
Neural Inverse Kinematics

Raphael Bensadoun¹ Shir Gur¹ Nitsan Blau¹ Tom Shenkar¹ Lior Wolf¹

Abstract

Inverse kinematic (IK) methods recover the parameters of the joints, given the desired position of selected elements in the kinematic chain. While the problem is well-defined and low-dimensional, it has to be solved rapidly, accounting for multiple possible solutions. In this work, we propose a neural IK method that employs the hierarchical structure of the problem to sequentially sample valid joint angles conditioned on the desired position and on the preceding joints along the chain. In our solution, a hypernetwork f recovers the parameters of multiple primary networks g_1, g_2, \dots, g_N , where N is the number of joints, such that each g_i outputs a distribution of possible joint angles, and is conditioned on the sampled values obtained from the previous primary networks $g_j, j < i$. The hypernetwork can be trained on readily available pairs of matching joint angles and positions, without observing multiple solutions. At test time, a high-variance joint distribution is presented, by sampling sequentially from the primary networks. We demonstrate the advantage of the proposed method both in comparison to other IK methods for isolated instances of IK and with regard to following the path of the end effector in Cartesian space.

1. Introduction

Given the joint angles, the position and orientation of the robot’s end-effector can be readily computed in a process called forward-kinematics. However, robotic planning and controls require mapping in the other direction, *i.e.*, from the end-effector’s Cartesian space coordinates to the joint positions. This inverse mapping is called Inverse Kinematic (IK). It is a nonlinear problem that often has multiple solutions (Craig, 2009).

¹Mentee Robotics. Correspondence to: Raphael Bensadoun <raphael@menteebot.com>, Shir Gur <shir@menteebot.com>.

For simple kinematic chains without much ambiguity, one can obtain analytical solutions for the IK problem. However, for the chain type that appears in robotic arms and other complex robots, one has to rely on numerical methods. In this work, we propose what is, as far as we can ascertain, the first deep learning solution that allows for multiple solutions.

Given a certain kinematic chain, one can readily obtain a training set consisting of pairs (x, y) of end-effector positions x and matching joint angles y , by sampling the latter and computing the former with forward kinematics. This straightforward way of obtaining the training set does not reflect the possibility of multiple solutions y , given the coordinates and orientation of x . In our framework, we employ a variational approach to the problem and sample, at inference time, from a distribution P_x that is conditioned on the vector of Cartesian-space specifications x .

Due to the structure of the kinematic chain, the IK problem can be seen as a hierarchical problem. Typically, the angle of the joint that is attached to the end-effector is uniquely determined by the location of all previous joints and the specifications in x . The previous joint may have multiple solutions given the joint angles that precede it. In general, as we move along the kinematic chain, from the fixed attachment point to the end effector, the number of possible configurations decreases. While this is true for any order in which we sequentially fix one joint, the kinematic chain is often equipped with a natural order, in which the first joints typically cause larger motions in Cartesian space.

In our model, this hierarchy is manifested by a sequential sampling of the joints from the distribution P_x . Namely, we parametrize P_x as a sequential process in which the joints are sampled one by one and the sampling of each joint is conditioned on the values obtained for the previous joints.

The IK problem is often characterized by a discontinuous solution space. While for a given y , we can expect to see multiple solutions y' that are close in the configuration space, there may be other solutions that obtain the desired position and orientation in x using a completely different configuration. Our model addresses this by employing Gaussian Mixture Models (GMMs) during sequential sampling.

The sequential sampling process, therefore, takes the following form. The distribution of the first joint is given as a GMM. The parameters of this GMM are computed by a

neural network, given the desired position and orientation of the end effector x . The angle of the first joint $y^{(1)}$ (this is the joint that is the most distant from the end effector) is sampled by this GMM and the set of possible configurations for the remaining joints is reduced. A second GMM is then inferred in a way that is conditioned both on x and on $y^{(1)}$. The second joint $y^{(2)}$ is sampled, and the process is repeated until all N joints are obtained.

In the framework we propose, the parameters of each GMM are obtained by a neural network that receives the preceding joint locations as inputs. Conditioning on the input x is obtained using a hypernetwork scheme, such that the parameters (weight matrices and biases) of the networks that provide the GMMs change dynamically, depending on x . This solution allows us to model the problem in a natural way, separating the conditioning on x from the conditioning on the sampled values.

Using the proposed neural IK solution, which we call IKNet, we present a path following method for recovering a sequence of joint location vectors given a sequence of smoothly varying end effector positions. This method runs online, such that at each time point the sampling of the joint angle depends on the preceding joints along the kinematic chain and on the angle of the same joint in the previous time step. The latter consideration ensures smoothness of the resulting path.

Our experiments demonstrate that IKNet outperforms a wide variety of IK methods, both optimization-based and learning-based. In the path following problem, our method generates multiple solutions, each more accurate and more stable than the single solution of the best baseline method. Additionally, we show that our probabilistic method displays robustness to noisy dimensions in the kinematic chain. Moreover, a relatively small number of examples is sufficient to finetune a trained model to perform well on a similar but unseen kinematic chain. Lastly, the representation learned by IKNet seems to help in learning other tasks.

2. Related Work

IK methods can be divided into analytical and numerical methods. Analytical methods (Raghavan & Roth, 1993; Diankov, 2010) provide a globally optimal solution, and in many situations multiple solutions, in an efficient and reliable way. However, the availability of analytical solutions is limited to models of limited complexity. Iterative (numerical) IK methods (Buss, 2004) update the vector of joint angular parameters through nonlinear optimization until convergence. A particular case is that of steerable needles, for which an optimization-based IK method was presented by Sears & Dupont (2007).

Attempts to apply machine learning methods to IK include

the application of One-Class SVM (Schölkopf et al., 2002) by Bócsi et al. (2011). D’Souza et al. (2001) applied locally weighted projection regression (Vijayakumar & Schaal, 2000; Klanke et al., 2008) to this problem. In another work, De Angulo & Torras (2008) have addressed IK with Parametrized Self Organizing Maps (Walter & Ritter, 1996).

