In Proceedings of the IEEE-RAS International Conference on Humanoid Robots (Humanoids), Osaka, Japan, to appear December 2012

Bayesian Calibration of the Hand-Eye Kinematics of an Anthropomorphic Robot

Uwe Hubert, Jörg Stückler, and Sven Behnke

Abstract— **We present a Bayesian approach to calibrating the hand-eye kinematics of an anthropomorphic robot. In our approach, the robot perceives the pose of its end-effector with its head-mounted camera through visual markers attached to its end-effector. It collects training observations at several configurations of its 7-DoF arm and 2-DoF neck which are subsequently used for an optimization in a batch process. We tune Denavit-Hartenberg parameters and joint gear reductions as a minimal representation of the rigid kinematic chain. In order to handle the uncertainties of marker pose estimates and joint position measurements, we use a maximum a posteriori formulation that allows for incorporating prior model knowledge. This way, a multitude of parameters can be optimized from only few observations. We demonstrate our approach in simulation experiments and with a real robot and provide in-depth experimental analysis of our optimization approach.**

I. INTRODUCTION

The majority of state-of-the-art methods for visual robot control assume a precise calibration of the robot kinematics with respect to task-relevant sensors. We propose a novel method for calibrating the hand-eye kinematics of an anthropomorphic robot. Many previous calibration methods require tedious learning procedures in which many training examples need to be collected, or restrict the calibration to only few parameters to avoid overfitting problems if only a small amount of training data is available. Limiting the set of parameters, however, also assumes that most parameters of the kinematic structure are precisely known a priori. This in turn requires precise manufacturing of the robot and rigid mechanical structures, the latter being difficult to achieve during a long life-time of the robot.

Our approach requires only a small amount of training data to optimize a multitude of parameters from an initial guess of the model. We achieve this by a Bayesian maximum a posteriori (MAP) formulation of the optimization problem that allows for incorporating prior model knowledge *during* the optimization process in order to guide the search efficiently.

Like in many other approaches (e.g., [1], [2]), our robots perceive visual markers attached to their bodies for the calibration. These measured poses are compared to the expected poses of the end-effector according to a parametric model of the hand-eye kinematics. We attach a checkerboard pattern to the end-effector of the robot (see Fig. 1). The robot sweeps its workspace and collects training examples within only a few minutes. The training examples are then used in a MAP [3] optimization process to refine an initial guess of the robot model. Our approach considers measurement uncertainties for the visual marker pose.

All authors are with the Autonomous Intelligent Systems Group, Computer Science Institute VI, University of Bonn, 53113 Bonn, Germany `nassnn at gmail.com, stueckler at ais.uni-bonn.de, behnke at cs.uni-bonn.de`

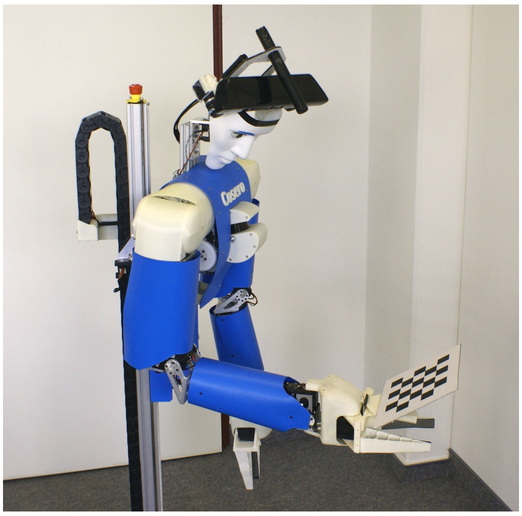

Fig. 1: Our robot gazing at a visual calibration marker attached to its end-effector.

In experiments, we compare the efficiency of our MAP optimization approach with the standard maximum likelihood (ML) method. We evaluate our method using our anthropomorphic robots Dynamaid and Cosero [4] that we constructed from lightweight aluminum parts and off-the-shelf Robotis Dynamixel actuators. We demonstrate that our optimization approach is well suited to calibrate our robots, and that our approach requires only little training effort.

II. RELATED WORK

Research on the identification and calibration of robot kinematics has a long tradition (e.g., see [5] for an early survey). Within the taxonomy of Hollerbach and Wampler [6], our approach belongs to the category of open-loop methods, in which a bundle of link observations is used to optimize the parameters for the mismatch between the expected pose of the links and their measurements.

