Calibration of Error Distributions in Robot Kinematics for Increased Precision in Manipulation Tasks

Gerstewitz, Tim and Lehner, Peter and Burkhard, Lukas and Topalovic, Natalija and Albu-Schaeffer, Alin and Leidner, Daniel and Roa, Maximo A.

Citation Notice

```
@ARTICLE{11488285,
      author={Gerstewitz, Tim and Lehner, Peter and Burkhard, Lukas and Topalovic, Natalija and Albu-Schaeffer, Alin and Leidner, Daniel and Roa, Maximo A.},
      journal={IEEE Robotics and Automation Letters},
      title={Calibration of Error Distributions in Robot Kinematics for Increased Precision in Manipulation Tasks},
      year={2026},
      volume={11},
      number={6},
      pages={7174-7181},
      doi={10.1109/LRA.2026.3685447}
}
```

Calibration of Error Distributions in Robot Kinematics for Increased Precision in Manipulation Tasks

Tim Gerstewitz[1], Peter Lehner[1], Lukas Burkhard[1], Natalija Topalovic[1], Alin Albu-Schäffer[1], Daniel Leidner[2], Máximo A. Roa[1]

***Abstract*—The accuracy of robotic forward kinematics is commonly improved by calibration. However, most calibration methods only take deterministic errors, such as inaccurate geometry and unknown stiffnesses, into account and neglect errors with stochastic characteristics, including joint friction, gear backlash and component wear. In order to incorporate these effects, this paper presents a calibration procedure for a probabilistic forward kinematics model which identifies both systematic errors and error sources whose combined effects are modeled as distributions within a kinematic chain. Additionally, we present a method for using the identified distributions to enhance task-relevant accuracy by minimizing the end-effector uncertainty resulting from such distributions in a task-specific way. Finally, this idea is demonstrated in an experiment with a research rover tasked with picking up a payload box from a lander. Here, we show that end-effector error in task-relevant directions can be reduced by 40% by choosing low-uncertainty over high-uncertainty configurations.**

***Index Terms*—Manipulation Planning, Kinematics, Calibration and Identification**

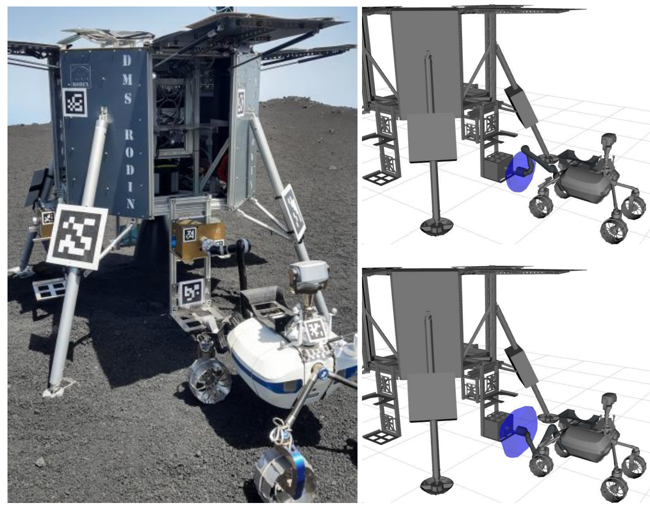

Fig. 1. Visualization of our core idea: By calibrating error distributions inside a kinematic chain, we can select pre-grasp configurations with lower pose uncertainty, thereby reducing the mean task-relevant error. Left: an example application where where the research rover LRU2 [2] picks up a payload box from a lander during DLR's ARCHES mission [3]. Right: one standard deviation of the EE pose is visualized in blue (enlarged by factor 5) with typical uncertainties within the kinematic chain.

I. Introduction

LIGHTWEIGHT and elastic robotic arms commonly need to be calibrated to reduce the influence of inaccuracies in their geometry, stiffnesses of the manipulator materials as well as elastic effects in the robot structure. However, in outdoor use-cases and especially for space exploration scenarios, additional error sources such as joint friction, gear backlash, component wear and environmental influences can also degrade forward kinematics (FK) accuracy by introducing configuration-dependent errors into the kinematic chain. Such errors could become critical in future space exploration missions, in which exploration rovers are expected to autonomously manipulate various objects, for example to prepare landing sites for human astronauts.

We argue that commonly-used FK calibration is insufficient to address most configuration-dependent error sources as it models errors within a kinematic chain as systematic, constant quantities. While this is true for some effects, particularly geometric inaccuracies, it is much harder to describe and incorporate effects such as joint friction or component wear. We adopt the view that the combined influence of such non-systematic effects can be modeled as error probability distributions inside the kinematic chain. Recent work in [1] has introduced a probabilistic kinematics model where such a "noisy" manipulator is used as an additional sensor for improving pose estimation. However, the error locations were identified from system experience and their magnitudes were determined through experiments specifically designed for individual cases. In order to provide a unified procedure for calibrating manipulators featuring both systematic and probabilistic errors, this letter contributes the following:

- An algorithm for the calibration of error probability distributions within a kinematic chain based on a moment-matching formulation.
- An algorithm to choose pre-grasp configurations which minimize the end-effector (EE) pose uncertainty in task-relevant directions to reduce mean task-relevant errors on those pre-grasp configurations, as shown in Fig. 1.
- Validation of both algorithms on a real exploration rover system.

Manuscript received: November, 4, 2025; Revised February, 18, 2026; Accepted April, 9, 2026.

This paper was recommended for publication by Editor Lucia Pallottino upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by Helmholtz Association project iFOODis (contract number KA2-HSC-06) and by Helmholtz Association project ALARO (contract number KA-TCP-22).

[1]T. Gerstewitz, P. Lehner, L. Burkhard, N. Topalovic, A. Albu-Schäffer and M. A. Roa are with the Institute of Robotics and Mechatronics, German Aerospace Center (DLR), 82234 Wessling, Germany {firstname.lastname@dlr.de}.

[2]D. Leidner is with SEAL Robotics GmbH and was with DLR at the time of development {firstname.lastname@seal-robotics.com}.

Digital Object Identifier (DOI): see top of this page.

II. Related Work

Errors in robotic manipulators are caused by multiple factors: [4] categorizes them as geometric inaccuracies (assembly, manufacturing), environmental influences (component wear, thermal deformation) and unknown model parameters for joint offsets, elasticity, friction, gear backlash, etc.

Yet, existing FK calibration procedures mostly focus on identifying geometric inaccuracies and unknown stiffnesses. These procedures use the kinematic Jacobian with respect to the desired parameters to approximate the mapping of parameter errors to EE displacements and invert this relation to identify these parameters. In [5], this is done only for geometric parameters in a least-squares sense based on the Denavit-Hartenberg (DH) convention. The authors of [6] and [7] formulate the problem recursively through an extended Kalman filter and an additional particle filter, respectively. Additionally calibrating joint stiffnesses in the model is performed in [8].

