure that can be performed automatically without additional external sensors or tools is a significant advantage, especially for complex robotic systems. Here, we show that the kinematics of a multi-fingered robotic hand can be precisely calibrated only by moving the tips of the fingers pairwise into contact. The only prerequisite for this is sensitive contact detection, e.g., by torque-sensing in the joints (as in our DLR-Hand II) or tactile skin. The measurement function for a given joint configuration is the distance between the modeled fingertip geometries, but the actual measurement is always zero. In an in-depth analysis, we prove that this contact-based calibration determines all quantities needed for manipulating objects with the hand, i.e., the difference vectors of the fingertips, and that it is as sensitive as a calibration using an external visual tracking system and markers. We describe the complete calibration scheme, including the selection of optimal sample joint configurations and search motions for the contacts despite the initial kinematic uncertainties. In a real-world calibration experiment for the torque-controlled four-fingered DLR-Hand II, the maximal error of 17.7 mm can be reduced to only 3.7 mm.

Web: <https://dlr-alm.github.io/2023-humanoids-contact/>

I. INTRODUCTION

Autonomous robots that robustly perform dextrous manipulation tasks in the real world generally require precise models of the system’s kinematics. Regarding multi-fingered robotic hands, one crucial class is planning algorithms for finding an optimal grasp for a given 3D model of an object [1, 2]. This task depends on the precise placement of the fingers on the object’s surface. Another class is methods for dextrous in-hand manipulation. Here, only recent modern deep reinforcement learning algorithms trained in simulation have enabled dexterity close to human performance. Primarily when performed in a purely tactile setting [3, 4, 5], i.e., without cameras, where only joint angles (and tactile measurements, e.g., via torque-sensing) are used, robust zero-shot sim2real transfer requires a precise kinematics model with a maximal error of a few millimeters.

It is desirable to have a quick calibration procedure that runs entirely automatically and is self-contained, i.e., it does not need any additional external sensors or tools to obtain a precise kinematic model.

A. Related Work

A classical and often-used approach for calibrating kinematic chains and trees is visual tracking of the end-




Fig. 1. DLR-Hand II [6] with thumb and index finger in contact for different poses. The kinematic tree of the whole hand is indicated in orange, and the three active joints plus the fourth passive joint are drawn together with the dimensions of the fingers. The other two fingers are far extended to allow for a large shared workspace of the current pair. For the whole calibration, all six finger pairs are measured and calibrated jointly.

effector(s). For a robotic hand, this visual approach with an external tracking system was demonstrated by Lee et al. [7]. However, adding visual markers to the many end-effectors is pretty cumbersome, especially as more than one marker is usually required per fingertip due to the mutual occlusions of the many fingers in the small workspace of a hand. Using an electromagnetic tracking system to measure the position and orientation of the fingertips [8] is less prone to occlusions. However, the requirement for additional hardware hinders the ease of use of such an approach.

Another calibration procedure that avoids the need for a tracking system uses geometric constraints on the kinematic chain. An early work constrained the motion of an end-effector on a plane to calibrate a robotic arm [9]. Since then, there have been examples of using mechanical fixtures [10], relative calibration techniques [11], and precise reference plates [12] to calibrate a robotic arm without a vision system. Bennett and Hollerbach [13] applied these ideas to a robotic hand and calibrated the multi-finger Utah-MIT hand by rigidly connecting two fingertips with a plate, resulting in a closed-loop kinematic chain. A downside of all those approaches is that the mounting procedure can damage the fingers and takes time, especially if it needs to be repeated for each finger pair.

A particular case of the relative geometric calibration

* First two authors contributed equally.

¹DLR Institute of Robotics & Mechatronics, Germany

²Deggendorf Institute of Technology, Germany

³Technical University of Munich, Germany

Contact: johannes.tenhumberg@dlr.de, leon.sievers@dlr.de

techniques does not require mechanical fixtures to enforce the constraints but relies on self-contact. The humanoid robot iCub was intensely used for such studies. First in simulation [14] and then on the actual hardware together with other sensing modes [15]. Roncone et al. [16] calibrated the full DH parameters of the iCub robot using self-touch with a tactile skin with 4200 sensing points. Besides identifying the parameters of a given kinematic chain, self-contact can also be used to incrementally update the kinematic scheme [17] or identify the layout of an artificial tactile skin [18].

B. Contributions

In prior work, an external camera system was first used to calibrate an elastic kinematic model of the humanoid Agile Justin [19]. Then, this approach became more accessible by using the robot’s internal RGB camera to calibrate the complete kinematic tree [20]. This paper extends the calibration to the DLR-Hand II [6, 21] (see Fig. 1) of the robot and introduces a contact-based approach that does not rely on anything other than the kinematic tree structure itself. The main contributions of this work are:

- We show, to our knowledge for the first time, that the kinematics of a multi-fingered hand can be calibrated using only contact measurements of (all) finger pairs, i.e., without any external tools like cameras or markers.
- An in-depth theoretical analysis of the calibration problem with redundant parameters, including proof that the quantities relevant for manipulation tasks, i.e., the difference vectors between the fingertips, are entirely determined by our contact measurement scheme.
- A comparison of calibration by contact measurements with calibration using a visual tracking system (theoretical analysis and simulation experiments).
- A complete scheme for executing a contact-based calibration for a general multi-fingered hand, including optimal selection of finger configurations and planning of search motions to deal with the initial uncertainties in the kinematics.
- Real-world experiments with the contact calibration of the DLR-Hand II resulting in a reduction of the maximal error from 17.7 mm for the nominal kinematic to only 3.7 mm for the full DH parameters calibration.