A straightforward application of neural networks and other regression techniques to map between the vector x of the end effector’s position to the vector y of joint locations (El-Sherbiny et al., 2018; Duka, 2014). Such methods fail to model the entire solution space for a given x and are also, as we show empirically, less accurate. (Csiszar et al., 2017) propose to heuristically divide the dataset to reduce ambiguities.

An important challenge for learning-based IK methods is to perform modeling online, with a given setup, and not based on a large training set (Rolf et al., 2010a;b; Baranes & Oudeyer, 2013). Such methods employ frameworks such as the one by Moulin-Frier et al. (2014). In Sec. 5.4 we experiment with finetuning an existing model to model, with a relatively few samples, a model that deviates from it.

IK-based learning methods have been applied for computer graphics purposes, in order to obtain a more natural motion (Grochow et al., 2004; Huang et al., 2017). Learning in such cases is based on motion capture and other sources of data. In contrast to these methods, our method is aimed at finding all possible solutions for a given kinematic chain and the data we employ is synthetic data generated by this chain.

In order to perform sequential sampling, we employ a series of networks g_k , which are all conditioned on the pose vector x . To this end, we employ a hypernetwork (Ha et al., 2016). The hypernetwork scheme has two components: a primary network g , which outputs the computation result, and a hypernetwork f , which is used for conditioning on some input. The weights of network g are not learned directly. Instead, they are provided as the output of network f . Therefore, the weights of g are dynamic and vary based on the input of f . Hypernetworks have been used for RNNs since their inception, but we are not aware of any other application for a series of primary networks.

While there exist alternative ways of conditioning, such as passing x as an additional input to each network g_k , hypernetworks provide a modular solution, in which the capacity of the conditioning network can be increased, while employing relatively shallow primary networks (Galanti & Wolf, 2020). This is useful in our IK framework, in which data is generated synthetically. In our experiments, we present a baseline that conditions on the end effector position without using a hypernetwork and demonstrate the advantage of employing hypernetwork-based conditioning in the context

of IK. The hypernetwork structure that we employ performs hierarchical sampling. In the context of image generation, [Bensadoun et al. \(2021\)](#) have combined a hypernetwork with a hierarchical sampling to obtain multiple valid solutions.

3. Method

In this section, we describe a method for learning the mapping from an end-effector position \mathbf{x} to the distribution of the joint angles P_x , thereby enabling the sampling of joint angles $\tilde{\mathbf{y}} \in \mathbb{R}^N \sim P_x$, such that applying $\tilde{\mathbf{y}}$ to the N -joint kinematic model results in the end-effector at position \mathbf{x} .

At training time, we are given a set of pairs of vectors $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}$, in which every vector $\mathbf{x}^{(i)}$ is matched with a single vector $\mathbf{y}^{(i)} \in S_x$, where S_x is a set of possible matching y-space vectors for the vector \mathbf{x} . Note that our formulation allows for the existence of indices i, j such that $\mathbf{x}^{(i)} = \mathbf{x}^{(j)}$, but $\mathbf{y}^{(i)} \neq \mathbf{y}^{(j)}$. To index the vector $\mathbf{y}^{(i)} \in \mathbb{R}^N$, we use the superscript $\mathbf{y}^{(i)}$ to denote the i -th sample, and the subscript \mathbf{y}_k to denote the k -th joint.

In the IK problem, x is the position of the robot’s end-effector, and $\mathbf{y} \in \mathbb{R}^N$ is the vector of joint angles, where N is the number of joints in the kinematic chain. The end effector position can include its location or both its location and orientation. In the latter case, the number of plausible IK solutions decreases.

Every valid Cartesian position \mathbf{x} has one or more matching joint configurations \mathbf{y} , which collectively form the set S_x . Since the forward mapping is one to one, $S_x \cap S'_x = \emptyset$ for $\mathbf{x} \neq \mathbf{x}'$. However, our method does not employ this fact.

Our goal is to learn to map every vector \mathbf{x} to a conditional distribution P_x , such that the likelihood of every $\mathbf{y} \in S_x$ is high and, conversely, low for $\mathbf{y} \notin S_x$.

Due to the hierarchical structure of the kinematic chain, we parametrize P_x such that sampling a vector $\tilde{\mathbf{y}}$ from this distribution is done sequentially

$$\tilde{\mathbf{y}}_1 \sim p_x^1 := p_x(\mathbf{y}_1) \quad (1)$$

$$\tilde{\mathbf{y}}_2 \sim p_x^2 := p_x(\mathbf{y}_2 | \tilde{\mathbf{y}}_1) \quad (2)$$

...

$$\tilde{\mathbf{y}}_k \sim p_x^k := p_x(\mathbf{y}_k | \tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_{k-1}) \quad (3)$$

... and

$$P_x := [p_x^1, p_x^2, \dots, p_x^N]. \quad (4)$$

Namely, the first element $\tilde{\mathbf{y}}^1$ is sampled first from the distribution $p_x(\mathbf{y}^1)$, then the second element, $\tilde{\mathbf{y}}^2$, in a way that is conditioned on the first, $\tilde{\mathbf{y}}^1$, using the distribution $p_x(\mathbf{y}_2 | \tilde{\mathbf{y}}_1)$, and so on. This way of sampling is natural for kinematic chains, as mentioned in Sec. 1.

Specifically, we model each part of the P_x distribution

p_x^1, \dots, p_x^N as a Gaussian mixture model (GMM). This way, the distribution is able to capture sets S_x that have a discontinuous shape, with multiple regions. For m GMM components, the distribution p_x^k is thus parameterized by a vector $\mathbf{m}_x^k \in \mathbb{R}^{3m}$ capturing the mean, variance, and mixture coefficient of each of the m components.