Recent work focused on the calibration of the kinematic chain with respect to sensors such as cameras and lasers in which task-relevant information is perceived later. Nickels [7], for example, optimized the hand-eye kinematic chain of the NASA Robonaut. This method observes a spherical tool held in the hand of the robot to measure the position of the end-effector in the head-mounted camera. From multiple observations, Nickels improves Denavit-Hartenberg

(DH) parameters [8] and joint gear reductions using a least squares method. He reports an increase of the positioning accuracy in the sensor frame from 13.75 cm to 1.72 cm. Pradeep et al. [1] calibrate sensor frames of cameras and a 3D laser with respect to the kinematic chain of the PR2 robot. Similar to our approach, they observe a checkerboard pattern attached to the end-effector. They tune joint angle offsets and the relative pose of the sensors in the attachment links with a maximum likelihood approach that considers measurement uncertainties. Birbach et al. [2] calibrate the intrinsic parameters of both cameras, their frames relative to the head, an inertial measurement unit in the head, and the joint angle offsets and elasticities of the hand-head chains of the robot Rollin' Justin. They observe a point marker on the wrist with the stereo cameras in the head. The approach incorporates all measured sensor data in a holistic batch calibration using a least squares approach. An initial guess about the parameters is used to provide a search region for the marker and a starting point for the least squares optimization.

Calibration methods are also part of many body schema learning approaches (such as [9], [10], [11], [12]), e.g., to refine initial estimates of the body scheme. In [11], four degrees of freedom (DoF) of a kinematic chain of a humanoid robot are recursively refined in an active learning approach. The approach requires about 600 observations to converge to an accurate estimate of the body schema. Hart and Scassellati [12] identify the kinematic structure of a humanoid robot using Circle Point Analysis [13]. They tune DH parameters and joint gear reductions through least squares optimization.

Our approach differs from the above work by including the full set of DH parameters of the kinematic chain as well as gear reductions within a *maximum a posteriori* optimization framework that considers the uncertainties of the visual measurements together with prior information on the robot kinematics. Our choice of parameter set is minimal, and by incorporating prior knowledge not only as an initial guess but also as prior regularizing information during the optimization process, our method only requires a small amount of training examples.

III. METHOD

We split the calibration process into two parts: an exploration phase in which the robot moves its arm and head to various joint states and scans the visual marker at its end-effector, and an optimization phase in which the robot model is refined with the collected training data.

We implement our calibration method on our service robots Cosero and Dynamaid [4]. The hand-eye kinematic chains of our robots only contain rotational joints (see Fig. 2). However, since we optimize all DH parameters, also linear joints are supported by our method.

A. Exploration Phase

In the exploration phase, the robot collects training examples for the optimization process. We attach a visual marker to the end-effector that the robot perceives in its head-mounted camera.

The robot drives its arm to a series of target poses. It adjusts its gaze using the pan-tilt neck to keep the visual markers visible. For gaze control, we exploit the prior robot model assuming that it is sufficiently accurate to keep the visual marker within the image. One training example comprises the visual measurement of the end-effector pose and the measured joint state of the 9-DoF kinematic chain from head to end-effector.

In order to increase the amount of training samples, the robot moves its head to five different poses for each arm configuration and brings the visual marker to the center and the corners of the image. In each head pose, we add small random values to the target states of the kinematic chain.

B. Optimization Phase

We apply the MAP method to refine the a priori model of the hand-eye kinematic chain of our robots. We calibrate the DH parameters (link length, link offset, link twist, and joint angle) and joint gear reductions of the 9 joints in the kinematic chain. Let $\mathcal{Z} = \{z_i\}_{i=1}^n$ be the training set of n end-effector pose observations in given joint states $\mathcal{X} = \{x_i\}_{i=1}^n$. The MAP estimate $\hat{\theta}_{\mathrm{MAP}}$ of the kinematic parameters is given by

$$\hat{\theta}_{\mathrm{MAP}}(\mathcal{Z}, \mathcal{X}) = \underset{\theta}{\operatorname{argmax}}\; p(\mathcal{Z}|\theta, \mathcal{X})\, p(\theta|\mathcal{X}),$$

where θ denotes the kinematic parameters to optimize. For global robot models, we assume independence of the parameters from the joint states such that $p(\theta|\mathcal{X}) = p(\theta)$.

In contrast to maximum likelihood (i.e., also least squares) formulations, we include prior knowledge about the robot model through the prior $p(\theta)$. We assume this prior to be normal distributed with mean $\overline{\theta}$ and covariance Σ_θ, i.e.,

$$p(\theta) = \eta \exp\left(-\tfrac{1}{2}\left(\theta - \overline{\theta}\right)^T {\Sigma_\theta}^{-1} \left(\theta - \overline{\theta}\right)\right),$$

where the normalization constant η is independent of θ.

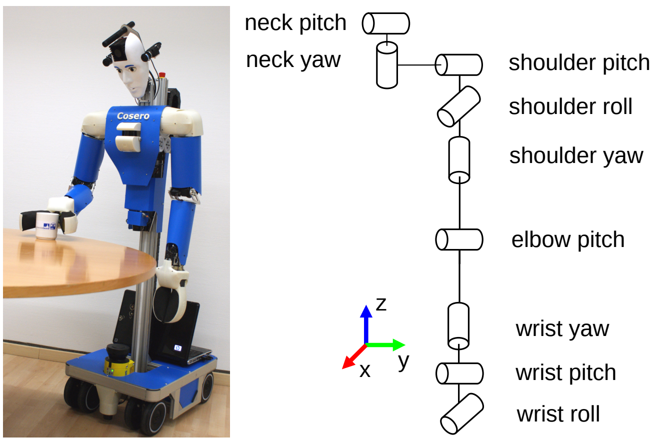

Fig. 2: Schematics of the hand-eye kinematic chains of our robots Cosero and Dynamaid.