II. ROBOT MODEL

The forward kinematics of a robot maps from the configuration space of generalized joint angles to the robots’ physical pose in the cartesian workspace. Both for robotic arms and hands, an accurate model of the robot’s kinematics is essential for precise and collision-free motions. A well-fitted model is also required to perform challenging grasping and manipulation tasks.

DH parameters are widely used to describe this kinematic model of the robot. In this formulation, four values $\rho_i = [d_i, r_i, \alpha_i, \theta_i]$ describe the connection between two consecutive frames of the robot:

$${}^{i-1}T_i = \text{Rot}_x(\alpha_i) \cdot \text{Trans}_x(r_i) \cdot \text{Rot}_z(\theta_i) \cdot \text{Trans}_z(d_i) \quad (1)$$




Fig. 2. The scheme shows the actual contact measurement on the hardware (black) and the same joint configuration applied to the robot model (gray). The fingers are in penetration for the current set of calibration parameters. This error (red) should be minimized over the calibration process.

The joints q_i are treated as offsets to θ_i or d_i in (1) depending on the type of joint. This minimal representation with two translational and two rotational parameters is enough to describe an arbitrary robot.

The frame of the end-effector ($i = E$) relative to the robot’s base ($i = 0$) is calculated by applying the transformations in series:

$${}^0T_E = {}^0T_1 \cdot {}^1T_2 \cdot \dots \cdot {}^{E-1}T_E \quad (2)$$

For robots with a kinematic tree structure with multiple end-effectors $E_k, k = 1 \dots N_E$, equation (2) holds for each branch. The forward kinematics maps from joint configurations $q \in \mathcal{Q}$ to all the N_F frames of the robot

$$f(q, \rho) = F = [{}^0T_1, {}^0T_2, \dots, {}^0T_{N_F}]. \quad (3)$$

Each frame F_i describes a full 6D pose $(F_{i,x}, F_{i,r}) \in \mathbb{R}^3 \times SO(3)$.

III. CALIBRATION FOR IN-HAND MANIPULATION

A. Identification

The goal of the calibration is to identify this kinematic model. In our case, the central parameters of the forward model are the DH parameters ρ . Nevertheless, in general, there can be additional parameters that have to be estimated jointly with the kinematic parameters. Examples are camera intrinsics or additional frames to close the measurement loop. We denote the combination of all calibration parameters Θ . These parameters of the measurement function h can be identified using a dataset $D = \{(q^{(n)}, y^{(n)})\}_{n=1 \dots N}$ of corresponding pairs between the robot’s configuration $q^{(n)}$ and the measurement $y^{(n)}$. We formulate the identification as a single combined least-squares problem based on all measurements and all body parts to minimize the error $\epsilon = y - h(q, \Theta)$ between measurements and model predictions. To find the optimal parameters Θ^* , we use the maximum a posteriori (MAP) approach:

$$\min_{\Theta} \left[\sum_n \frac{1}{\sigma_m^2} \left\| y^{(n)} - h(q^{(n)}, \Theta) \right\|^2 + \Delta \Theta^T \Lambda_p^{-1} \Delta \Theta \right] \\ \text{with } \Delta \Theta = \Theta - \Theta_p$$

This approach employs a Gaussian distribution beforehand, characterized by a mean Θ_p and a diagonal covariance matrix $\Lambda_p = \text{diag } \sigma_p^2$. The uncertainty of the calibration parameters is modeled via σ_p . Furthermore, we assume the usual Gaussian distribution with a mean of zero and a standard deviation of σ_m for the measurement noise. By




Fig. 3. The scheme shows the three measurement functions discussed in Section III. *Top (Section III-B)*: The task measurement function h_t measures the relative positions of the end-effectors. *Middle (Section III-D)*: The contact measurement function h_c measures the scalar distance between two end-effectors. *Bottom (Section III-C)*: The cartesian measurement function h_v uses an external tracking system to measure the absolute position of the end-effectors.

incorporating the prior, this method is regularized, ensuring a minimum exists, even if there are redundancies in the measurement or kinematic model.

B. Task Measurement Function

We aim to calibrate this kinematic model of the hand for dextrous in-hand manipulation. To move the fingers in a controlled manner, as demonstrated by Pitz et al. [4] with the cube, an accurate model for the relative positions of the fingertips to each other is necessary. To relate closely to this task, we chose our desired measurement function to measure the relative positions directly. The calibration process should then minimize the error of this function.