Next to the DH convention, there are other approaches to model kinematics, most notably the Product-of-Exponentials (POE) formula. With calibration based on this FK formula, the work presented in [9] achieves higher EE accuracy solely by considering geometric offsets, while [10] additionally considers joint stiffnesses. However, as for the methods in the last paragraph, these approaches neglect error sources beyond geometry and stiffnesses.

Recently, machine learning methods have also been proposed for model error calibration. A prominent example is [11] which proposes to first identify geometric errors using a POE-based calibration, before learning configuration-dependent residual errors using Gaussian processes (GP). In [12] and [13], the complete residual error of a model is learned using GPs based on a POE and DH formulation, respectively. However, these approaches do not model how error distributions within the kinematics propagate, making them hard to interpret.

Lastly, works on uncertainty representation on rigid-body motions have explored describing error sources in kinematic chains. The authors of [14] propose modeling uncertain geometric parameters as intervals which are compounded along the kinematic chain based on a POE formulation and also give an algorithm for estimating the intervals' widths. However, the generated EE pose bounding box can be prone to overestimation. By contrast, the authors of [15] explore a model with Gaussian-distributed errors and give a numerical example of possible end-effector poses under heavy gear backlash. Yet, no procedure for distribution estimation from only EE measurements is given. Similarly, [1] has recently published a probabilistic robot model in a state estimation context that improves pose estimation accuracy. However, the error locations are identified from system experience and their magnitude is determined with experiments tailored to individual error cases. Estimating the location and scale of errors modeled as distributions in a unified manner as well as differentiating them from systematic, deterministic effects is therefore an open problem to the best of our knowledge.

III. Background

The probabilistic robot kinematics model presented in [1] models a kinematic chain as a series of uncertain relative transformations $\mathcal{X}$ living in the Special Euclidean group $SE(3)$. A POE model formulation, originally presented in [16], is used to express a transformation $\mathcal{X}_{i,i+1} \in SE(3)$ between a link-joint pair at frame $\mathcal{F}_i$ and the subsequent pair at frame $\mathcal{F}_{i+1}$ as a link transformation $\mathcal{X}_{\text{link}} \in SE(3)$ followed by a joint transformation $\mathcal{X}_{\text{joint}} \in SE(3)$, i.e.,

$$\mathcal{X}_{i,i+1} = \mathcal{X}_{\text{link}}\mathcal{X}_{\text{joint}} = e^{\hat{\mathbf{p}}_i} e^{\hat{\mathbf{s}}_i q_i}. \tag{1}$$

Here, both transformations are expressed through the exponential map $e^{(\cdot)}$, which maps the twists $\hat{\mathbf{p}}_i \in \mathfrak{se}(3)$ and $\hat{\mathbf{s}}_i \in \mathfrak{se}(3)$ from Lie algebra to Lie group.

As the Lie algebra is a vector space formed by taking the tangent space to the group, twists describe the nonlinear $SE(3)$ manifold in locally linear coordinates. In [1], the authors propose modeling uncertain transformations through a random, local perturbation in these coordinates, such that the transformation from above becomes

$$\mathcal{X}_{i,i+1} = \mathcal{X}_{\text{link}}\mathcal{X}_{\text{random}}\mathcal{X}_{\text{joint}} = e^{\hat{\mathbf{p}}_i} e^{\hat{\boldsymbol{\xi}}_i} e^{\hat{\mathbf{s}}_i q_i}, \quad \boldsymbol{\xi}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma_i}). \tag{2}$$

With this formulation, the local perturbation $\hat{\boldsymbol{\xi}}_i$ is zero-mean Gaussian-distributed around $\mathcal{X}_{\text{link}}$, as shown in Fig. 2. The complete kinematic chain of a manipulator with k link-joint pairs then reads as

$$\mathcal{X}_{0,k+1} = \left[\prod_{i=1}^{k} e^{\hat{\mathbf{p}}_i} e^{\hat{\boldsymbol{\xi}}_i} e^{\hat{\mathbf{s}}_i q_i}\right] e^{\hat{\mathbf{p}}_{k+1}}, \quad \boldsymbol{\xi}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma_i}). \tag{3}$$

Here, the mean pose of the forward kinematics can be attained simply by setting all $\boldsymbol{\xi}_i$ to zero, while intermediate (and EE for $i = k+1$) uncertainties are given as

$$\boldsymbol{\Sigma}_{0,i} = \mathbf{Ad}_{\bar{\mathcal{X}}_{i,i-1}} \boldsymbol{\Sigma}_{0,i-1} \mathbf{Ad}^T_{\bar{\mathcal{X}}_{i,i-1}} + \boldsymbol{\Sigma}_i. \tag{4}$$

In this expression, the adjoint matrices $\mathbf{Ad}_{\mathcal{X}}$ of a transformation are formed as (for details, see [1] and [17]):

$$\mathbf{Ad}_{\mathcal{X}} = \begin{bmatrix} \mathbf{R} & [\mathbf{t}]_\times \mathbf{R} \\ \mathbf{0} & \mathbf{R} \end{bmatrix} \in \mathbb{R}^{6\times 6}, \tag{5}$$

where $\boldsymbol{R} \in SO(3)$ is a rotation matrix encoding the rotational part of the transformation, $\boldsymbol{t} \in \mathbb{R}^3$ encodes the translation, and $[\mathbf{t}]_\times \in \mathbb{R}^{3\times 3}$ forms a skew-symmetric matrix from $\boldsymbol{t}$'s entries.

Therefore, evaluating the EE uncertainty requires an upwards recursion from the robot base to compose individual pose uncertainties via the adjoints of the corresponding elements. By scaling uncertainties via the adjoint matrices during this recursion, the model fulfills the intuitive notion that rotational errors are scaled by $\boldsymbol{t}$ when propagating to the EE where they can cause both translational and rotational errors. On the other hand, translational errors in the kinematic chain can only cause translational offsets at the end-effector.

IV. Methods

Information about EE pose uncertainty could be used to select approach configurations for object manipulation which minimize this uncertainty in directions relevant to a certain

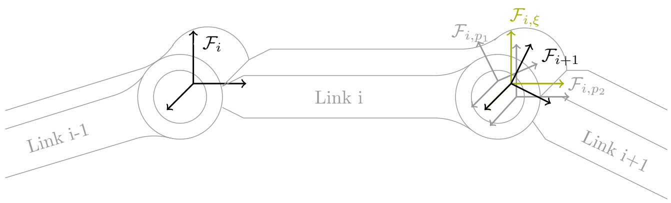

Fig. 2. Two realizations of random offset ξ_i from true frame $\mathcal{F}_{i,\xi}$ to perturbed frame $\mathcal{F}_{i,p}$. We assume perturbations are caused by the combined effects of difficult-to-model factors such as friction or gear backlash, hence choosing a probabilistic model. Adapted from [1].

task. However, this requires differentiating deterministic effects affecting a robot's accuracy from effects modeled as random before identifying the covariances inside the kinematic chain creating this EE pose uncertainty.