For the data likelihood $p(\mathcal{Z}|\theta, \mathcal{X})$ we assume independence between the individual training examples, i.e.

$$p(\mathcal{Z}|\theta, \mathcal{X}) \approx \prod_{i=1}^{n} p(z_i|\theta, x_i).$$

Each observation of the end-effector marker is explained by the noisy measurement of the end-effector pose $z = f(x, \theta) + \epsilon$ as given by the forward kinematics (FK) $f(x, \theta)$. The noise in the measurement is mainly due to inaccuracies of the visual marker detection in the camera images. We model the measurement noise normal distributed $\epsilon \sim \mathcal{N}(0, \Sigma_z)$ with zero mean and covariance Σ_z which is obtained empirically from a series of marker measurements.

Overall, the MAP estimate of the kinematic parameters is

$$\hat{\theta}_{\mathrm{MAP}}(\mathcal{Z}, \mathcal{X}) \approx \underset{\theta}{\operatorname{argmax}}\ \eta_A \exp\left(-A(\theta)\right) \eta_B \exp\left(-B(\theta)\right),$$

where

$$\begin{aligned} A(\theta) &:= \tfrac{1}{2}\left(\theta - \overline{\theta}\right)^T {\Sigma_\theta}^{-1} \left(\theta - \overline{\theta}\right), \\ B(\theta) &:= \sum_{i=1}^{n} \tfrac{1}{2}\left(z_i - f(\theta, x_i)\right)^T {\Sigma_z}^{-1} \left(z_i - f(\theta, x_i)\right), \end{aligned}$$

and η_A and η_B are normalization constants.

Taking the logarithm, we have

$$\hat{\theta}_{\mathrm{MAP}}(\mathcal{Z}, \mathcal{X}) \approx \underset{\theta}{\operatorname{argmax}}\ \log(\eta_A) - A(\theta) + \log(\eta_B) - B(\theta),$$

where we ignore the constant terms and define $L(\theta) := A(\theta) + B(\theta)$ to arrive at

$$\hat{\theta}_{\mathrm{MAP}}(\mathcal{Z}, \mathcal{X}) \approx \underset{\theta}{\operatorname{argmin}}\ L(\theta).$$

We solve this minimization problem by gradient descent on L for the kinematic parameters θ. Since $\Sigma = \Sigma^T$, we obtain the derivative of L for the parameters θ as

$$\nabla_\theta L = {\Sigma_\theta}^{-1}\left(\theta - \overline{\theta}\right) + \sum_{i=1}^{n} \nabla_\theta f(\theta, x_i)^T \Sigma_z^{-1} \left(z_i - f(\theta, x_i)\right).$$

We use resilient backpropagation (Rprop) [14] to implement fast and robust gradient descent. Since the prior introduces a bias for small amounts of training data, we use it to guide the search for the actual parameters and adapt it towards the estimate during the optimization. We replace the a priori model by the current a posteriori model after convergence or a fixed number of iterations and reiterate the whole gradient descent optimization process until convergence or a maximum number of prior replacements.

IV. EXPERIMENTS

We evaluate our approach in simulation and with the real robot. The simulation has been implemented using the Gazebo framework.

Since we do not have ground truth available for the real robot, we evaluate convergence and accuracy properties of our method in simulation. The simulation also allows for conducting time-expensive and otherwise hardware-wearing experiments. We will also report on experiments with the real robot.

TABLE I: Pose error of the a priori models on the simulated test data set.

	avg error		**max error**	
model	**pos [mm]**	**ori [°]**	**pos [mm]**	**ori [°]**
unmodified	0.00	0.00	0.00	0.00
moderate	74.76	15.72	274.55	41.69
strong	296.65	42.26	883.78	108.88

Our optimization approach contains hyper-parameters in the form of the covariance matrices Σ_θ and Σ_z. In simulation experiments, we determined $\Sigma_\theta = 0.1^2\, I$ as a good setting for the prior covariance and that the prior should be replaced with the current a posterior estimate after each 70 iterations. The accuracy of the pose measurements Σ_z has been determined empirically from a series of pose measurements. It is $0.001\,\mathrm{m}^2$ and $0.002\,\mathrm{rad}^2$ for the positional and rotational components, respectively.