Let N be the dimensionality of the vectors \mathbf{y} . The method employs a hypernet structure, in which the hypernet f maps its input \mathbf{x} to the set of parameters of N primary networks g_1, g_2, \dots, g_N . The mapping between \mathbf{x} and P_x and the sampling from this distribution is formulated as follows:

$$[\theta_1, \theta_2, \dots, \theta_N] = f(\mathbf{x}) \quad (5)$$

$$\mathbf{m}_x^1 = g_1(\theta_1) \quad (6)$$

$$\tilde{\mathbf{y}}_1 \sim p_x^1 \quad (7)$$

$$\mathbf{m}_x^2 = g_2(\theta_2, \tilde{\mathbf{y}}_1) \quad (8)$$

$$\tilde{\mathbf{y}}_2 \sim p_x^2 \quad (9)$$

$$\mathbf{m}_x^3 = g_3(\theta_3, \tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2) \quad (10)$$

$$\tilde{\mathbf{y}}_3 \sim p_x^3 \quad (11)$$

...

$$\mathbf{m}_x^N = g_N(\theta_N, \tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2, \dots, \tilde{\mathbf{y}}_{N-1}) \quad (12)$$

$$\tilde{\mathbf{y}}_N \sim p_x^N \quad (13)$$

where for the primary networks g_k , the network parameters are mentioned explicitly as the first input parameter, and p_x^k is the GMM distribution with the parameters \mathbf{m}_x^k . Given an input sample \mathbf{x} , the hypernet f returns, in Eq. 5, the parameters of the N primary networks. Then, in Eq. 6, the GMM parameters of the distribution for the first element of $\tilde{\mathbf{y}}$ are obtained through the primary network g_1 . Subsequently, in Eq. 7, a value is obtained from this distribution. Conditioned on the sampled value, the parameters of the GMM of the second element in $\tilde{\mathbf{y}}$ are obtained (Eq. 8), and this value is sampled (Eq. 9). The process continues until all N values of the vector $\tilde{\mathbf{y}}$ are obtained, each conditioned on the previous elements.

Fig. 1 illustrates the proposed methods, where given a query input \mathbf{x} for f , as a desired end-effector position, the network produced N GMMs, as the number of joints, that we can sample from. Each sampled solution will be mapped to the end-effector position using forward kinematics.

3.1. Training

During training, we learn only the parameters of the hypernetwork f . The parameters of the primary networks g_k are given by f and change based on the input to this network. In the training procedure, we employ a teacher forcing scheme, in which the values of the training sample $\mathbf{y}^{(i)}$ are employed instead of sampling.

The loss term we employ during training maximizes the log

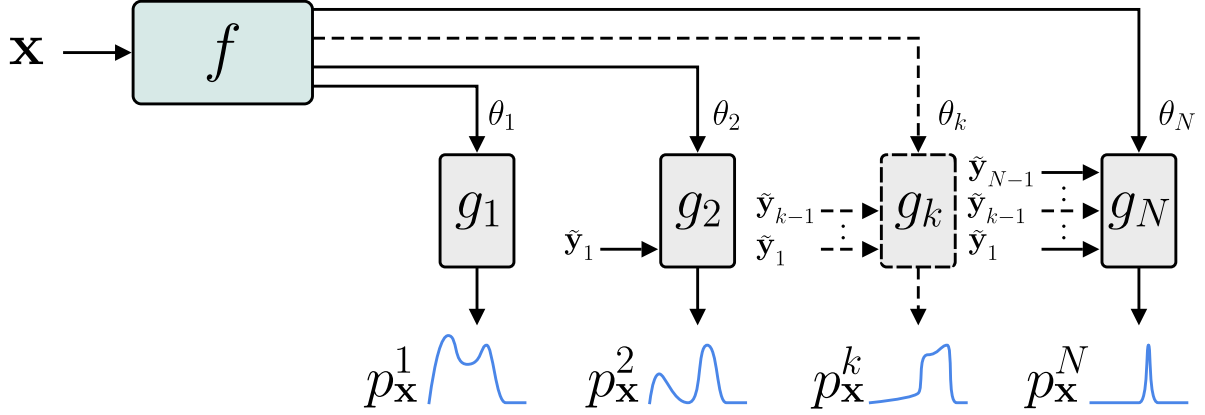


Figure 1. Illustration of the proposed method. Given a query input \mathbf{x} , which is the desired end-effector position, the hypernet f maps its input to the set of parameters of N primary networks g_1, g_2, \dots, g_N , where N is the number of joints in the kinematic chain. The output $\tilde{\mathbf{y}}_k$ of g_k is a GMM $p_{\mathbf{x}}^k$, modeling the solutions distribution for joint k , and the input for each g_k are all the previously sampled $\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_{k-1}$.

likelihood of the training samples

$$\mathcal{L} = - \sum_i \log \prod_{k=1}^N p_{\mathbf{x}^{(i)}}^k \left(\mathbf{y}_k^{(i)} | \mathbf{y}_1^{(i)}, \dots, \mathbf{y}_{k-1}^{(i)} \right) \quad (14)$$

where $\mathbf{y}_1^{(i)}, \dots, \mathbf{y}_{k-1}^{(i)}$ are the ground truth values, *i.e.*, no sampling takes place, and instead of sequential sampling as in Eq. 7,9,11, and 13, the values of $\mathbf{y}^{(i)}$ are used. The training sample $\mathbf{x}^{(i)}$ is manifested through the distributions $p_{\mathbf{x}^{(i)}}^k(\mathbf{y}_k | \dots)$, which are based on the GMM parameters $\mathbf{m}_{\mathbf{x}^{(i)}}^k$ (Eq. 6,8,10, and 12).

3.2. Architecture

The network f is a 4-layer fully-connected network. Each linear layer has a dimension of 1024, with ReLU and batch-norm following each layer. The last layer of f is followed by N projection layers that map the last dimension of 1024 to the vector of weights, θ_k , for each network g_k .

The networks g_k take as an input the weights produced by f and a sequence of joint angles $\mathbf{y}_1, \dots, \mathbf{y}_{k-1}$. Each network is composed of three linear layers with a hidden dimension of 256, and ReLU activation between the layers. The output is a vector of $3m$ elements, where a subset of m elements denote the prior for each GMM component.