For N_E fingers, $N_E - 1$ distance vectors define the whole set. One can choose one fingertip E_1 as the basis and compute the positions relative to it for the remaining fingers

$$y^k = h_t^k(q, \rho) = f(q, \rho)_{E_k, x} - f(q, \rho)_{E_1, x} \quad (4)$$

One can then concatenate the measurements for the pairs to get the combined measurements $h_t = [h_t^2, \dots, h_t^{N_E}]$, $h_t^k \in \mathbb{R}^3$. This results in $3 \cdot (N_E - 1)$ data points per robot pose for the three spatial directions and N_E end-effectors.

Note that (4) is our theoretically desired measurement function, which measures all the information we care about for precise in-hand manipulation. For example, we can not detect a translational offset of all end-effectors like a cartesian tracking system could. However, such a shift does not change the fingers' relative behavior and is irrelevant to our task.

Task Test Set: Besides the task measurement function, the test set is central to evaluating the calibration quality. We want a high accuracy across the whole cartesian workspace. Therefore, the test set for evaluating all our experiments should be uniformly distributed in the cartesian workspace. We sample random joint configurations for each finger and map them to the end-effector positions via the forward kinematics f . In the cartesian workspace, we create a fine

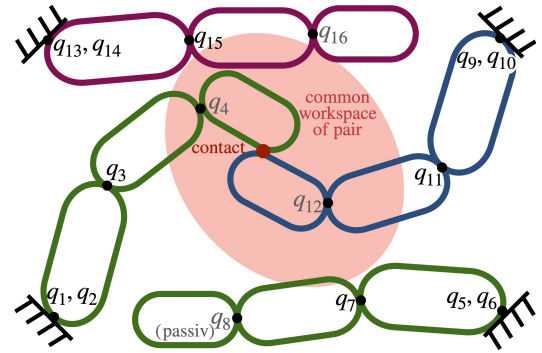


Fig. 4. The graphic shows the contact calibration approach for a four-fingered hand. For pairwise contact, two fingers always need to move out of the shared workspace of the current finger pair to ensure that self-collisions are avoided and the available configuration space for contact detection is well used.

grid and draw one of the grid cells uniformly, and in the second step, we draw one joint configuration that lies in this grid cell. This sampling strategy ensures a uniform distribution over the workspace, which does not hold for configurations sampled randomly in the configuration space.

After defining a suitable measurement function and test set for our task, we discuss the actual measurement methods that can be applied to the hardware. First, an external camera system that tracks the cartesian position of visual markers followed by our pairwise contact measurements. Note that there are also additional considerations, especially the independence of particular infrastructure and the low complexity of the measurement setup, speaking in favor of contact-based measurements.

C. Cartesian Measurement Function

One option to perform the actual measurements is to mount visual markers on the kinematic tree and use an external tracking system to collect cartesian measurements of these markers. Assuming that the markers are fixed at position m_{E_k} relative to end-effector E_k , the cartesian measurement function

$$y^k = h_v(q, \Theta)^k = {}^cT_0 \cdot f(q, \Theta)_{E_k} \cdot m_{E_k} \quad (5)$$

describes how a marker moves dependent on the joint configuration q and the calibration parameters Θ . Suppose one wants to calibrate the forward kinematics jointly for multiple end-effectors. In that case, one can combine the measurements for each of the N_E markers to a combined measurement function $h_v = [h_v^1, \dots, h_v^{N_E}]$, $h_v^k \in \mathbb{R}^3$. While, in general, the measurements with an external camera can be made without constraining the robot directly, one still needs to account for self-collision and a clear view of the markers. These additional constraints reduce the possible space in which measurements can be collected.

D. Contact Measurement Function

Another option to perform the actual measurement is to use contact information. We describe this procedure in more detail in Section V. In general, the corresponding




Fig. 5. This figure shows the ordered eigenvalues for different measurement setups to analyze the sensitivity. The evaluation was done with the nominal kinematic of the DLR-Hand II; for a more generic hand, see the right plot. The task measurement function h_t is blue, and our contact measurement function h_c is red. Furthermore, we show three modes. For the one where all the pairs are calibrated simultaneously (\bullet), the kernels of both measurement functions have the same size. The same is true for the calibration with three fingers (\blacktriangle). However, the kernel sizes differ when just a single pair (\bullet) is considered. The light gray vertical lines indicate the maximal number of parameters for each mode.

measurement function measures the distance between parts of the robot. One needs the exact geometry of the two bodies to compute the distance d^u between a pair of bodies $u = (E_k, E_l)$ on the kinematic tree. The distance is only dependent on the relative position and orientation of the bodies

$$d^u = d^u(E_k T_{E_l}) = d^u(f(q, \rho)_{E_k}^{-1} \cdot f(q, \rho)_{E_l}) = d^u(q, \rho). \quad (6)$$

This relative frame $E_k T_{E_l}$ is directly computable from the forward kinematics $f(q, \rho)$.