A. Simulation Results

We create synthetic test and training data sets in simulation. The test data set and one training data set (denoted as *full*) consist of 10,000 randomly generated joint states each in the full range of joint angles $(-\pi, \pi]$. The joint angles are sampled from a uniform distribution. The *teach-in* training data set contains 370 joint states within the joint angle limits of the kinematic chain that have been teached in with the real robot. For the training data sets, the sampled joint angles are modified with Gaussian noise with a standard deviation of 0.005 rad. In order to simulate visual measurement inaccuracies, we add Gaussian noise to the pose of the end-effector that we obtain from the FK of the actual model. For this purpose, we set the standard deviation of the linear and rotational pose dimensions to 0.012 m and 0.04 rad, respectively. The resulting training sets have an average error of about 11 mm and 0.035 rad between the actual and the measured end-effector pose. The maximum errors range from 36 mm to 48 mm and 0.12 rad to 0.19 rad.

We compare the resulting models of MAP and ML for three different prior assumptions on the models. The first prior (*unmodified*) assumes perfect knowledge about the actual model. For the other models we add *moderate* and *strong* random changes to the actual model that we draw from normal distributions. The end-effector pose errors on the test data set for these a priori models are shown in Table I.

From Table II we can see that the unmodified and moderate a priori models converge to a posteriori models with similarly small errors on the test set when calibration is performed on the full training set with our MAP approach. For both a priori models, the optimization converges to very similar kinematic chains. This results in the same error values for the percision given in Table II. The parameters differ at most by 0.04° in orientation, 0.19 mm in lengths, and 0.03 % in gear ratios from the actual model.

TABLE II: Pose error of the MAP models in simulation when trained on the full training data set.

	avg error		max error	
model	**pos [mm]**	**ori [°]**	**pos [mm]**	**ori [°]**
unmodified	0.35	0.07	0.92	0.18
moderate	0.35	0.07	0.92	0.18

TABLE III: Pose error of the ML and our MAP approach in simulation when trained on the teach-in training data set.

	avg error		max error	
model	**pos [mm]**	**ori [°]**	**pos [mm]**	**ori [°]**
moderate, ML	9.49	1.55	35.74	4.35
moderate, MAP	9.42	1.54	35.66	4.33
strong, ML	9.47	1.55	35.58	4.35
strong, MAP	9.42	1.53	35.98	4.30

Table III shows the results of the ML and MAP methods on the teach-in training dataset. Both approaches converge to similar models. The individual deviations of the parameters from the actual model for the moderately noisy prior can be seen in Fig. 3, 4, and 5.

B. *Results with the Real Robot*

We apply our method to calibrate the hand-eye kinematics of our service robot Dynamaid. We compare our MAP calibration method with the ML approach. Afterwards, we demonstrate the applicability of our calibration in an experiment in which the robot points to specific positions with a laser pointer. We also measure the repeatability accuracy of the robot mechanism and the pose estimation method to relate the uncertainty of the mechanism with the calibration results.

1) Comparison between ML and MAP Calibration: For the evaluation of our calibration method with our robot Dynamaid, we generated a test and a training data set. The training set consists of 597 examples that originate from 128 teach-in poses. The testing set contains 327 joint states from 71 different teach-in poses. In Table IV we give the average and maximum pose errors of the prior kinematic model of Dynamaid on these datasets. For the optimization

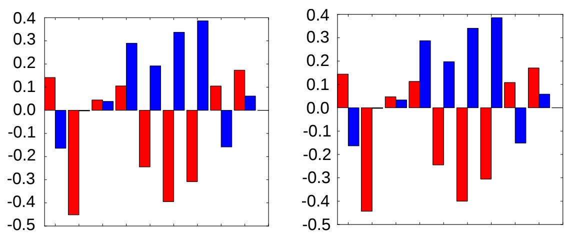

(a) ML (b) MAP

Fig. 3: Joint angle (red) and link twist (blue) deviations of the trained models from the actual model on the teach-in dataset in °. The parameters are arranged from neck pitch (left) to wrist roll (right).

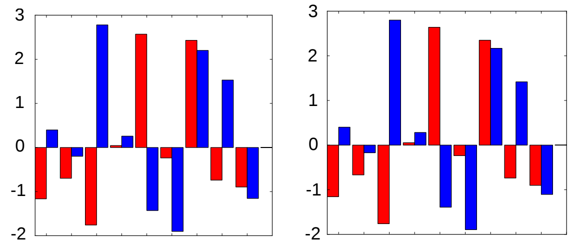

(a) ML (b) MAP

Fig. 4: Link length (red) and link offset (blue) deviations of the trained models from the actual model on the teach-in dataset in mm. The parameters are arranged from neck pitch (left) to wrist roll (right).