A. Forward Kinematics Model

For this work, we model geometric inaccuracies and unknown elasticities in the manipulator as deterministic errors in the kinematic chain. Since the robot model in Eq. 3 does not yet take stiffnesses into account, we combine it with the compliant POE formula proposed in [10]. Here, all joints are assumed to be elastic around their rotation axis $\boldsymbol{\omega}_i$, leading to a deflection $\tilde{q}_i$ of the subsequent link given by

$$\tilde{q}_i = c_i \boldsymbol{\omega}_i^T (\mathbf{p}_{i,k}^{\text{cg}} \times (m_{i,k}\mathbf{g}_i)), \tag{6}$$

where c_i is the compliance coefficient and $\times$ denotes the cross product. The induced torque $\mathbf{p}_{i,k}^{\text{cg}} \times (m_{i,k}\mathbf{g}_i)$ on this joint consists of the acceleration $\mathbf{g}_i$ acting on the combined subsequent masses $m_{i,k}$ hanging on the joint as well as the masses' combined center of gravity $\mathbf{p}_{i,k}^{\text{cg}}$ in the ith frame [10]. The compliant, probabilistic POE model we use is then

$$\mathcal{X}_{0,k+1} = \left[\prod_{i=1}^{k} e^{\hat{\mathbf{p}}_i} e^{\hat{\boldsymbol{\xi}}_i} e^{\hat{\mathbf{s}}_i q_i} e^{\hat{\mathbf{s}}_i \tilde{q}_i}\right] e^{\hat{\mathbf{p}}_{k+1}}, \ \boldsymbol{\xi}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma_i}). \tag{7}$$

B. Calibration for Deterministic Effects

With this model, the deterministic effects of geometric and elastic offsets must first be identified. This can be achieved with the calibration approach outlined in [10]. Based on a dataset of recorded EE poses in different configurations, the error between model and measurements can be mapped to parameter errors via the FK Jacobian using a first-order approximation. In a POE model, it can be shown that the Jacobian $\mathbf{A}$ is composed of individual adjoint matrices [10]. For a single measurement n, the relation between model error $\boldsymbol{y}^{(n)}$ and parameter errors $\mathbf{x}$ is given by [10]

$$\mathbf{y}^{(n)} = \mathbf{A}^{(n)}\mathbf{x}, \tag{8}$$

where

$$\begin{aligned}
\mathbf{y}^{(n)} &= \log(\mathcal{X}_{0,k+1}^{a} \mathcal{X}_{0,k+1}^{-1})^V \\
\mathbf{A}^{(n)} &= \left[\mathbf{Ad}_{\mathcal{X}_{0,\mathbf{p}_1}},\ \mathbf{Ad}_{\mathcal{X}_{0,\tilde{q}_1}}\mathbf{s}_1,\ \mathbf{Ad}_{\mathcal{X}_{0,\mathbf{p}_2}},\ \ldots,\ \mathbf{Ad}_{\mathcal{X}_{0,\mathbf{p}_{k+1}}}\right] \\
\mathbf{x} &= \left[\partial\mathbf{p}_1, \quad \partial c_1, \quad \partial\mathbf{p}_2, \quad \partial c_2, \quad \ldots, \quad \partial\mathbf{p}_{k+1}\right]^T.
\end{aligned} \tag{9}$$

Here, $\mathcal{X}_{0,k+1}^{a}$ is the actual, measured end-effector pose, $\mathcal{X}_{0,k+1}^{-1}$ is the nominal model's EE pose, $\mathbf{A}^{(n)}$ is the measurement's Jacobian, and $\partial\mathbf{p}_i$ and ∂c_i are the to-be-calibrated corrections to the geometry and to the compliance coefficients. These parameters can be identified iteratively [9]:

$$\begin{aligned}
\mathcal{X}_{0,k+1}^{w} &= f(\mathbf{x}_{\text{nom}} + \mathbf{x}_w), \quad \mathbf{x}_w = \mathbf{A}_w^{\dagger}\mathbf{y}_w \\
\mathbf{y}_{w+1} &= \log(\mathcal{X}_{0,k+1}^{a} \mathcal{X}_{0,k+1}^{w^{-1}})^V
\end{aligned} \tag{10}$$

Here, f is the FK function according to Eq. 7, while $\dagger$ is the pseudoinverse. Additionally, $\mathbf{y}_w$ and $\mathbf{A}_w$ are stacked versions of $\mathbf{y}^{(n)}$ and $\mathbf{A}^{(n)}$ for N measurements in the wth iteration. In line 1, the FK parameters before calibration $\mathbf{x}_{\text{nom}}$ are corrected and yield new nominal EE poses $\mathcal{X}_{0,k+1}^{w}$. These lead to new FK errors $\mathbf{y}_{w+1}$ serving as inputs for the next iteration. The procedure stops once the error $||\mathbf{y}_{w+1} - \mathbf{y}_w|| < \epsilon$ falls below a threshold ϵ [9].

C. Calibration for Random Effects

After Eq. 10 has converged, residuals

$$\boldsymbol{r} = [\boldsymbol{r}^{(1)}, ..., \boldsymbol{r}^{(N)}]^T, \qquad \boldsymbol{r}^{(n)} = \log(\mathcal{X}_{0,k+1}^{a} \mathcal{X}_{0,k+1}^{\text{cal}^{-1}})^V \tag{11}$$

are still present between the measured EE poses and the calibrated poses $\mathcal{X}_{0,k+1}^{\text{cal}^{-1}}$. A typical calibration assumption is that these residuals are solely due to measurement noise; however, under the assumption of error distributions within the kinematics, the residuals are due to measurement noise as well as these distributions. Our core idea is to exploit this insight to calibrate the covariance matrices of effects assumed to be random in the kinematic chain.