If the bodies at E_k and E_l have simple geometric forms, one can directly compute the distance d_m . In the case of the DLR-Hand II, the fingertips are perfect capsules in the contact area and, therefore, straightforward to compute. The more general case is that the geometries are given as arbitrary meshes. In this case, one has to use, for example, algorithms like GJK to compute the distance.

The contact measurement function for the pair u can now be written as

$$y^u = h_c(q, \Theta)^u = d^u(f(q, \rho)). \quad (7)$$

This function can only measure the scalar distance between two body pairs. With N_E end-effectors on the kinematic tree, there are in total $N_U = \binom{N_E}{2}$ pairs. The combined measurement function is in this case $h_c = [h_c^1, \dots, h_c^{N_U}]$, $h_c^k \in \mathbb{R}$

The contact measurement adds hard constraints to the data collection (see Section V-A). As only configurations in contact can be measured, the available configuration space is drastically reduced. Therefore, the remaining subspace can be quite different from the cartesian task test set described in Section III-B. The following section discusses how those different measurement functions relate and how many parameters they can identify. Furthermore, we discuss the influence of the different calibration and test sets on the sensitivity.




Fig. 6. The plot is structured equivalently to Fig. 5 but for a more generic hand kinematic without parallel axes or mounting frames. Here, in theory, all 64 parameters for the entire hand (\bullet) can be identified. However, when only two fingers are considered (\bullet), the contact measurement function h_c is still less sensitive than the task measurement function h_t . This speaks in favor of a holistic calibration of the whole kinematic tree.

IV. PROBLEM ANALYSIS

A. Sensitivity Analysis

A single contact measurement (7) measures less information than a cartesian tracking system (5). The tracking system can measure the absolute position of the markers in the workspace. The contact can only measure the relative distance between two parts of the kinematic chain. In other words, the contact can measure the joint angles where the distance between the two bodies is zero.

The question is whether this reduced information in the measurement function leads to non-identifiable parameters. We can answer this question directly by looking at the jacobians of the different measurement functions

$$J^s = \frac{\partial h(q^s, \Theta)}{\partial \Theta} \Big|_{\Theta_0}. \quad (8)$$

Assuming that each measurement is d -dimensional, concatenating those matrices for each measurement leads to the combined jacobian $\mathbf{J} = [J^1, \dots, J^{N_D}]$ with dimensions $(N_D \cdot d) \times N_\Theta$. We can investigate which parameters are identifiable by the measurement function h by looking at the nullspace of $\mathbf{J}^T \mathbf{J}$. The size of the nullspace marks how many of the model parameters Θ can not be identified. Furthermore, from the eigenvectors corresponding to the eigenvalues close to zero, one can identify the parameters (or sets of parameters) which can not be measured.

When one deals with two distinct measurement functions, as we described with the desired task h_t and our actual contact measurement function h_c , the necessary condition is that

$$\text{kernel}(\mathbf{J}_c^T \mathbf{J}_c) \subseteq \text{kernel}(\mathbf{J}_t^T \mathbf{J}_t). \quad (9)$$

If this condition is satisfied, one can identify all the relevant parameters to the task, defined by h_t . Note that this is less strict than demanding that both kernels are zero and applies in general to distinct measurement functions for calibration and evaluation. We allow for unidentifiable parameters if they do not influence our desired measurement function.

B. Optimal Experimental Design

Following Carrillo et al. [22], we use task D-optimality to select appropriate samples for measuring. However, we have

two key differences in our setup. First, we have a theoretical desired task measurement function h_t and an actual measurement function h_c we can apply on the hardware. Ultimately, we want a good fit for the desired measurement function. Second, the contact measurement reduces the configuration space quite drastically. Still, we want a good fit and high accuracy across the whole workspace. Therefore we also have two sets here. The desired test set is uniformly sampled across the whole cartesian test set and one actual test set, which can be measured via contact detection. We generalize the optimality framework to account for those mismatches between measurement functions and distributions. Carrillo et al. [22] showed that the central equation decouples and the task D-optimality can efficiently be computed with

$$O_D = \frac{1}{l} \det \left(\text{cov}(\Theta) \sum_{s=1}^{N_D} J_t^{sT} J_t^s \right). \quad (10)$$

The mean over $J_i^T J_i$ is constant for a given test set of size N_D and a desired task measurement function h_t . The covariance over the calibration measurements can be estimated using the actual chosen measurement function and the specific calibration set. Let $S = \{s_i\}_{i=1}^{N_D}$ be a subset of a larger calibration set and $\mathbf{J}_c = [J_c^{s_1}, \dots, J_c^{s_{N_D}}]$ the combined jacobians corresponding to those measurements. Then the covariance is given by

$$\text{cov}(\Theta) = \mathbf{J}_c^T \text{diag}(\sigma_m^2) \mathbf{J}_c + \text{diag}(\sigma_p^2) \quad (11)$$

and transforms the uncertainty in the measurements σ_m and the priors σ_p into an uncertainty in the parameters.