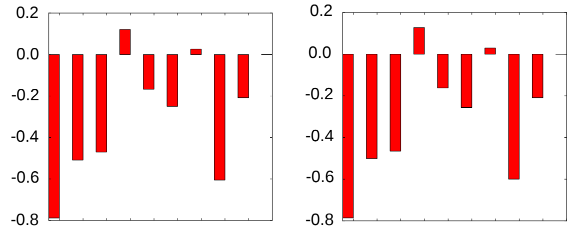

(a) ML (b) MAP

Fig. 5: Gear reduction deviations of the trained models from the actual model on the teach-in dataset in %. The parameters are arranged from neck pitch (left) to wrist roll (right).

of Dynamaid, we found the following variances for the prior parameter distribution: $0.1\,\mathrm{rad}^2$ for the joint angles, $0.01\,\mathrm{rad}^2$ for the link twists, $0.01\,\mathrm{m}^2$ for the link lengths and offsets, and 0.001 for the gear reductions.

For training the models in Table IV, we used the following convergence criteria: We consider gradient descent converged, if the deviations from the average pose error within the last 20 iterations do not exceed $10^{-5}\,m$ in position and $10^{-5}\,rad$ in orientation. The optimization is also aborted, if a maximum number of 3000 iterations has been reached. In the MAP approach, we replace the a priori model after convergence of gradient descent, and restart the optimization. This process is reiterated 40 times or until gradient descent converges within the first 30 iterations after the replacement of the a priori model. On our training set, MAP requires a total amount of 917 iterations and the prior is replaced 15 times. The number of iterations for ML is 407.

The results in Table IV demonstrate that our MAP method could significantly improve the kinematic model of Dynamaid. Compared to the ML approach, our method provides better generalization on the test data set. The good performance of ML on the training data indicates that it overfits the training set for the price of a good result on the test set.

In Fig. 6, 7, and 8 we compare the parameters of the learned ML and MAP models with the a priori model. It can be seen that the approaches tend to different solutions. While ML adjusts gear reduction parameters stronger than our MAP approach, MAP stronger adapts link length and offset parameters. This effect is explained by the a priori

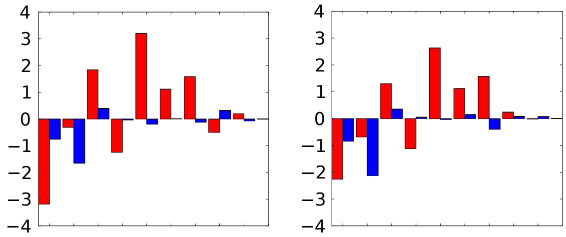

(a) ML (b) MAP

Fig. 6: Joint angle (red) and link twist (blue) deviations of trained models from the actual model on the teach-in dataset in °. The parameters are arranged from neck pitch (left) to wrist roll (right).

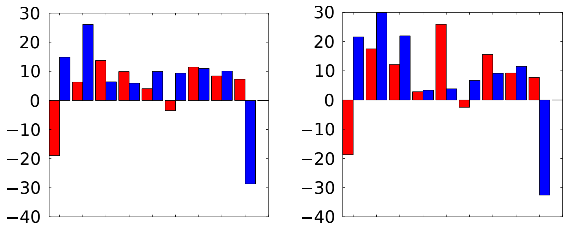

(a) ML (b) MAP

Fig. 7: Link length (red) and link offset (blue) deviations of trained models from the actual model on the teach-in dataset in mm. The parameters are arranged from neck pitch (left) to wrist roll (right).

knowledge about the precision of the initial guess of the gear reductions that we incorporate in our MAP approach. In ML, no such information is available, and the optimization method can run into a different local minimum in this high-dimensional parameter space.

We further evaluate the performance of MAP and ML through cross-validation and local training. We merge training and test set into a union set with 924 samples.

In order to evaluate the quality of the approaches for training on specific amounts of data that are randomly distributed in the assessed workspace of the robot, we conduct a variant of cross-validation on the union set. The union set is split randomly into n partitions which are individually used for training. The remaining samples to each partition are used as a test set for this partition. For each partition

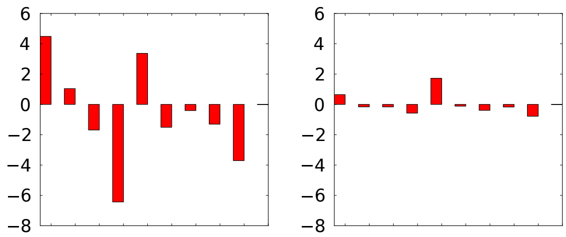

(a) ML (b) MAP

Fig. 8: Gear reduction deviations of trained models from the actual model on the teach-in dataset in %. The parameters are arranged from neck pitch (left) to wrist roll (right).