In the following, we assume that all covariance matrices are diagonal matrices and all off-diagonal entries are zero. This assumption implies that stochastic errors are independent at their source location within the kinematic chain. Nonetheless, since they combine with other covariances away from their source location due to the adjoint-mixing property of the PPOE model (Eq. 4), it is still possible to approximate error effects with correlations. Note, that the derivations and resulting procedure could also be applied to full covariance matrices, but experimental validation is left for future work. We first reformulate the recursive covariance propagation from Eq. 4. Using properties of the adjoint matrix, namely $\mathbf{Ad}_{\mathcal{X}^{-1}} = \mathbf{Ad}_{\mathcal{X}}^{-1}$ [17, Eq. (33)] and $\mathbf{Ad}_{\mathcal{X}\mathcal{Y}} = \mathbf{Ad}_{\mathcal{X}}\mathbf{Ad}_{\mathcal{Y}}$ [17, Eq. (34)], we can write

$$\begin{aligned}
\boldsymbol{\Sigma}_E &= \boldsymbol{\Sigma}_n + \mathbf{Ad}_{E,k}\boldsymbol{\Sigma}_{0,k}\mathbf{Ad}_{E,k}^T + \\
&= \boldsymbol{\Sigma}_n + \mathbf{Ad}_{E,k}(\mathbf{Ad}_{k,k-1}\boldsymbol{\Sigma}_{0,k-1}\mathbf{Ad}_{k,k-1}^T + \boldsymbol{\Sigma}_{k-1})\mathbf{Ad}_{E,k}^T \\
&= \cdots = \boldsymbol{\Sigma}_{\text{n}} + \sum_{j=1}^{k} \mathbf{Ad}_{j,E}^{-1}\ \boldsymbol{\Sigma}_j\ \mathbf{Ad}_{j,E}^{-T}.
\end{aligned} \tag{12}$$

Here, we have used the shorthand $\mathbf{Ad}_{E,k} = \mathbf{Ad}_{\bar{\mathcal{X}}_{E,k}}$ for the transformation from end-effector k+1th link to kth link, while $\boldsymbol{\Sigma}_E$ and $\boldsymbol{\Sigma}_n$ denote the covariance at the EE and of the measurement noise, respectively.

From this formulation, it can be clearly seen that both the uncertainty of the observation process as well as intra-chain probabilistic errors contribute to the end-effector uncertainty.

1) Moment Matching Perspective: Eq. 12 therefore forms a model on how second moments of error distributions propagate, which we can exploit for calibration using a method of moments approach [18, p. 117]. This can be seen as an extension to traditional calibration methods which match first moment error information like Eq. 10 (the error means, what we referred to as deterministic effects) to second moments.

To do so, we approximate the end-effector covariance in the ith configuration $\mathbf{\Sigma}_E^{(i)}$ using the corresponding residual $\boldsymbol{r}^{(i)}$ by setting $\mathbf{\Sigma}_E^{(i)} = \boldsymbol{r}^{(i)}\boldsymbol{r}^{(i)^T}$. We can then substitute this approximation into Eq. 12 and vectorize the resulting expression:

$$\operatorname{vec}\left(\mathbf{I}_6\mathbf{\Sigma}_n\mathbf{I}_6 + \sum_{j=1}^{k}\mathbf{Ad}_{j,E}^{(i)^{-1}}\ \mathbf{\Sigma}_j\ \mathbf{Ad}_{j,E}^{(i)^{-T}}\right) = \operatorname{vec}\left(\boldsymbol{r}^{(i)}\boldsymbol{r}^{(i)^T}\right)$$
$$\Leftrightarrow \underbrace{\left[\mathbf{I}_{36},\ \dots\ \left(\mathbf{Ad}_{j,E}^{(i)}\otimes\mathbf{Ad}_{j,E}^{(i)}\right)^{-1},\ \dots\right]}_{\mathbf{B}^{(i)}}\boldsymbol{\theta} = \underbrace{\operatorname{vec}\left(\boldsymbol{r}^{(i)}\boldsymbol{r}^{(i)^T}\right)}_{\mathbf{z}^{(i)}}. \tag{13}$$

Here, we have lumped together all covariance entries in the vector $\boldsymbol{\theta} = [\operatorname{vec}(\mathbf{\Sigma}_n), \operatorname{vec}(\mathbf{\Sigma}_1), ..., \operatorname{vec}(\mathbf{\Sigma}_k)]$, where the vectorization operation $\operatorname{vec}(\cdot)$ stacks matrix columns into a long column vector. In addition, the operator $\otimes$ denotes the Kronecker product between two matrices, which has a useful relation with the $\operatorname{vec}(\cdot)$ operation applied here [19, p. 60, Eq. (520)]).

The deterministic calibration's measurement residuals $\mathbf{z} \in \mathbb{R}^{6N\times 1}$ are propagated back into the kinematic chain through the matrix $\mathbf{B} \in \mathbb{R}^{6N\times 6(k+1)}$. By simply inverting $\mathbf{B}$ however, we neglect the fact that covariance matrices have to be *positive semi-definite (psd)*, putting a constraint on the possible values for their entries. Therefore, the calibration problem for our covariances reads as:

$$\begin{aligned} &\min_{\boldsymbol{\theta}}\ ||\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\theta}||_2 \\ \text{s. t.}\quad &\operatorname{unvec}(\boldsymbol{\theta}_j)\ \text{psd}\quad j = 1,\dots,k+1 \end{aligned} \tag{14}$$

The $\operatorname{unvec}(\cdot)$ operation reassembles a matrix from its stacked version created through vectorization. To solve this calibration problem, we adopt a *Projected Gradient Descent (Proj-GD)* procedure: specifically, after updating the covariance entries $\boldsymbol{\theta}$, we reassemble them into their respective matrices before projecting these matrices into the psd set. For an arbitrary matrix it can be shown that the closest psd matrix (with respect to the Frobenius norm) is the one that has its negative eigenvalues replaced with zeros [20]. Overall, our procedure for calibrating covariance matrices of random effects consists of the following two steps

$$\begin{aligned} \boldsymbol{\theta}_{w+1} &= \boldsymbol{\theta}_w - \alpha\boldsymbol{B}^T(\boldsymbol{B}\boldsymbol{\theta}_w - \boldsymbol{z}) \\ \boldsymbol{\theta}_{w+1} &= \operatorname{Proj}_{>0}(\boldsymbol{\theta}_{w+1}). \end{aligned} \tag{15}$$

Here, α is a scaling parameter, w denotes an iteration counter and $\operatorname{Proj}_{>0}$ is the projection back onto the set of psd matrices (see [20, Theorem 2]).

2) Maximum-Likelihood (ML) Perspective: In the following, we will show that the same estimator for covariance entries can be derived using a ML formulation given the first moments (deterministic effects) of the error distributions. We first assume that all residuals are independent from each other. In addition, we assume that each measurement's residual after systematic calibration is a sample drawn from the EE distribution in this measurement's manipulator configuration. When each measurement is obtained in a different arm configuration, each sample stems from a different EE distribution due to the configuration-dependent propagation rule in Eq. 12. *Importantly, this means that each measurement's residual is a sample from a different distribution and that the residuals are therefore not identically distributed.* Hence the likelihood for N residual errors consists of N different distributions, leading to

$$\begin{aligned} \mathcal{L}(\mathbf{\Sigma}_n,\mathbf{\Sigma}_1,\dots,\mathbf{\Sigma}_k) &= \prod_{i=1}^{N}\mathcal{N}(0,\mathbf{\Sigma}_E^{(i)}) \\ &= \prod_{i=1}^{N}(2\pi)^{-3}\det(\mathbf{\Sigma}_E^{(i)})^{-\frac{1}{2}}\exp\left(-\frac{1}{2}\boldsymbol{r}^{(i)^T}\mathbf{\Sigma}_E^{(i)^{-1}}\boldsymbol{r}^{(i)}\right). \end{aligned} \tag{16}$$