In order not to explicitly select an optimal value for the parameter m , we set it at a very high value of $m = 50$, and make sure that the vector of priors is sparse. Specifically, the relevant (*i.e.*, not mean- or variance-related) m values produced by each g_k undergo a sparsemax (Martins & Astudillo, 2016) operation.

4. Path Following

As an application of the IK network, we present a method for recovering a sequence of joint locations $\mathbf{Y} = [\mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^n]$, given a desired path of end-effector locations $\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n]$. Ideally, given a many-to-one situation, one would like to obtain multiple different sequences, each of which should depict a smooth path in the joint location space, which matches (by applying forward kinematics) the desired sequence \mathbf{X} .

To achieve this, we employ a path following method. At each point along the desired path, we have the joint angles in the previous time-step $\bar{\mathbf{y}}^{t-1}$, and the desired end-effector position \mathbf{x}^t . By employing our network (f then g_1), we obtain a GMM distribution $p_{\mathbf{x}^t}^1$ for the first joint at time t .

To maintain a smooth transition, we sample $\bar{\mathbf{y}}_1^t$ (*i.e.*, the first joint of the current time step) from $p_{\mathbf{x}^t}^1$, subjected to the neighborhood Ω of $\bar{\mathbf{y}}_1^{t-1}$ of radius r , which we denote by $\Omega(\bar{\mathbf{y}}_1^{t-1}, r)$:

$$\bar{\mathbf{y}}_1^t \sim p_{\mathbf{x}^t}^1 |_{\Omega(\bar{\mathbf{y}}_1^{t-1}, r)} \quad (15)$$

This way, we maximize the smoothness of the path, while sampling from the learned distribution of the joint location that is conditioned on the end effector position in the current time-step. Thus emphasizing GMM components that are closer to the the joint angles in the previous time-step.

We now have the angle of the first joint at time t , denoted by $\bar{\mathbf{y}}_1^t$. We repeat the process using the same path following procedure, this time applied to $p_{\mathbf{x}^t}^2 = g(\theta_2, \bar{\mathbf{y}}_1^t)$. After this, the process is iterated for the rest of the joints. In our experiments we choose $r = 0.1$ radian.

Table 1. Results for the two 2D chains.

	Mean distance (cm)	Accuracy
2-joints	0.6 ± 0.6	$97.6\% \pm 4.7\%$
4-joints	0.8 ± 0.6	$95.6\% \pm 3.7\%$

5. Experiments

Our experiments check the performance of the IK method for isolated poses as well as for paths. We also present an experiment evaluating the representation learned by IKNet.

Metrics We compare the accuracy of the results using the following scores: (1) the mean Euclidean distance over the sample between $FK(\bar{\mathbf{y}})$ and \mathbf{x} , where $\bar{\mathbf{y}}$ is the obtained solution and FK stands for the forward kinematics, (2) accuracy, which is the ratio of solutions for which the Euclidean distance is less than a set threshold.

5.1. Evaluation on a 2D chain

In order to illustrate the ability of our method to capture the distribution of solutions, we train and evaluate with two settings of 2-dimensional chains – two joints (2J) and four joints (4J), where each joint is limited to 180 degrees for each direction. We train our model on a dataset of 20K random (reachable) points, and test on a different set of 1K (reachable) points. Tab. 1 show results for mean distance error and accuracy, where we measure accuracy as the percentage of points that are up to 2cm from the end-effector.

In the case of the 2J-chain, we notice, analytically, that the number of solutions for the first and second joints are 2 and 1, respectively. In the case of the 4J-chain, there is an infinite number of solutions for a given end-effector, since the kinematic chain has more degrees of freedom than the end-effector.

In Fig. 2 we illustrate 100-sampled solutions for each chain, and for a given end-effector position. Fig. 2(b,d) show the learned GMMs for each joint, and Fig. 2(a,c) show the sampled chain layouts, where opacity represents the probability of the layout, and starting / end-effector positions are illustrated with circle and X symbols, respectively.

As can be seen, the network learns to model the two layout solutions for a given point in the 2J-chain. The GMM of the first joint is collapsed into two main means, and the GMM of the second joint is collapsed into a single mean, both with low variance. In the 4J-chain, we can observe that the first joint’s angles are spread according to its GMM distribution, while the third and fourth joints collapsed to the same distribution of results as in the 2J-chain.

5.2. Comparison with IK methods

In the following section, we present the experiment setting used for evaluating our method, termed IKNet, against well-established IK methods, as well as machine-learning baselines. We use three different robotic arms, with different levels of complexity, as our benchmark kinematic chains.

5.2.1. BASELINES

Numerical methods We experiment with three types of optimization-based methods, (i) the Damped least squares Jacobian (Wampler, 1986) – which optimizes the end-effector position using the Jacobian of the model, (ii) the IK software package IKPy (Manceron, 2021) – which optimizes the position using ‘L-BFGS-B’ (Zhu et al., 1997) optimizer, and (iii) a differentiable model of the forward-kinematics package, DiffNEA, by Sutanto et al. (2020) – which optimizes the joint angle to minimize the L2 distance between the current and desired end-effector position.

Each method was initialized with multiple starting points in order to obtain multiple solutions, and ran for the same amount of time per sample during the optimization step for a fair comparison.

Learning methods We construct three different network models that capture different aspects of our method. First, we build an MLP network with depth 3 and width 1024, which takes as an input the desired end-effector, and outputs all joint angles at once. This baseline is incapable of generating multiple solutions. Second, instead of modeling conditional distributions sequentially, we use a network f with the same architecture as our hypernetwork to output mean vectors, covariance matrices (via Cholesky decomposition) and selection weights, to model the distribution as a mixture of multivariate Normal random variables, sampling all joints at once.

Last, we experiment with a recurrent neural network (RNN) architecture for modeling the sequential properties of the distribution of joint solutions. The RNN architecture is composed of a shared weights GRU (Cho et al., 2014) module for all joints, and an independent MLP part for each joint before and after the GRU module. This is done in order to model the unique part of each joint. To reflect the angles of the preceding joints, while reducing the repetition in the solution space, we project the end-effector position to the current coordinate system, reflecting the location of the joint after the preceding joints have determined its location.