(10) lets us compute how well different calibration sets are suited to minimize the error of the desired measurement function over a desired test set. This criterion can be used to choose a good set of suitable poses for measuring. Besides reducing the overall size of the necessary calibration set, this selection criterion also counteracts the mismatch in the measurement functions and the calibration and test sets. We report the results of the sensitivity analysis and the optimal selection criterion in Section VI-B

V. CONTACT MEASUREMENT PROCEDURE

A. Sample Generation

Our approach is to collect pairwise contact measurements for the tree structure of the hand. A key difference for contact-based calibration is that one does not know the exact joint configuration, which will be measured beforehand. Generally, the measurement process collects pairs (q, y) to calibrate a function $g(q, \theta) = y$. For vision-based measurements, one collects the cartesian position y_i for selected joints q_i . On the other hand, for contact-based measurements, the y_i is known a priori; the contact is, by definition, $y_i = 0$. The contact measurement delivers the exact configuration q_i , which leads to contact.

Therefore, contact measurement must also include a search to find the exact configuration in which the contact happens. We tackle this problem by generating for each measurement point a trajectory along which contact will probably happen.




Fig. 7. Red: L^2 norm of $\Delta\tau$. $\Delta\tau$ denotes the difference between the torque signal measured by the passive finger before the approach started, τ_0 , and the currently measured torque. Blue: L^2 norm of Δq . Δq denotes the difference between the position signal measured by the passive finger before the approach started and the currently measured position. The contact is detected when the torque threshold $\tau_t = 0.1$ Nm is exceeded. Note how the passive finger's joint angles can change after the active finger starts to drive causing vibrations in the system. The torque sensors only measure noise until the contact occurs.

We define this path by its endpoints. The start configuration is far from contact with the nominal robot model, and the end configuration is deep in penetration. Between those endpoints, we then detect contact for the specific finger pair.

Our approach to generating these search trajectories for the contact measurements consists of multiple steps. The first step is to find the volume in the workspace which both end-effectors can reach. We save a large ($n=100000$) set of configurations that fall in this intersection for both fingers. The next step is to randomly choose one configuration for finger A and check for which configurations of finger B the tips collide. The threshold for collision one chooses here strongly depends on the robot model's initial uncertainty. Continuing from this pair of configurations, we choose which finger should be static and which should move into contact. For the moving finger, we then sample an additional random configuration as the start point of the measurement drive.

We repeat this procedure for all six finger pairs. For each pair, the rest of the kinematic structure should move as far away as possible from the combined workspace of the current end-effector pair (see Fig. 4). The goal is to obstruct the measurements as little as possible and allow for diverse joint configurations in this constrained setting.

One additional problem is the small form factor of a robotic hand, especially compared to the errors of the uncalibrated system. For a robotic arm with a total reach of one meter, an error of a few centimeters does not change the measurement setup. However, in our case, the DLR-Hand II with its four fingers, the ratio between error in the forward kinematics and the robot's actual size is much larger. We measured uncalibrated errors up to 17.7 mm. That equals roughly 10% of the workspace and is also about the fingertip size. The consequence is that even if, in simulation for the nominal kinematic, the two fingers touch close to the center, it is still possible that the measurement fails on the hardware. Therefore one needs to account for this high uncertainty while generating samples and safe trajectories to collect those samples. Furthermore, robust contact detection is crucial for reliable measurements on the actual hardware.




Fig. 8. Convergence of mean and maximal cartesian error ϵ_v on the uniform cartesian test set for different calibration sets for the contact measurement function h_c . Randomly chosen sets (red) converge slower than the sets designed according to task D-optimality (10). The greedy strategy is green, and the detmax[23] strategy is blue. The distribution mismatch between calibration and test set can explain why the gap between random and selected samples remains significant even for larger calibration sets.

B. Contact Detection

It is essential to detect the precise joint angles in which the contact force between the fingertips is as small as possible. The DLR-Hand II is equipped with an output side torque sensor for each of the twelve active degrees of freedom. Before the passive finger gets approached, the offset, τ_0 , is set to the measured torque, $\tau_m(t)$. When the active finger moves and the change in torque is larger than the threshold τ_t , a contact is detected (see Fig. 7).

VI. EXPERIMENTS

As the model of the forward kinematics for the DLR-Hand II, we include all four DH parameters per joint, resulting in $N_\Theta = N_{\text{DoF}} \times 4$ calibration parameters. The hand has four fingers. Each finger has three active and one passive joint. This results in $16(= 4 \times 4)$ DH parameters per finger and $64(= 4 \times 16)$ parameters for the whole hand. Furthermore, each finger has three parallel joints, and the ring and middle finger are mounted with the same orientation.

A. Sensitivity Analysis

To answer the question of whether our contact measurement function h_c can identify all the parameters relevant for the task measurement function h_t , we compute the nullspace of $\mathbf{J}^T \mathbf{J}$ as described in Section IV-A. We evaluate the respective Jacobians \mathbf{J}_c and \mathbf{J}_t for the nominal robot kinematic. For the actual measurement function h_c , we use 100 configurations per pair, which were generated as described in Section V-A. For the task measurement function h_t , we use 100 configurations per finger drawn from the cartesian test set.