In order to calibrate the covariances, the idea is to maximize this likelihood with respect to the covariance entries of elements within the kinematic chain (as well as the measurement noise). This is equivalent to minimizing the negative log-likelihood (NLL) [18, p. 8]:

$$\begin{aligned} &\underset{\mathbf{\Sigma}_n,\mathbf{\Sigma}_1,\dots,\mathbf{\Sigma}_k}{\arg\max}\ \mathcal{L}(\mathbf{\Sigma}_n,\mathbf{\Sigma}_1,\dots,\mathbf{\Sigma}_k) = \arg\min_{\boldsymbol{\theta}} -\log(\mathcal{L}(\boldsymbol{\theta})) = \\ &= \arg\min_{\boldsymbol{\theta}}\sum_{i=1}^{N}\left[\log(\det(\mathbf{\Sigma}_E^{(i)})) + \boldsymbol{r}^{(i)^T}\mathbf{\Sigma}_E^{(i)^{-1}}\boldsymbol{r}^{(i)}\right]. \end{aligned} \tag{17}$$

As before, we have lumped together all covariance entries in the vector $\boldsymbol{\theta}$. We find an optimum of this NLL via a gradient descent procedure. To derive the gradient, we use the chain rule to first find the derivative of $f(\boldsymbol{\theta}) = -\log(\mathcal{L}(\boldsymbol{\theta}))$ with respect to the covariance entries, before finding the derivative of $\mathbf{\Sigma}_E^{(i)}$ with respect to $\boldsymbol{\theta}$. According to [19, p. 9, Eq. (49)], we obtain

$$\frac{\partial}{\partial\mathbf{\Sigma}_E^{(i)}}\log(\det(\mathbf{\Sigma}_E^{(i)})) = \mathbf{\Sigma}_E^{(i)^{-T}}, \tag{18}$$

as well as [19, p. 10, Eq. (61)] for the second expression:

$$\frac{\partial}{\partial\mathbf{\Sigma}_E^{(i)}}\mathbf{r}^{(i)^T}\mathbf{\Sigma}_E^{(i)^{-1}}\mathbf{r}^{(i)^T} = -\mathbf{\Sigma}_E^{(i)^{-T}}\mathbf{r}^{(i)}\mathbf{r}^{(i)^T}\mathbf{\Sigma}_E^{(i)^{-T}}. \tag{19}$$

With these expressions, we now have

$$\frac{\partial f(\boldsymbol{\theta})}{\partial\boldsymbol{\theta}} = \sum_{i=1}^{N}\left[\mathbf{\Sigma}_E^{(i)^{-T}}\left(\mathbf{I}_6 + \mathbf{r}^{(i)}\mathbf{r}^{(i)^T}\mathbf{\Sigma}_E^{(i)^{-T}}\right)\left(\frac{\partial\mathbf{\Sigma}_E^{(i)}}{\partial\boldsymbol{\theta}}\right)\right] \overset{!}{=} \mathbf{0} \tag{20}$$

for the gradient which must be zero at the optimum. Since Eq. 12 is linear with respect to the covariance matrices and hence also linear in $\boldsymbol{\theta}$, the derivative $\partial\mathbf{\Sigma}_E^{(i)}/\partial\boldsymbol{\theta}$ is constant and non-zero. Therefore

$$\mathbf{I}_6 + \mathbf{r}^{(i)}\mathbf{r}^{(i)^T}\mathbf{\Sigma}_E^{(i)^{-T}} \overset{!}{=} \mathbf{0} \tag{21}$$

must hold for an optimum in the NLL $f(\boldsymbol{\theta})$. Substituting the covariance propagation relation from Eq. 12 and vectorizing the resulting equation recovers the same estimator, Eq. 13, as obtained via the moment-matching approach.

D. Choosing Configurations based on Uncertainty

Knowing the magnitude of uncertainties in the kinematic chain could be useful for selecting configurations that increase the reliability of a given task by minimizing EE pose uncertainty. However, in many robotic tasks, certain pose components require more precise EE positioning than others. It is desired to prioritize these dimensions with higher relative weighting for uncertainty. Based on this idea of task-dependent weighting of dimensions, we propose a metric to evaluate task approach configurations with regard to their uncertainty. By multiplying the predicted EE covariance $\mathbf{\Sigma}_E$ with a weighting matrix $\mathbf{W}$, we obtain the task-relevant covariance $\mathbf{\Sigma}_E^*$:

$$\mathbf{\Sigma}_E^* = \mathbf{W}\mathbf{\Sigma}_E\mathbf{W}^T \quad \text{with} \quad w_{ii} \in [0, 1], \quad w_{ij} = 0 \tag{22}$$

Here, the weighting matrix encodes the importance of being accurate in certain EE pose components relative to others. The weights are set by a user and should reflect the relative accuracy requirements of the task at hand: for example, for a peg-in-hole task, weights in directions orthogonal to the hole should be chosen higher than those in the direction of the hole since aligning accurately with the hole is often more critical than insertion depth.

With this task-relevant EE uncertainty, we can choose a configuration $\mathbf{q}^* \in \mathcal{Q}$ from a set of possible pre-grasp configurations $\mathcal{Q}$ that has smaller uncertainty along the desired dimensions by minimizing $\mathbf{\Sigma}_E^*$'s trace:

$$\mathbf{q}^* = \underset{\mathbf{q}\in\mathcal{Q}}{\operatorname{argmin}}\, \operatorname{tr}(\mathbf{\Sigma}_E^*) = \underset{\mathbf{q}\in\mathcal{Q}}{\operatorname{argmin}}\, \sigma_{ii}^2. \tag{23}$$

This set of possible approaches will only hold multiple possibilities if the manipulator has at least one degree of freedom more than the task requires. However, if the task can be achieved with multiple configurations, then Q could be produced by repeatedly running a sampling-based motion planner. We choose to minimize the trace in Eq. 23 because it is defined as the sum of eigenvalues, which means minimization shrinks uncertainty in the principal covariance directions uniformly. This is not the case for a configuration metric based on determinants or maximum eigenvalue, which would prefer to minimize only one direction.

V. Experiments

To validate our proposed methods, we performed experiments with DLRs Lightweight Rover Unit 2 (LRU2) [2] shown in Fig. 3a. This rover features a *Jaco2* manipulator by *Kinova* [21] made from carbon fiber which introduces flexibilities in the kinematics. Additionally, the arm is mounted at the rear of the base on a carbon fiber plate which bends depending on the arm's configuration during manipulation, thus introducing a configuration-dependent uncertainty in the transformation chain from rear-facing camera to the Jaco2 arm that is not captured by an existing extrinsic camera calibration. As such, we hypothesize that the probabilistic and compliant FK model presented in Eq. 7 is an appropriate description.