5.2.2. KINEMATIC CHAINS (DATASETS)

We demonstrate our method on 3 different kinematic chains, which differ in their scale and degrees of freedom (DoFs) (Fig. 3). All demonstrated kinematic chains are redundant,

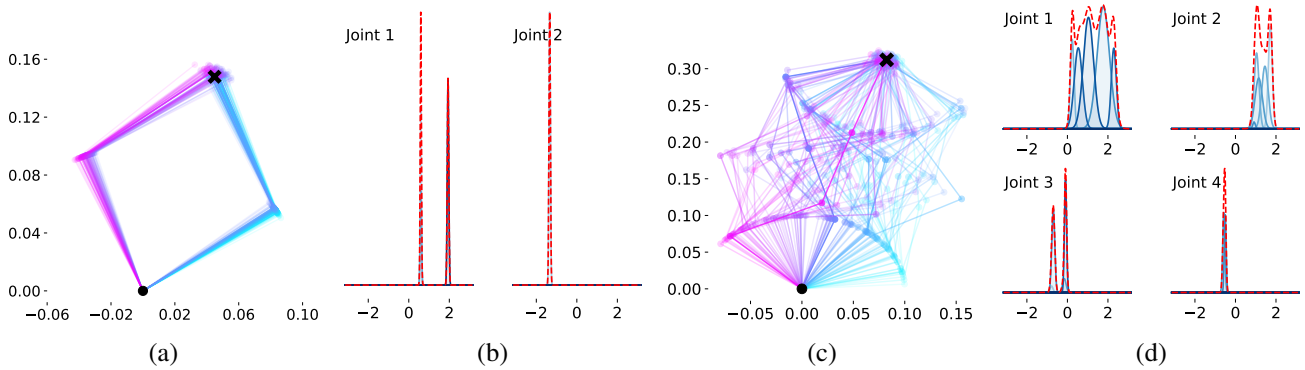


Figure 2. Illustration of the two 2D chains settings. We present 100 solutions for each chain and for a given end-effector position. (b,d) show the learned GMMs for each joint, and (a,c) show the sampled chain layouts, where opacity represents the probability of the layout, and starting / end-effector positions are illustrated with circle and X symbols, respectively.

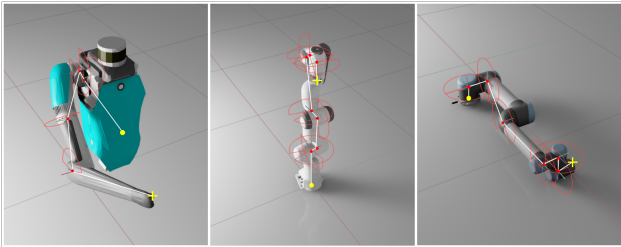


Figure 3. Illustration of the three kinematic chains used in our experiments. From left to right, Digit arm as obtained from the DigitRobot.jl repository <https://github.com/adubredu/DigitRobot.jl> – containing 4 DoFs, Franka panda – containing 7 DoFs and UR5 – containing 6 DoFs. The joints axis are illustrated in red.

i.e the number of DoFs is greater than the task space dimension (for the 3D inverse position kinematic task employed in the experiments, this dimension is 3).

Digit arm Digit by Agility robotics is a humanoid, bipedal robot, made for work in environments designed for humans. Digit’s upper torso is integrated with two 4-DoFs arms aimed for basic manipulation and object carrying tasks. As a source for this model, we employ the public repository by Alphonse Adu-Bredu that is available at <https://github.com/adubredu/DigitRobot.jl>. According to this resource, Digit’s arm DoF angular ranges are $[(-1.3, 1.3), (-2.5, 2.5), (-1.75, 1.75), (-1.35, 1.35)]$ radians.

UR5 UR5 by Universal Robots is an industrial, flexible, lightweight, 6-DoFs robotic arm with working radius of up to 85.0cm. The DoFs angular range are $[(-3.14, 3.14), (-3.14, 3.14), (-2.5, 2.5), (-3.14, 3.14), (-3.14, 3.14), (-3.14, 3.14)]$ radians.

Franka Franka Emika Panda is a 7-DoFs programmable robotic arm with working radius of up to 85.5cm. The DoF angular ranges are $[(-2.9, 2.9), (-1.76, 1.76), (-2.9, 2.9), (-3.07, -0.07), (-2.9, 2.9), (-0.02, 3.75), (-2.9, 2.9)]$ radians.

5.2.3. RESULTS

The results are presented in Tab. 2 and depict mean and Standard Deviation on a test set of reachable arm locations. As can be seen, IKNet outperforms both the optimization based baselines and the learning based baselines in terms of mean distance, with the exception of IKPy outperforming somewhat better than our method on average, but with much higher variance.

Accuracy is measured at a 10cm threshold. Our method outperforms all baselines in this metric. Lastly, we measure the actual runtime per method and report results for 100 executions. For this purpose, the learning based methods were run on a CPU. As can be seen, our method is considerably more efficient than the optimization based methods. While the runtime cannot be taken at face value, since it is implementation-dependent, our method has an inherent advantage since it does not iterate.

5.3. Path following

We next evaluate the path following methods presented in Sec. 4. This experiment does not involve any training and evaluation sequences \mathbf{X} of end effector positions were generated by moving the Digit arm smoothly. Each sequence is of length 50.

As a baseline to the path following method we propose, we employ IKPy, which is the best baseline method found in our single-position experiments. As an iterative method, IKPy finds a solution that is close to the starting point and

Table 2. Inverse kinematic results on three kinematic chains. Our results are provided as the average of 100 samples and not based on the most likely solution, which would improve IKNet results. For each method we present the mean across the test set as well as the Standard Deviation. As can be seen, our method achieves best accuracy and standard deviation across all robots, and best accuracy for Digit arm and UR5. For Franka panda, our method is compatible with the results of IKPy, but with better standard deviation.