Fig. 5 shows the results of this analysis for the DLR-Hand II. The task measurement function h_t is drawn in blue, and our contact measurement function h_c is in red. All the pairs are calibrated simultaneously for the rightmost mode (●). Both measurement functions have 56 eigenvalues larger than $1 \cdot 10^{-6}$. Analyzing the eigenvectors confirms that the kernel of our contact measurement function h_c is wholly included in the task measurement function h_t kernel. Therefore, we can identify all parameters relevant to the task. The parallel axis of the fingers can explain the $8(= 2 \times 4)$




Fig. 9. The plot is designed equivalently to Fig. 8, but for the cartesian measurement function h_v . The overall convergence of the error ϵ_v is quicker because the absolute cartesian position yields more information per sample. Furthermore, the difference between the random approach and the dedicated selection criteria is smaller in this setting. One reason is that the calibration set for the external tracking system is closer to the task test set.

unidentifiable parameters. Each finger has 3 parallel axes along which a shift can be adjusted without influencing the end-effector, resulting in an increased size of nullspace by 2 per finger and 14 identifiable parameters per finger.

The leftmost mode (●) shows the calibration of just a single pair. Considering the calibration of one pair, from the $32(= 2 \times 16)$ total parameters, $28(= 2 \times 14)$ should be identifiable. However, the critical insight is that measuring between two chains yields less information than measuring between three or more chains. This holds for the actual scalar distance measurement (red / h_c), where 26 parameters are identifiable, and the task measurement of the relative positions (blue / h_t), where one can identify 27 parameters. However, one can measure even less with the scalar distance function when restricting the measurements to a single pair. An intuition is that two end-effectors can move on a sphere around each other without changing the scalar distance between them. Therefore, in this case, one can not identify all the parameters relevant to the task by pure contact measurements.

This invariance generally resolves when adding a third chain (▲) to the picture, favoring a holistic calibration of the full kinematic tree. Fig. 5 shows the same analysis but for a generic four-fingered hand without parallel axes or mounting frames. In theory, all parameters can be identified for the entire hand. However, in praxis, eigenvalues below 10^{-6} still indicate poor sensitivity.

B. Optimal Experimental Design

We conducted an extensive simulation study to verify our analysis in Section IV. We use the DH formalism to parametrize the forward kinematics as described in Section II and apply noise on the nominal DH parameters to obtain new robot models. For the rotational DH parameters α and θ we applied sampled uniformly $\pm 5^\circ$ and for the translational parameters DH d and r we used uniform noise ± 5 mm.

In this fashion, we created 100 different kinematics to ensure a broad distribution of models. Next, we simulated the data collection step and collected measurements for the actual contact measurement function h_c , the actual cartesian measurement function h_v , and the task measurement function h_t . On average, the models deviated 21mm from the nominal kinematic on the uniform cartesian test set.

TABLE I
CALIBRATION RESULTS IN MM

	Actual Meas. Fun.			Task Meas. Fun.
	mean	std	max	mean
nominal	6.07	3.90	17.70	8.01
calibrated joint offsets	1.04	0.82	5.13	1.36
calibrated full DH	0.72	0.58	3.69	0.89




Fig. 10. Comparison of contact error ϵ_c distribution for the different calibration models. The calibration with all DH parameters can reduce the maximal error to 3.7 mm (see also Table I)

Fig. 8 show the results of the optimal sample selection introduced in Section IV-B. The error on the cartesian test set is drawn over the calibration set size for different selection strategies. Randomly chosen sets (red) converge slower than the sets designed according to task D-optimality (10). We compare a greedy strategy (green), which at each step adds the sample s_i , which improves the task D-optimality most against the DETMAX algorithm [23] (blue). This procedure tries to swap samples in an existing calibration set to improve the task D-optimality. Both selection strategies outperform the random approach, significantly reducing the mean error to 0.1mm with a set of 300 measurements.

The distribution of configurations for the contact measurements differs from the uniform distribution in the cartesian workspace on which we evaluate the calibrations. This distribution mismatch can explain why the gap between random and selected samples remains significant even for larger calibration sets. Our approach using the task D-optimality as a selection criterion does account for this shift directly and improves the contact calibration with its strict constraint for data generation significantly.

Fig. 9 shows the same analysis but for the cartesian measurement function h_v with an external tracking system. The overall convergence is quicker because the absolute cartesian position yields more information per sample. Furthermore, the difference between the random approach and the dedicated selection criteria is smaller in this setting. One reason is that the calibration set for the external tracking system is closer to the task test set. While proving overall that the selection over task D-optimality is suited to select good calibration sets, those results show that this approach is particularly viable when there is a substantial mismatch between the task measurement function and its test set versus the actual measurement function and its corresponding calibration set.