A. Calibration Experiment

1) Setup: First, we calibrated the rover's manipulator and drove the Jaco2 to 154 different configurations in the field of view of the rover's rear-facing camera. This workspace region also corresponds to the region associated with picking a payload box from the lower shelf of the RODIN lander used in DLR's ARCHES space-analogue mission [3]. We used AprilTags [22] attached to the end-effector to detect the resulting pose with multiple tags as described in [23]. This is shown in Fig. 3a and Fig. 3b. Additionally, we performed an outlier removal step for the AprilTag measurements. While the AprilTags on the camera facing side of the Jaco2's flange are usually well visible, the tags on the side of the flange are often hardly visible, which can lead to offsets in orientation in particular (see Fig. 4a). We calculated the standard deviation of each pose component (x, y, z, roll, pitch, yaw) across the dataset and removed all poses where at least one component exceeded two component standard deviations. The threshold was chosen empirically to balance between outliers skewing the estimation while still retaining enough data for stable estimates. This reduced the number of EE poses in the dataset from 154 to 119.

Afterwards, we split the dataset into training and test splits using a $90/10$ ratio. To investigate the variability of results, we performed this train-test-split 10 times: we randomly permuted the order of configurations before splitting, thus providing different training and test sets for each split. For these perturbations, we fixed the random number generator's seed to a sequential number between 1 and 10.

For each of these 10 splits, we subsequently performed the calibration procedure for systematic geometric and compliant effects using the respective training sets as inputs to Eq. 10. During this calibration, we estimated the geometric offsets of all link-joint pairs in the kinematic chain. Moreover, we assumed that only the manipulator's base joint as well as its first joint are compliant, since these carry almost all of the arm's weight and are therefore subjected to the highest load. To additionally take into effect the flexible carbon fiber plate the Jaco2 is mounted on, we also introduced two virtual compliant joints orthogonal to the real joint's rotation axis for the base joint and the first joint. These are inserted into the kinematic chain before the real joint and model the base plate's configuration-dependent deformation. To mitigate coupling with geometric parameters, we removed columns from matrix $\mathbf{A}$ corresponding to elastic joints that did not contribute to the matrix rank.

Next, we used the residual errors of the calibration for deterministic effects to calibrate for the diagonal entries of all link-joint pairs' covariance matrices and those of the noise using the optimization parameters shown in Tab. I. Note that we removed 8 columns of matrix $\mathbf{B}$ prior to optimization as they did not contribute to the matrix rank. The condition number of $\mathbf{B}$ then ranged from 240 to 244, indicating that the estimation is not highly sensitive to the training data.

2) Results: The results of the calibration for systematic effects are shown Tab. II. During 9 out of 10 train-test-splits, the mean translational and the mean rotational error on the test

(a) Left: The rover in an example configuration while collecting the dataset. Right: Coordinate frames used during calibration and the rough field-of-view of the camera (blue).

(b) Left: Jaco2 configuration with the AprilTag marker ring visible. Right: Same configuration with overlaid detection and offset w. r. t. nominal FK shown (dotted).

Fig. 3. Calibration experiment setup for LRU2's Jaco2 arm.

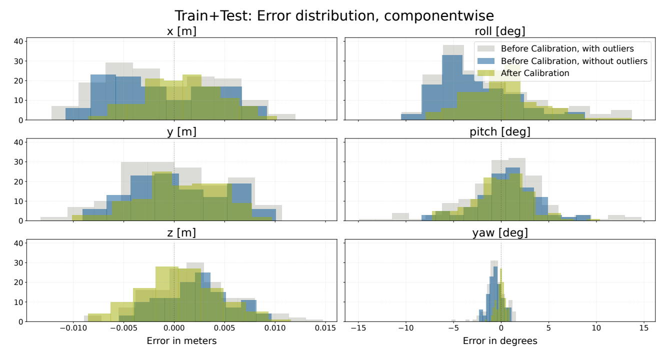

(a) Pose error component distributions of seed 1's combined training and test set with camera frame as reference frame.

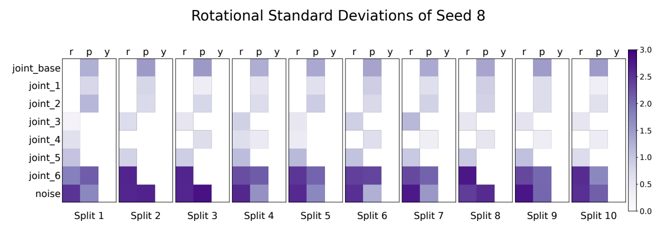

(b) Calibrated standard deviations for random offsets within the Jaco2's kinematic chain.

Fig. 4. Overview of results of the calibration experiment for LRU2's Jaco2 arm.

Optimization Parameters

Initial $\mathbf{\Sigma}_i$	α	Stopping Criteria	Runtime
Trans.: 0.001m Rot.: 0.0085rad	0.0025	$\|\|\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\theta}_{w+1}\|\|_2 - \|\|\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\theta}_w\|\|_2 < 1e-10$	$7.56 \pm$ 0.53s

TABLE I
OPTIMIZATION PARAMS. FOR COVARIANCE CALIBRATION.

Mean Error on Test Set in mm / deg

	Seed	1	2	3	4	5	6	7	8	9	10
Init.	Trans.	8.40	6.28	7.46	7.49	8.65	7.33	8.15	7.36	7.93	7.13
	Rot.	4.88	5.61	4.22	5.91	5.68	5.14	5.42	5.47	3.73	5.62
Cal. Ours	Trans.	7.10	4.95	5.91	7.38	7.44	7.71	5.57	6.31	7.62	6.48
	Rot.	5.85	4.06	2.70	4.97	4.13	4.92	4.80	4.30	3.30	3.89
Cal. GP	Trans.	11.50	10.46	9.36	9.81	10.25	9.18	9.90	10.9	10.52	8.53
	Rot.	5.41	4.82	3.27	4.80	5.21	4.73	5.20	4.33	3.11	4.62

TABLE II
MANIPULATOR ERROR COMPARISON ON EACH SEED'S TEST SETS BEFORE AND AFTER RUNNING SYSTEMATIC EFFECTS CALIBRATION.

sets decreased significantly. Moreover, we show the EE error distributions of all pose components for the representative train-test-split 1 in Fig. 4a. Here, all component distributions become roughly zero-centered and Gaussian-shaped with a smaller base for the distributions after calibrating.

In addition, we compared our calibration approach to the proposed approach in [13], where GPs are used as a configuration-dependent model for the FK error in each dimension. We trained these GPs on our respective training sets per split using the kernel function described in [13]. The resulting performance on the test set is shown in Tab. II. Compared to our procedure, this approach lead to higher errors in both translational and rotational components.