	Distance (cm)	Accuracy	Runtime (s)
DIGIT - Arm (4 DoFs)			
Damped least squares Jacobian	8.3 ± 19.6	80.4% ± 15.9%	0.0400
IKPy	7.5 ± 16.0	79.0% ± 13.6%	0.0850
DiffNEA	28.7 ± 20.9	21.8% ± 6.5%	0.1100
MLP	12.3 ± 1.2	56.2%	0.0002
RNN	12.7 ± 11.0	53.5% ± 24.2%	0.0090
Multivariate GMMs	4.4 ± 3.7	92.3% ± 12.7%	0.0020
IKNet	2.3 ± 1.7	99.5% ± 1.3%	0.0070
UR5 (6 DoFs)			
Damped least squares Jacobian	10.9 ± 22.3	76.2% ± 40.9%	0.2000
IKPy	6.5 ± 15.4	81.7% ± 9.9%	0.1200
DiffNEA	36.4 ± 24.5	15.3% ± 4.8%	0.1100
MLP	59.1 ± 0.9	2.6%	0.0002
RNN	5.7 ± 8.9	84.7% ± 7.4%	0.0090
Multivariate GMMs	5.0 ± 3.8	89.8% ± 15.5%	0.0050
IKNet	2.8 ± 2.1	98.8% ± 1.9%	0.0120
Franka (7 DoFs)			
Damped least squares Jacobian	4.4 ± 13.3	88.7% ± 6.8%	0.1100
IKPy	2.1 ± 7.6	91.2% ± 6.3%	0.1650
DiffNEA	22.8 ± 18.4	29.4% ± 8.9%	0.1100
MLP	57.7 ± 1.12	5.4%	0.0003
RNN	4.9 ± 6.1	84.7% ± 7.4%	0.0200
Multivariate GMMs	6.5 ± 4.6	82.1% ± 16.8%	0.0100
IKNet	3.1 ± 2.6	98.0% ± 2.0%	0.0300

can, therefore, be applied sequentially to the positions along the end effector’s path to obtain a smooth trajectory in the joint space.

Comparing the Euclidean error in the end effector positions, the paths we generate obtain a similar average error to that of IKPy (3.0cm). However, for our method this is an average over 100 different paths and not the results for the best path, while for IKPy it is the result for the single path that starts at the ground truth joint position and varies smoothly (an ideal setting for IKPy). When selecting the generated path with the highest fidelity among all 100 generated paths (still considerably faster than running IKPy), we obtain an average error of 1.8 cm.

Fig. 4 presents the multiple trajectories obtained per a single target sequence \mathbf{X} . As can be seen, the generated paths

present a high degree of variability.

5.4. Robustness and few-shot learning

Since our model is learning-based and since it employs a sampling-based approach, it can naturally model noisy robot dimensions. To demonstrate this we created a set of 55 4-jointed 3D arms that differ by at most $\pm 20\%$ in the length of each segment of the chain. We then learn one IKNet per arm and one based on the entire data. At test time, we sample 10 new random arms and evaluate the 55 single arm models and the one trained on the entire dataset.

Out of the 55 random models, some are more similar than others and are ranked by the accuracy (the error ranges between 2.1-3.6cm, on average). The model that we train on the entire data is, on average, at the 75th percentile of the

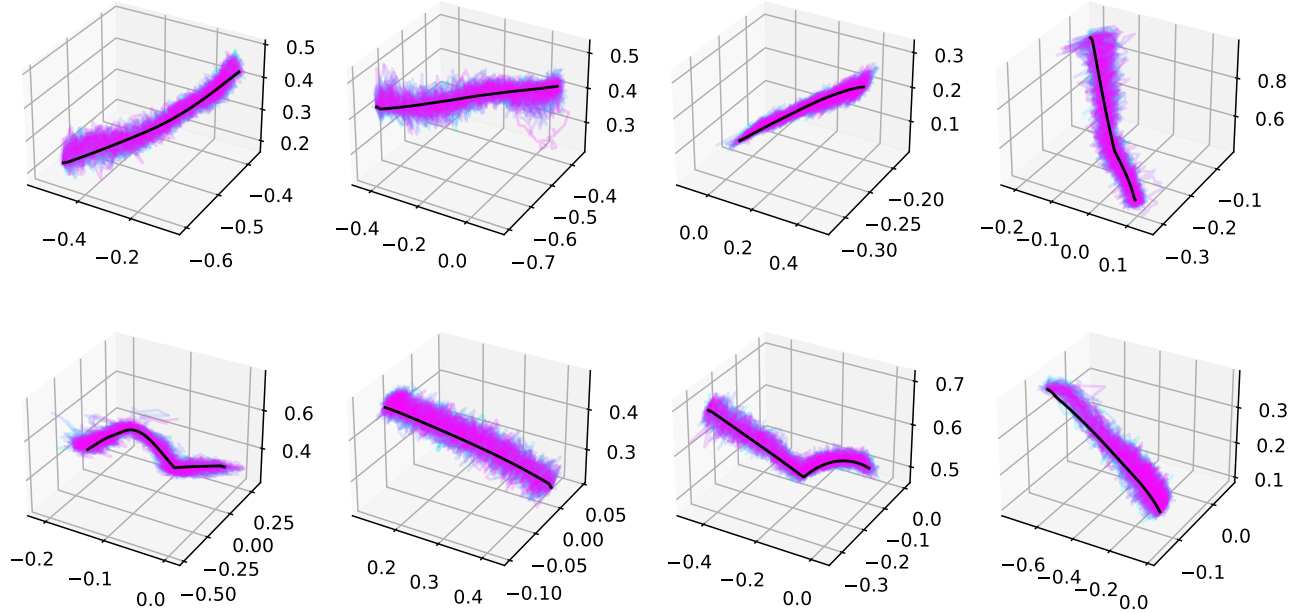


Figure 4. Path following. Given a sequence of end effector locations, we use a path following method in order to recover multiple possible trajectories. The axis of the plots represent the angles along the first three joints of the Digit arm.

results (standard deviation 6%), which is a clear indication that the unified model learns a robust solution that matches many random arms.