TABLE IV: Pose error of a priori model, ML, and our MAP approach when calibrating Dynamaid.

	avg error		max error	
model, set	**pos [mm]**	**ori [°]**	**pos [mm]**	**ori [°]**
prior, training	36.45	7.37	73.94	14.89
MAP, training	13.76	3.02	**29.66**	10.30
ML, training	**13.33**	**2.84**	30.65	**10.09**
prior, test	45.58	8.30	85.99	13.48
MAP, test	**14.28**	**3.00**	**50.98**	9.87
ML, test	16.93	3.10	66.79	**9.70**

size, we measure the average and maximum pose error of the learned models for the n partitions on the corresponding test sets. Table V shows results for various partition sizes. If the methods did not converge on specific partitions, we also show results without these partitions (marked by *). It can be seen, that ML and our MAP method yield quite similar average errors on the test sets, while MAP has lower maximum errors in most cases. For little training data (23 samples), MAP obtained smaller average errors. The results indicate better generalization by the MAP approach, especially for smaller training datasets.

In a second set of experiments, we assessed how well the learned models of both approaches generalize to untrained areas of the workspace. We randomly chose n samples from the union set, and trained models on the $k = 924/n$ closest samples in the space of joint angles. The remaining samples are used as test sets. The results of our experiments for different n are shown in Table VI. In these experiments, MAP clearly outperforms ML with regard to average and maximum error for $n \geq 10$, i.e., $k \leq 92$ on the test sets. For smaller n, MAP still yields smaller maximum errors on the test sets.

2) Model Application Experiment: We demonstrate the benefit of a model calibrated with our approach in the following experiment. We attach a laser pointer to the end-effector of the robot with which the robot shall point at the corners of a checkerboard in front of it. The checkerboard is placed at several poses relative to the robot, and the positioning accuracy between pointed and target positions is measured. The robot holds its end-effector in a distance of about 25 cm to the checkerboard and points to the four outer corners of the checkerboard. Since we are interested in the best available MAP model for this experiment, we train it on the union set. The model consists of the DH parameters of the kinemaitic chain, which represent the FK of the arm. With this information, the inverse kinematics (IK) of a target pose is determined.

First, we determine the repeatability accuracy of the pointing process. We measure the repeatability accuracy of the robot mechanism, if the pose of the checkerboard is estimated only once before the robot repeatedly points to a location. If the pose is measured at each trial, we also include the accuracy of the visual pose estimation process. Table VII summarizes the repeatability accuracies in both cases. The accuracy can be used as a reference, which accuracy is

TABLE V: Cross-validation (n partitions, k samples) of ML and MAP with the union set (924 samples in total). 40 (23)*: omitted 2 MAP runs (exceeded 40 prior replacements), omitted 5 ML runs (exceeded 3000 iterations).

		avg error		max error	
set, n (k)	**method**	**pos [mm]**	**ori [°]**	**pos [mm]**	**ori [°]**
training, 2 (426)	MAP	13.30	2.94	**39.34**	**10.16**
	ML	**13.15**	**2.75**	41.75	10.54
training, 5 (184)	MAP	13.15	2.93	**35.54**	10.54
	ML	**12.93**	**2.72**	36.35	**10.48**
training, 10 (92)	MAP	12.84	2.88	37.08	10.38
	ML	**12.63**	**2.69**	**35.00**	**9.94**
training, 20 (46)	MAP	12.30	2.80	32.79	**9.18**
	ML	**12.12**	**2.63**	**31.50**	14.00
training, 40 (23)	MAP	**13.32**	**4.20**	**197.44**	**118.54**
	ML	20.76	8.92	361.90	176.71
training, 40 (23)*	MAP	10.98	2.60	29.52	**9.53**
	ML	**10.76**	**2.44**	**27.47**	10.88
test, 2 (426)	MAP	13.56	2.96	**38.76**	**10.71**
	ML	**13.48**	**2.79**	42.72	11.35
test, 5 (184)	MAP	13.73	3.01	**45.19**	**11.16**
	ML	**13.62**	**2.84**	49.20	12.55
test, 10 (92)	MAP	14.05	3.04	**49.67**	**12.58**
	ML	**14.01**	**2.91**	57.53	13.92
test, 20 (46)	MAP	**14.60**	3.11	**55.17**	**13.25**
	ML	14.85	**3.04**	80.15	14.36
test, 40 (23)	MAP	**20.08**	**5.47**	543.33	**179.96**
	ML	28.44	10.17	**454.69**	179.98
test, 40 (23)*	MAP	**16.61**	**3.31**	**75.69**	**13.85**
	ML	17.35	3.51	111.34	18.60

achievable at all with the robot mechanism and the sensor setup in this experiment.

Finally, we report on the accuracies of the laser pointing experiment with the a priori model and our MAP estimate in Table VIII. Our proposed method clearly improves the accuracy in this experiment.

C. Local Models

We evaluate if it is beneficial to train local models for different parts of the configuration space. We choose the smaller test set of 327 teach-in joint states in this experiment. For each joint state, we train a localized MAP model by weighting each training example by the distance in configuration space using a Gaussian kernel function.