Results for the subsequent calibration for the diagonal elements of covariance matrices can be seen in Fig. 4b. We show the square root of these elements, that is, standard deviations for all frames where non-zero elements were identified. During the calibration, we represented translational uncertainty components in meters and rotational components in radians with a scaling factor of 1 applied between them. A key observation is that the identified uncertainties grow in magnitude the closer their origin frame is to the end-effector, which is an intuitive and expected result. Additionally, the calibration procedure identifies almost entirely rotational uncertainties within the manipulator's kinematic chain regardless of the train-test-split used. Only for split 10 a small translational uncertainty of roughly 1mm is identified at the base of the kinematic chain.

B. Noise Experiment

To validate the estimated covariances, we performed an additional experiment to characterize the observed measurement noise: we set the Jaco2 to 5 different configurations and ran the Apriltag detection 10 times for each of these configurations.

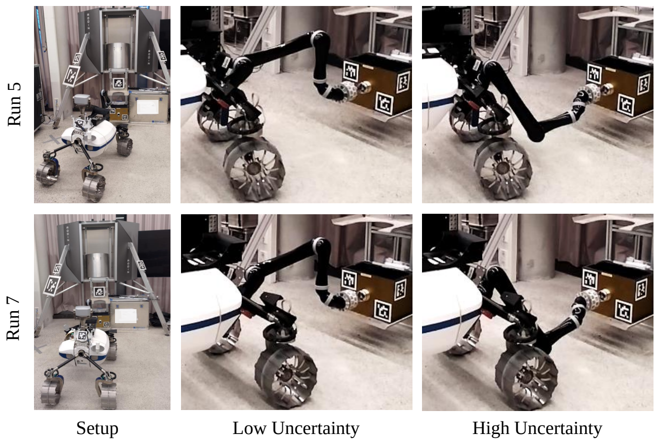

Fig. 5. Visualization of measurement runs 5 and 7. Left: different rover base poses in front of the lander. Middle+Right: found approach configurations with lowest and highest uncertainty in the EE's $x-y$-plane.

Afterwards we subtracted the mean across these 10 detections before calculating the standard deviation per pose component across all 50 observations.

1) Results: The standard deviations across these 50 configurations were 0.27mm, 0.397mm and 1.114mm for x, y and z components, and 4.309deg, 2.216deg and 0.414deg in roll, pitch and yaw directions respectively.

C. Manipulation Experiment

In the third experiment we investigated the applicability of this calibrated uncertainty information to choose approach configurations which minimize possible EE deviations in directions relevant for payload box picking.

1) Setup: We show the experiment in Fig. 5. Here, a lander mock-up can be seen in the background, with a payload box placed on the lower shelf. For a successful payload pickup, the Jaco2's docking interface has to fit to the boxes' silver docking adapter, seen in the payload's center in the images in Fig. 5. During this process, the alignment of docking interface and adapter in x and y directions of the EE frame (last frame in Fig. 3a) are most relevant, along with a relatively low error in the roll and pitch directions.

A single experiment run consisted of driving the rover manually in front of the lander, with its manipulator facing towards the shelf. The rover then estimated its pose in the lander frame using the AprilTags visible on the shelves and the supporting structure. Based on this information, we planned 10 approach configurations for the Jaco2 manipulator 5cm in front of the payload boxes' docking adapter. Next, we evaluated these configurations with respect to their predicted uncertainty in the EE's task-relevant x, y, roll and pitch directions, using the covariance matrices from split 8 since they are the most representative of all splits. For this evaluation, we used the weighting matrix $\mathbf{W} = \mathrm{diag}([0.4, 0.4, 0.0, 0.1, 0.1, 0.0])$ in Eq. 22 and again represented covariance entries in meters and radians with a scaling factor of 1 applied between them. Afterwards, we planned and executed a motion to each of these 10 approach configurations, returning to a neutral home configuration in between. Once the manipulator had settled in these approaches, we used the AprilTag detection again to estimate the actual EE pose with the back-facing camera. Overall, we repeated this experiment ten times, manually positioning the rover in a different base pose each time. Fig. 5 displays this base pose change for measurement runs 5 and 7.

Fig. 6. Norm of end-effector error in $x-y$-plane during all ten runs box-picking measurement runs.

Correlation coefficient of predicted metric (Eq. 23) and observed x-y-error norm

Seed	1	2	3	4	5	6	7	8	9	10
Ours	0.00	-0.07	0.66	-0.05	0.16	-0.04	0.68	0.28	0.19	0.67
GP	-0.04	0.48	0.60	-0.14	-0.05	0.35	-0.09	0.47	-0.36	-0.09

TABLE III
CORRELATION OF PREDICTED UNCERTAINTY WITH OBSERVED ERROR BETWEEN OUR AND THE GP APPROACH [13].

2) Results: For these 10 measurement runs, the norm of the observed error in the task-relevant EE $x-y$-plane is shown in Fig. 6. The configurations with the lowest predicted uncertainty in task-relevant directions in EE frame produced a lower error in x and y dimensions than the configurations with the highest predicted uncertainty in 7 out of 10 cases. In addition, while the 10 lowest uncertainty configurations produced a mean translational error $\bar{e} = \frac{1}{10}\sqrt{({}^{E}\Delta x)^2 + ({}^{E}\Delta y)^2}$ of 2.495mm (${}^{E}\Delta x$ denotes x difference in EE frame), the highest uncertainty configurations lead to a mean translational error of 4.133mm during our measurements runs, realizing a roughly 40% decrease in mean task-relevant FK error. We show the correlation coefficient between the predicted uncertainty metric Eq. 23 with the observed task-relevant error across all approaches during the runs in Tab. III. We compared this correlation coefficient with the predicted uncertainty from the GP method [13], where we assigned each GP's predicted variance to the corresponding entry on the diagonal of the EE covariance and used Eq. 23 to assign a metric to each approach.

VI. DISCUSSION

Through the systematic calibration, we have improved the Jaco2's mean FK error to 6.65mm in translational error across our ten test sets (see Tab. II). This is not the case for the GP-based calibration model [13] which we attribute to our relatively small training set causing the GPs to overfit. The observation that pose error distributions become zero centered and Gaussian after calibrating for deterministic effects (see Fig. 4a) not only underlines this mean error reduction, but is also important for the subsequent uncertainty calibration, as this requires Gaussian-distributed EE errors.

Moreover, our analysis indicates that calibrated standard deviations increase in magnitude the closer their origin frame is to the EE. This seems plausible, as the uncertainties' influence gets scaled by the manipulator length through the adjoints. The calibration also identifies almost only rotational uncertainties within the kinematics. This also seems reasonable as random rotational offsets cause both translational and rotational effects further down the kinematic chain, whereas the influence of translational effects remains translational. Our subsequent comparison against sampled noise indicates that we can capture the dominant nature of roll and pitch noise in the AprilTag detection through the covariance calibration. Especially the pitch uncertainty magnitude appears to be captured well, while some magnitude of the measurement uncertainty in roll and z directions is wrongly attributed to be within the kinematic chain. For the noise in z direction, this is again plausible, as translational EE offsets can be caused by random rotational offsets.