Additionally, we expect our method to have an advantage in scenarios in which one needs to learn to perform IK from a few real measurements. To test this, we propose to leverage transfer learning in order to tackle a very common real-world scenario, in which a given robot deviates from the specifications.

We learn an IK model \mathcal{M}_A for a 3D kinematic chain with 4 joints (“arm A”) based on 100K samples. Then, we sample a new arm (“arm B”) in which the segment lengths vary randomly by up to $\pm 20\%$ from the original design. We then fine-tune \mathcal{M}_A based on samples from arm B, obtaining model \mathcal{M}_B .

The results are presented in Fig. 5. Applying \mathcal{M}_A on the test set of arm B yields an error of 3.62cm (dashed blue line). Training on 100K samples for arm B yields 1.84cm (dashed blackline). With 1K samples, the finetuned model \mathcal{M}_B obtains a similar error of 1.92cm, while a model trained from scratch with the same number of samples has a much higher error of 9.26cm. Evidently, with a relatively small number of training points, one reaches the same level of accuracy that can be obtained from 100K samples of arm B.

5.5. Learning Seq2Seq mapping

Lastly, in order to demonstrate that the representation learned by IKNet is powerful, we consider the task of learning to map a single sequence of effector poses $\mathbf{X} = [x^1, x^2, \dots, x^{100}]$ to a sequence of joint angles $\mathbf{Y} = [y^1, y^2, \dots, y^{100}]$. For this seq2seq task, we employ a Transformer (Vaswani et al., 2017).

The transformer is trained on 10k pairs of matching sequences, each of length 100, from the UR5 chain. We compare two variants. In the first, the transformer is trained

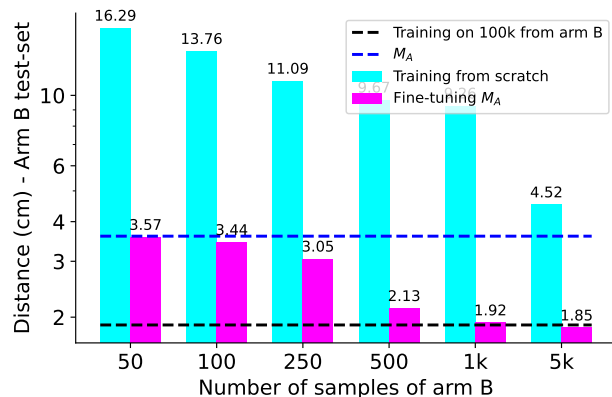


Figure 5. The number of training points during the training (magenta) or fine-tuning (cyan) phase of arm B vs. the forward kinematics error (cm) on the test-set of arm B.

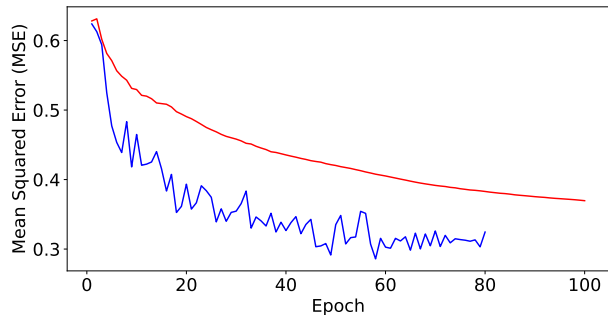


Figure 6. Convergence plots when training a sequence to sequence transformer model based either on the sequence of the end effector positions (red) or its embedding by network f (blue).

to map \mathbf{X} to \mathbf{Y} directly. In the second, we employ representation obtained by f' , which is the mapping between an input of network x and the activations of the last layer of f , before the N projections to the weights of the networks g_k . In this case, the sequence to sequence problem learned is between $f'(\mathbf{X}) = [f(x^1), f(x^2), \dots, f(x^{100})]$ and \mathbf{X} .

We employ a 6-layers deep Transformer encoder, with 8 attention heads. We also use an embedding of size 128, as our hypernetwork embedding size. When training with raw inputs, a trainable linear layer projecting the 3-dimensional input to 128 dimensions is applied, in order to allow for fair comparison.

In order to evaluate the results, we employ a test set of 100 trajectories of (\mathbf{X}, \mathbf{Y}) , each of length 100. The convergence graphs for both methods are presented in Fig. 6. The plots depict the mean squared Euclidean error on the test set per epoch. We stopped training the transformer model with the IKNet representation after 80 epochs, since it converged faster. Therefore, the two plots are of different lengths. Evidently, employing the IKNet representation leads to faster convergence and to an overall lower error.

6. Discussion

The most successful baseline methods are designed to provide one IK solution, given an initial position from which optimization starts. In contrast, variability is natural in our model, which is a feed-forward model and not an iterative optimization model.

The runtime of the method is a good indicator of its applicability. Reductio ad absurdum, without any limit on the runtime or the number of restarts, almost all optimization methods would reach optimal results and fully characterize the solution space. In the experiments, we relied on the default parameters provided in the implementation of each method. In the same vain, it is possible to use an optimization method on top the solutions provided by IKnet and

obtain negligible MSEs. We avoid this in order to provide the raw results as returned by this neural model.

The more degrees of freedom the kinematic chain has, the easier it is for an optimization method to find a single valid vector of joint angles, given the end effector’s location. However, modeling the ambiguity, which is the task solved by our network, becomes more difficult. In the Digit arm and UR5 experiment, where there are 4 and 6 DOF, respectively, our advantage over the baselines is more pronounced than it is for Franka with 7 DOF.

Characterization of the entire probability distribution can also help achieve what is called robust inverse kinematics (Sinha & Chakraborty, 2019). In this setting, one would like to select the IK solution that is the most stable with respect to errors in the joint angles. A robust solution is one such that the entire ball in joint angle space (where y resides) leads to end effector positions within a certain tolerance of the desired position and angle (x).