For a variance of $0.3\,\mathrm{rad}^2$ for the Gaussian kernel, the average errors of the MAP models is higher than the global model (see horizontal lines in Fig. 9). The error increases with the Euclidean distance from the reference joint state. If we set the variance of the Gaussian kernel to $0.1\,\mathrm{rad}^2$, the average error over the local models slightly improves on the global model (see Fig. 9). From about $30°$ distance in joint angles from the reference state, the average error of the local models grows beyond the error of the global model. Note, that only few local models exhibit the large maximal errors that can be seen in Fig. 9. We conclude that localized learning is prone to overfitting and does not provide a significant improvement for our robot platform.

TABLE VI: Results of ML and MAP for training on k nearest neighbors of n samples from the union set (924 samples in total). 20 (46)*: omitted 4 MAP runs (exceeded 40 prior replacements), omitted 1 ML run (exceeded 3000 iterations). 40 (23)*: omitted 14 MAP runs (13 exceeded 40 prior replacements, 1 exceeded 3000 iterations), omitted 8 ML runs (exceeded 3000 iterations).

		avg error		max error	
set, n (k)	**method**	**pos [mm]**	**ori [°]**	**pos [mm]**	**ori [°]**
training, 2 (426)	MAP	13.18	2.97	**37.96**	**10.81**
	ML	**12.81**	**2.73**	39.03	11.45
training, 5 (184)	MAP	12.80	2.73	**31.09**	9.84
	ML	**12.30**	**2.52**	32.60	**9.31**
training, 10 (92)	MAP	12.43	2.89	**31.71**	10.00
	ML	**11.88**	**2.66**	32.42	**9.22**
training, 20 (46)	MAP	**10.61**	**2.68**	**26.45**	**9.58**
	ML	17.21	8.75	326.25	168.77
training, 20 (46)*	MAP	10.59	2.57	**26.45**	9.58
	ML	**10.19**	**2.42**	27.00	**9.08**
training, 40 (23)	MAP	**11.58**	**3.79**	**191.09**	**90.86**
	ML	26.46	14.89	445.61	179.22
training, 40 (23)*	MAP	9.12	2.65	25.75	**10.77**
	ML	**8.35**	**2.18**	**24.64**	10.92
test, 2 (462)	MAP	16.12	3.54	**53.79**	**10.26**
	ML	**15.41**	**3.17**	59.91	10.32
test, 5 (184)	MAP	**17.40**	3.66	**76.96**	**11.43**
	ML	18.27	**3.34**	81.82	14.12
test, 10 (92)	MAP	**19.72**	**3.73**	**86.37**	**12.54**
	ML	21.59	3.82	104.85	16.19
test, 20 (46)	MAP	**24.14**	**4.15**	**122.20**	**15.37**
	ML	40.41	11.27	598.50	179.98
test, 20 (46)*	MAP	**23.86**	**4.06**	**122.20**	**15.00**
	ML	28.58	4.76	141.39	22.10
test, 40 (23)	MAP	**37.49**	**7.95**	**560.38**	**179.87**
	ML	69.25	19.92	1010.24	180.02
test, 40 (23)*	MAP	**30.28**	**4.73**	**152.27**	**21.05**
	ML	38.22	6.95	221.89	47.49

TABLE VII: Repeatability accuracy in the laser pointing experiment.

mode	**avg error [mm]**	**max error [mm]**
estimate pose once	3.5	8
reestimate pose each time	5.6	12

TABLE VIII: Position errors in the laser pointing experiment.

model	**avg error [mm]**	**max error [mm]**	**std. dev. [mm]**
prior	34.4	60	12.3
MAP	22.2	40	10.3

V. DISCUSSION

Our simulation experiments demonstrate that the amount of training samples and their distribution determine the

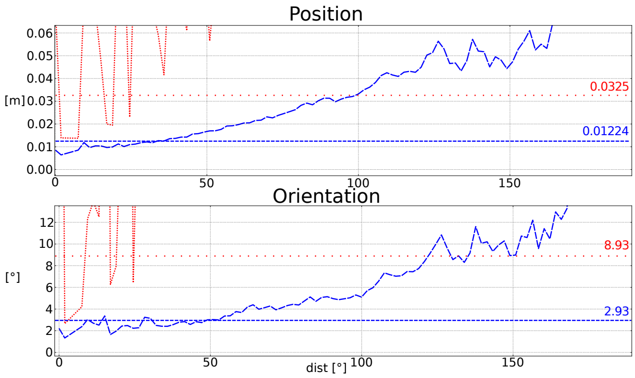

Fig. 9: Localized Learning. Top: Average (blue, dashed) and maximum (red, dotted) position error over all 327 local models. We also show average (blue, dashed) and maximum (red, dotted) errors of the global models for reference as horizontal lines. Bottom: Average and maximum orientation errors (same color coding as for position).

quality of an a posteriori model. For 10,000 joint states distributed in $(-\pi, \pi]$ the results starting from different a priori models match each other for the same training set. It is to be expected that an increase in data leads to better results. If the generated samples are only taken within the bounds of the allowed joint margins, then the outcome of the calibration differs more from the ground truth. Still, both kinds of methods (ML and MAP) yield similar results in simulation, in contrast to the results with the real robot.