Nonetheless, the results of our third experiment show that subsequently using kinematic uncertainty information can be beneficial for selecting pre-grasp configurations for object manipulation: in our scenario, this benefit amounts to a relative mean task-relevant error reduction of approximately 40% between lower-uncertainty and higher-uncertainty pre-grasps, indicating that we have captured at least some of the stochasticity in the manipulator during calibration. In addition, we show that using the information from covariance calibration to score pre-grasps with the presented uncertainty metric is *at worst* uncorrelated with the observed error but can lead to considerable error reduction. Based on our correlation analysis, this is not the case when using uncertainty information from a GP-based method. Finally, we attribute the observed dependency of some uncertainty elements on which calibration data was used, to local optima in the likelihood function of our calibration approach.

VII. Conclusion

We have presented an extension of well-known calibration methods for robot kinematics that separates error sources within a kinematic chain into deterministic effects and effects modeled as probability distributions. While we calibrate deterministic effects using known approaches, we have derived a novel moment-matching-based extension to additionally calibrate error distributions within a manipulator without the need for additional data or procedures. Moreover, we have defined a criterion to evaluate manipulation approach configurations with regard to end-effector pose uncertainty resulting from error distributions within the kinematic chain. In an experiment with DLR's LRU2 research rover, we combined both contributions to show that the proposed distribution quantification yields plausible values and can be used to decrease the mean error of manipulation approach configurations by selecting low-uncertainty configurations for a given task.

References

[1] L. Meyer, K. H. Strobl, and R. Triebel, "The probabilistic robot kinematics model and its application to sensor fusion," in *IEEE/RSJ Int. Conf. Intel. Robots, Sys.*, pp. 3263–3270, 2022.

[2] M. J. Schuster, S. G. Brunner, K. Bussmann, S. Büttner, A. Dömel, *et al.*, "Towards autonomous planetary exploration," *J. Intel. Robot. Sys.*, vol. 93, pp. 461–494, 2019.

[3] M. J. Schuster, M. G. Müller, S. G. Brunner, H. Lehner, P. Lehner, *et al.*, "The ARCHES space-analogue demonstration mission: Towards heterogeneous teams of autonomous robots for collaborative scientific sampling in planetary exploration," *IEEE Robot. Automat. Letters*, vol. 5, no. 4, pp. 5315–5322, 2020.

[4] B. Karan and M. Vukobratovic, "Calibration and accuracy of manipulation robot models - an overview," *Mech. Mach. Theory*, vol. 29, pp. 479–500, 1994.

[5] S. Hayati, K. Tso, and G. Roston, "Robot geometry calibration," *IEEE Int. Conf. Robot. Automat.*, pp. 947–951, 1988.

[6] I.-W. Park, B.-J. Lee, S.-H. Cho, Y.-D. Hong, and J.-H. Kim, "Laser-based kinematic calibration of robot manipulator using differential kinematics," *IEEE Trans. Mech.*, vol. 17, pp. 1059–1067, 2012.

[7] Z. Jiang, W. Zhou, H. Li, Y. Mo, W. Ni, and Q. Huang, "A new kind of accurate calibration method for robotic kinematic parameters based on the extended kalman and particle filter algorithm," *IEEE Trans. Industr. Electron.*, vol. 65, no. 4, pp. 3337–3345, 2018.

[8] J. Zhou, H. Nguyen, and H. Kang, "Simultaneous identification of joint compliance and kinematic parameters of industrial robots," *Int. J. Prec. Eng. Manufact.*, vol. 15, pp. 2257–2264, 2014.

[9] L. M. Chen, G. Yang, C. T. Tan, and S. H. Yeo, "Local POE model for robot kinematic calibration," *Mech. Mach. Theory*, vol. 36, no. 11, pp. 1215–1239, 2001.

[10] P. Y. Tao, G. Yang, Y. C. Sun, M. Tomizuka, and C. Y. Lai, "Product-of-exponential (POE) model for kinematic calibration of robots with joint compliance," in *IEEE Int. Conf. Adv. Intel. Mech.*, pp. 496–501, 2012.

[11] W. Jing, P. Y. Tao, G. Yang, and K. Shimada, "Calibration of industry robots with consideration of loading effects using product-of-exponential (POE) and gaussian process (GP)," *IEEE Int. Conf. Robot. Automat.*, pp. 4381–4385, 2016.

[12] W. Shen, G. Yang, T. Zheng, W. Y., K. Yang, Z. Fang., and Z. C., "An integrated accuracy enhancement method for cable-driven flexible continuum robot," *IEEE Int. Conf. Adv. Intel. Mechatron.*, pp. 1121–1126, 2019.

[13] E. Daş and J. W. Burdick, "An active learning based robot kinematic calibration framework using gaussian processes," in *IEEE Int. Conf. Robot. Automat.*, pp. 11495–11501, 2023.

[14] M. R. Pac and D. O. Popa, "Interval analysis of kinematic errors in serial manipulators using product of exponentials formula," *IEEE Trans. Automat. Sci. Eng.*, vol. 10, no. 3, pp. 525–535, 2013.

[15] Y. Wang and G. Chirikjian, "Error propagation on the euclidean group with applications to manipulator kinematics," *IEEE Trans. Robot.*, vol. 22, no. 4, pp. 591–602, 2006.

[16] K. Okamura and F. C. Park, "Kinematic calibration using the product of exponentials formula," *Robotica*, vol. 14, pp. 415 – 421, 1996.

[17] J. Solà, J. Deray, and D. Atchuthan, "A micro lie theory for state estimation in robotics," *ArXiv*, vol. abs/1812.01537, 2018.

[18] K. P. Murphy, *Probabilistic Machine Learning: An introduction*. MIT Press, 2022.

[19] K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook*. Tech. Univ. Denmark, 2012.

[20] N. Higham, "Computing a nearest symmetric positive semidefinite matrix," *Lin. Alg. App.*, vol. 103, pp. 103–118, 1988.

[21] Kinova, "Robotic arm — kinova - assistive technologies." https://www.kinovarobotics.com/product/jaco, March 2025.

[22] E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *IEEE Int. Conf. Robot. Automat.*, pp. 3400–3407, 2011.

[23] C. Nissler, S. Buettner, Z.-C. Marton, L. Beckmann, and U. Thomasy, "Evaluation and improvement of global pose estimation with multiple apriltags for industrial manipulators," in *IEEE Int. Conf. Emerging Tech., Factory Automat.*, pp. 1–8, 2016.