7. Conclusions

We present a neural IK model that can capture the inherent one-to-many ambiguity of the problem, while training on a dataset with one-to-one samples. The architecture consists of a single hypernetwork and a sequence of primary matrices. Making use of the hierarchical nature of the problem, each joint is sampled from a GMM that is conditioned on the samples performed for the previous joints in the kinematic chain.

Having an accurate feed forward model that supports multiple outputs has a few advantages, which we demonstrate. First, the solution is extremely efficient at run-time. Second, it entails an effective representation of the Cartesian positions. Third, one can use such a model to obtain multiple smooth paths. In addition, while not demonstrated here, having a differentiable model allows it to be incorporated as a module in a complex network during inference or as part of a loss at training time.

References

- Baranes, A. and Oudeyer, P.-Y. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1):49–73, 2013.
- Bensadoun, R., Gur, S., Galanti, T., and Wolf, L. Meta internal learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- Buss, S. R. Introduction to inverse kinematics with jacobian transpose, pseudoinverse and damped least squares

- methods. *IEEE Journal of Robotics and Automation*, 17 (1-19):16, 2004.
- Bócsi, B., Nguyen-Tuong, D., Csató, L., Schölkopf, B., and Peters, J. Learning inverse kinematics with structured prediction. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 698–703, 2011. doi: 10.1109/IROS.2011.6094666.
- Cho, K., Van Merriënboer, B., Bahdanau, D., and Bengio, Y. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- Craig, J. J. *Introduction to robotics: mechanics and control, 3/E*. Pearson Education India, 2009.
- Csiszar, A., Eilers, J., and Verl, A. On solving the inverse kinematics problem using neural networks. In *2017 24th International Conference on Mechatronics and Machine Vision in Practice (M2VIP)*, pp. 1–6, 2017. doi: 10.1109/M2VIP.2017.8211457.
- De Angulo, V. R. and Torras, C. Learning inverse kinematics: Reduced sampling through decomposition into virtual robots. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 38(6):1571–1577, 2008.
- Diankov, R. Automated construction of robotic manipulation programs. 2010.
- D’Souza, A., Vijayakumar, S., and Schaal, S. Learning inverse kinematics. In *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180)*, volume 1, pp. 298–303 vol.1, 2001. doi: 10.1109/IROS.2001.973374.
- Duka, A.-V. Neural network based inverse kinematics solution for trajectory tracking of a robotic arm. *Procedia Technology*, 12:20–27, 2014.
- El-Sherbiny, A., Elhosseini, M. A., and Haikal, A. Y. A comparative study of soft computing methods to solve inverse kinematics problem. *Ain Shams Engineering Journal*, 9(4):2535–2548, 2018.
- Galanti, T. and Wolf, L. On the modularity of hypernetworks. In *Advances in Neural Information Processing Systems 33*. Curran Associates, Inc., 2020.
- Grochow, K., Martin, S. L., Hertzmann, A., and Popović, Z. Style-based inverse kinematics. In *ACM SIGGRAPH 2004 Papers*, pp. 522–531. 2004.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Huang, J., Wang, Q., Fratarcangeli, M., Yan, K., and Pelachaud, C. Multi-Variate Gaussian-Based Inverse Kinematics. *Computer Graphics Forum*, 2017. ISSN 1467-8659. doi: 10.1111/cgf.13089.
- Klanke, S., Vijayakumar, S., and Schaal, S. A library for locally weighted projection regression. *Journal of Machine Learning Research*, 9:623–626, 2008.
- Manceron, P. ikpy: An inverse kinematics library aiming performance and modularity (v3.2). <https://github.com/Phylliade/ikpy>, 2021.
- Martins, A. and Astudillo, R. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International conference on machine learning*, pp. 1614–1623. PMLR, 2016.
- Moulin-Frier, C., Rouanet, P., and Oudeyer, P.-Y. Explauto: An open-source python library to study autonomous exploration in developmental robotics. In *4th International Conference on Development and Learning and on Epigenetic Robotics*, pp. 171–172. IEEE, 2014.
- Raghavan, M. and Roth, B. Inverse Kinematics of the General 6R Manipulator and Related Linkages. *Journal of Mechanical Design*, 115(3):502–508, 09 1993. ISSN 1050-0472. doi: 10.1115/1.2919218. URL <https://doi.org/10.1115/1.2919218>.
- Rolf, M., Steil, J. J., and Gienger, M. Bootstrapping inverse kinematics with goal babbling. In *2010 IEEE 9th International Conference on Development and Learning*, pp. 147–154. IEEE, 2010a.
- Rolf, M., Steil, J. J., and Gienger, M. Goal babbling permits direct learning of inverse kinematics. *IEEE Transactions on Autonomous Mental Development*, 2(3):216–229, 2010b.
- Schölkopf, B., Smola, A. J., Bach, F., et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Sears, P. and Dupont, P. E. Inverse kinematics of concentric tube steerable needles. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pp. 1887–1892, 2007. doi: 10.1109/ROBOT.2007.363597.
- Sinha, A. and Chakraborty, N. Computing robust inverse kinematics under uncertainty. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 59230, pp. V05AT07A065. American Society of Mechanical Engineers, 2019.
- Sutanto, G., Wang, A., Lin, Y., Mukadam, M., Sukhatme, G., Rai, A., and Meier, F. Encoding physical constraints

in differentiable newton-euler algorithm. In *Learning for Dynamics and Control*, pp. 804–813. PMLR, 2020.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Vijayakumar, S. and Schaal, S. Locally weighted projection regression: An $o(n)$ algorithm for incremental real time learning in high dimensional space. In *Proceedings of the seventeenth international conference on machine learning (ICML 2000)*, volume 1, pp. 288–293. Morgan Kaufmann, 2000.

Walter, J. and Ritter, H. Rapid learning with parametrized self-organizing maps. *Neurocomputing*, 12(2-3):131–153, 1996.

Wampler, C. W. Manipulator inverse kinematic solutions based on vector formulations and damped least-squares methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 16(1):93–101, 1986.

Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on mathematical software (TOMS)*, 23(4):550–560, 1997.