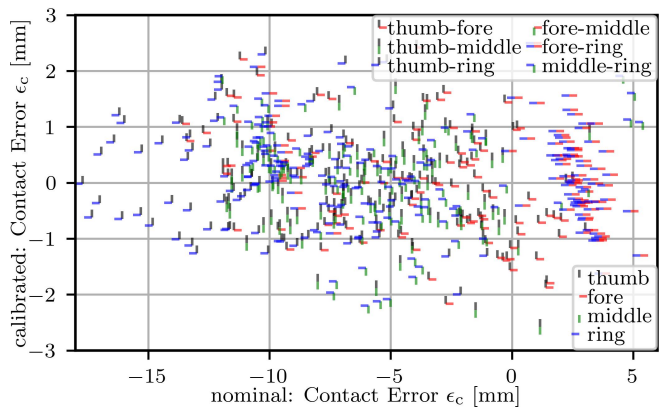


Fig. 11. This plot analyses the contact errors ϵ_c of the individual fingers and pairs before and after the calibration. Each finger has an individual marker defined by a color and an orientation. A measurement of a specific pair is then an overlay of those two finger markers. The x-axis shows the signed distance error before and the y-axis after the calibration, giving detailed insight into the error distribution of the hand.

C. Calibration of the real DLR-Hand II

Following our methodology described in Section IV and Section V-A, we calibrated the DLR-Hand II via pairwise contact measurements. Overall, we collected 300 samples to ensure we have a large enough set for calibration and evaluation. The training-test split was 80/20, and we used cross-validation to get the distribution over the whole dataset. Our findings revealed that 150 samples are sufficient for an accurate calibration. Collecting that data takes 9 minutes, making this automatic calibration procedure easy and quick to use.

The uncalibrated nominal forwards kinematic has a mean error of 6 mm over all 300 samples and a maximal error of up to 17.7 mm for the scalar distance measurement. The full error distribution over the signed distance function is shown in red in Fig. 10. Light blue is the calibrated model using only the joint offsets as calibration parameters, reducing the maximal error to 5.1 mm. The entire calibration model, including all DH parameters, is dark blue. This model reduces the maximal error further to 3.7 mm.

Fig. 11 analyses the errors of the individual fingers before and after the calibration. Each finger has an individual marker, and a measurement of a specific pair overlaps those two finger markers. The x-axis shows the signed distance error before and the y-axis after the calibration. Besides the significant overall error reduction, one can see that different finger pairs have other error distributions before and after the calibration. For example, the pair fore-ring is often slightly apart for the nominal model, while the thumb-ring pair is often in deep penetration.

Table I shows the detailed results of the different models. Additionally, the errors in the task space are given. After calibration, one can use (11) to transform the resulting calibration errors of the contact measurement function h_c into the errors of the task measurement function h_t . Starting from the calibration errors σ_m one can first apply \mathbf{J}_c to map to the uncertainties of the calibration parameters $\text{cov } \Theta$ and




Fig. 12. The DLR-Hand II in a finger contact pose from the testset *a*) with the measured joint angles q mirrored to a nominal *b*) and a calibrated *c*) model. The model's error is visibly reduced from a distance of 8.3 mm to a penetration of 1 mm, allowing dextrous in-hand manipulation.

then apply in a second step J_t to map to the task space. This results in a reduction of the mean error in the task space from 8 mm for the uncalibrated model to 0.9 mm after calibration.

Finally, Fig. 12 shows the improvement through our calibration procedure on the real DLR-Hand II. When mirroring a joint configuration q , which is in contact, onto the nominal and the calibrated robot model, one can see the sizeable uncalibrated error and the good fit after the calibration, enabling dextrous grasping and in-hand manipulation.

VII. CONCLUSIONS AND FUTURE WORK

Using a pairwise contact measurement approach, we calibrated the complex robotic DLR-Hand II with 12 active and 4 passive joints. This calibration approach has minimal requirements on the robotic hardware and needs no additional tools, making it easy to apply in different setups. One only needs a method to contact contacts (e.g., torque sensors in our case) and a mathematical model to describe the distance between the body pairs.

From an uncalibrated distance error of up to 17.7 mm (roughly equivalent to 10% of the overall size of the workspace), we could reduce maximal error to 3.7 mm and the average error to 0.7 mm. Looking at the desired task measurement function relevant to dextrous in-hand manipulation, we made a sensitivity analysis to confirm that the scalar distance information between individual finger pairs is enough to identify all relevant DH parameters. Furthermore, we showed that performing a joint calibration of multiple finger pairs yields more information than looking at just a single pair. We used task D-optimality to counteract the mismatch between our desired task measurement function and its corresponding cartesian test set versus the less informative, more constrained contact measurement approach. An exhaustive simulation study verifies this selection approach's effectiveness in balancing the calibration and the desired task.

In future work, we want to extend the calibration approach to the robot structure's elasticities, especially in the drivetrain and the fingertips. For this, forces are applied via the torque-controlled joints while in contact.