For the real robot, the training reduces error nearly to a third of the a priori model. It is interesting which expected variances have led to these results. The variance for the joint angles is the same that showed good results in simulation. It is much bigger than the variances of the other DH parameters. An explanation is, that this parameter encompasses possible servo offsets. The other DH parameters are better measurable manually or can be taken from a CAD model. In most joints, the servos are directly attached to the links without an additional transmission. If we assume good manufacturing standards of the servos, this explains the low values for the variances of the gear reductions.

Our experiments with the real robot demonstrate, that our MAP approach generalizes better than the ML method if only few training examples are available. By providing different variances for parameters in the a priori model, MAP allows for incorporating prior knowledge about the uncertainty of parameters and their scale. In contrast, ML treats all parameters equally, and hence may yield different solutions.

The increase in accuracy in the laser experiment is only one third compared to the a priori model which is less than the pose accuracy of the visual marker. This can be explained by the extension of the kinematic chain by approximately 25 cm and the conversion of orientation error to position errors. The repeatability accuracy is also remarkably good for a low-weight robot with an armlength of an adult human.

Localized learning seems not to be beneficial for our robot. Although there are small local dependencies, the error raises with distance, and training effort would rise significantly to cover the complete workspace such that the local models could have an advantage, compared to a global one.

VI. Conclusion

In this paper, we propose a Bayesian maximum a posterior (MAP) approach to the calibration of the hand-eye kinematics of an anthropomorphic robot. We include prior knowledge on the robot model during the optimization to efficiently guide the search for model parameters and to avoid overfitting problems. In experiments in simulation and with a real robot, we could demonstrate, that our method converges to accurate models. On our robot Dynamaid, our MAP approach generalizes better from few training examples than a maximum likelihood (ML) approach that has often been used in previous work. By this, our method can cope with many parameters but only requires little training effort.

In future work we will investigate active learning of the robot kinematics. In such an approach, the robot explores informative joint states, which could be implemented based on information gain in our probabilistic framework. In order to adapt the hand-eye kinematics continuously on-line, we could enhance the robot with smaller markers such as LEDs. This approach could be extended for tool-use, if the robot perceives the pose or the tip of the tool.

References

[1] V. Pradeep, K. Konolige, and E. Berger, "Calibrating a multi-arm multi-sensor robot: A bundle adjustment approach," in *Proc. of the Int. Symposium on Experimental Robotics (ISER)*, 2010.

[2] O. Birbach, B. Bäuml, and U. Frese, "Automatic and self-contained calibration of a multi-sensorial humanoid's upper body," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Saint Paul, USA, 2012.

[3] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006.

[4] J. Stückler, D. Holz, and S. Behnke, "RoboCup@Home: Demonstrating everyday manipulation skills in RoboCup@Home," *IEEE Robotics & Automation Magazine*, 2012.

[5] Z. Roth, B. W. Mooring, and B. Ravani, "An overview of robot calibration," *IEEE J. of Robotics and Automation*, vol. 3, no. 5, 1987.

[6] J. M. Hollerbach and C. W. Wampler, "The calibration index and taxonomy for robot kinematic calibration methods," *International Journal of Robotics Research*, vol. 15, no. 6, pp. 573–591, 1996.

[7] K. M. Nickels, "Hand-eye calibration for robonaut," NASA Johnson Space Center, Tech. Rep., 2003.

[8] J. Denavit and R. S. Hartenberg, "A kinematic notation for lower-pair mechanisms based on matrices." *Trans. of the ASME. Journal of Applied Mechanics*, vol. 22, pp. 215–221, 1955.

[9] M. Hersch, E. Sauser, and A. Billard, "Online learning of the body schema," *International Journal of Humanoid Robotics*, vol. 5, no. 2, pp. 161–181, 2008.

[10] J. Sturm, C. Plagemann, and W. Burgard, "Body schema learning for robotic manipulators from visual self-perception," *Journal of Physiology-Paris*, vol. 103, no. 3-5, pp. 220–231, Sept. 2009.

[11] R. Cantin-Martinez, M. Lopes, and L. Montesano, "Body schema acquisition through active learning," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Alaska, USA, 2010.

[12] J. W. Hart and B. Scassellati, "A robotic model of the ecological self," in *Proceedings of the IEEE/RSJ International Conference on Humanoid Robots (Humanoids)*, Bled, Slovenia, 2011.

[13] B. Mooring, M. Driels, and Z. Roth, *Fundamentals of Manipulator Calibration*. New York, NY, USA: John Wiley & Sons, Inc., 1991.

[14] M. Riedmiller, "Rprop-description and implementation details," University of Karlsruhe, Tech. Rep., 1994.