REFERENCES

- [1] D. Winkelbauer *et al.*, "A Two-stage Learning Architecture that Generates High-Quality Grasps for a Multi-Fingered Hand," in *International Conference on Intelligent Robots and Systems*, 2022.
- [2] —, "A two-stage learning architecture that generates high-quality grasps for a multi-fingered hand," in *Proc. Int. Conf. Intelligent Robots and Systems*, 2022.
- [3] L. Sievers, J. Pitz, and B. Bäuml, "Learning Purely Tactile In-Hand Manipulation with a Torque-Controlled Hand," in *International Conference on Robotics and Automation*, 2022.
- [4] J. Pitz, L. Röstel, L. Sievers, and B. Bäuml, "Dextrous Tactile In-Hand Manipulation Using a Modular Reinforcement Learning Architecture," in *International Conference on Robotics and Automation*, 2023.
- [5] L. Röstel, J. Pitz, L. Sievers, and B. Bäuml, "Estimator-coupled reinforcement learning for robust purely tactile in-hand manipulation," in *Proc. IEEE-RAS International Conference on Humanoid Robots*, 2023.
- [6] J. Butterfaß, M. Grebenstein, H. Liu, and G. Hirzinger, "DLR-Hand II: Next generation of a dextrous robot hand," in *International Conference on Robotics and Automation*, 2001.
- [7] S.-M. Lee, K.-D. Lee, S.-H. Jung, and T.-S. Noh, "Kinematic Calibration System of Robot Hands Using Vision Cameras," in *International Conference on Ubiquitous Robots and Ambient Intelligence*, 2013.
- [8] N. Tan, X. Gu, and H. Ren, "Simultaneous Robot-World, Sensor-Tip, and Kinematics Calibration of an Underactuated Robotic Hand With Soft Fingers," *IEEE Access*, 2018.
- [9] M. Ikits and J. Hollerbach, "Kinematic calibration using a plane constraint," in *International Conference on Robotics and Automation*, 1997.
- [10] M. Meggiolaro, G. Scriffignano, and S. Dubowsky, "Manipulator Calibration Using A Single Endpoint Contact Constraint," in *ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, 2000.
- [11] Y. Sun, D. J. Giblin, and K. Kazerooni, "Accurate Robotic Belt Grinding of Workpieces with Complex Geometries using Relative Calibration Techniques," *Robotics and Computer-Integrated Manufacturing*, 2009.
- [12] A. Joubair and I. A. Bonev, "Kinematic Calibration of a Six-Axis Serial Robot Using Distance and Sphere Constraints," *The International Journal of Advanced Manufacturing Technology*, 2015.
- [13] D. J. Bennett and J. M. Hollerbach, "Closed-loop Kinematic Calibration of the Utah-MIT Hand," *Experimental Robotics I*, 1990.
- [14] K. Stepanova, T. Pajdla, and M. Hoffmann, "Robot Self-Calibration Using Multiple Kinematic Chains—A Simulation Study on the iCub Humanoid Robot," *IEEE Robotics and Automation Letters*, 2019.
- [15] K. Stepanova *et al.*, "Automatic Self-Contained Calibration of an Industrial Dual-Arm Robot with Cameras Using Self-Contact, Planar Constraints, and Self-Observation," *Robotics and Computer-Integrated Manufacturing*, 2022.
- [16] A. Roncone, M. Hoffmann, U. Pattacini, and G. Metta, "Automatic Kinematic Chain Calibration Using Artificial Skin: Self-touch in the iCub Humanoid Robot," in *International Conference on Robotics and Automation*, 2014.
- [17] R. Zenha, P. Vicente, L. Jamone, and A. Bernardino, "Incremental Adaptation of a Robot Body Schema Based on Touch Events," in *International Conference on Development and Learning and Epigenetic Robotics*, 2018.
- [18] L. Rustler *et al.*, "Spatial Calibration of Whole-Body Artificial Skin on a Humanoid Robot: Comparing Self-Contact, 3D Reconstruction, and CAD-Based Calibration," in *International Conference on Humanoid Robots*, 2021.
- [19] J. Tenhumberg and B. Bäuml, "Calibration of an Elastic Humanoid Upper Body and Efficient Compensation for Motion Planning," in *International Conference on Humanoid Robots*, 2021.
- [20] J. Tenhumberg, D. Winkelbauer, D. Burschka, and B. Bäuml, "Self-Contained Calibration of an Elastic Humanoid Upper Body Using Only a Head-Mounted RGB Camera," in *International Conference on Humanoid Robots*, 2022.
- [21] Haiying Hu *et al.*, "Calibrating Human Hand for Teleoperating the HIT/DLR hand," in *International Conference on Robotics and Automation*, 2004.
- [22] H. Carrillo *et al.*, "On task-oriented criteria for configurations selection in robot calibration," in *International Conference on Robotics and Automation*, 2013.
- [23] T. J. Mitchell, "An Algorithm for the Construction of "D-Optimal" Experimental Designs," *Technometrics*, 2000.