SCIENCE ROBOTICS | RESEARCH ARTICLE

ARTIFICIAL INTELLIGENCE

Machine learning–driven self-discovery of the robot body morphology

Fernando Díaz Ledezma* and Sami Haddadin

The morphology of a robot is typically assumed to be known, and data from external measuring devices are used mainly for its kinematic calibration. In contrast, we take an agent-centric perspective and ponder the vaguely explored question of whether a robot could learn elements of its morphology by itself, relying on minimal prior knowledge and depending only on unorganized proprioceptive signals. To answer this question, we propose a mutual information–based representation of the relationships between the proprioceptive signals of a robot, which we call proprioceptive information graphs (π-graphs). Leveraging the fact that the information structure of the sensorimotor apparatus is dependent on the embodiment of the robot, we use the π-graph to look for pairwise signal relationships that reflect the underlying kinematic first-order principles applicable to the robot's structure. In our discussion, we show that analysis of the π-graph leads to the inference of two fundamental elements of the robot morphology: its mechanical topology and corresponding kinematic description, that is, the location and orientation of the robot's joints. Results from a robot manipulator, a hexapod, and a humanoid robot show that the correct topology and kinematic description can be effectively inferred from their π-graph either offline or online, regardless of the number of links and body configuration.

INTRODUCTION

Understanding how an embodied agent makes sense of its sensorimotor system and learns about its morphology is a question in psychology, neuroscience, and, recently, robotics. Nevertheless, giving robots the capability to learn the properties of their bodies has been a topic addressed in a relatively limited number of works. In general, this connects to the commonly made assumption that the actual morphology of the robot, or at the very least, its kinematics, is always known and requires only calibration routines. However, such routines require external metrology systems. In contrast, similar to an approach shown previously (*1*), we recognize the importance of a robot's capability to assess and continuously update the knowledge about its morphology autonomously. This capability implies that future embodied robotic agents will have to leverage their sensorimotor system's inherent structure to gradually develop an understanding of their body despite being initially oblivious to its physical characteristics. Incremental learning of the morphology would allow robots to adapt their parameters to reflect the changes in the body structure that could result from self-inflicted or externally inflicted actions—see, for example, Bongard *et al.* (*2*). In addition, incorporating complementary sensor modalities akin to touch would enable self-calibration approaches independent of vision systems (*3*). Immediate extensions could apply to robots with modular topologies that benefit from self-monitoring and calibration capabilities. Likewise, body morphology learning could be combined with knowledge-transfer methods to allow robots with similar topologies to boost the learning of their bodily structure. In summary, given the importance that learning the body schema has for robots (*4*, *5*), this work provides a framework for the incremental learning of two fundamental aspects of the body schema, the robot body topology and its kinematic description.

An embodied robotic agent must autonomously learn and refine its body schema as a model representation of the robot's morphology. The body schema is mainly built from proprioceptive information (*6*, *7*) and generally refers to an internal representation of the body, including the arrangement and geometry of its parts. Adaptive and self-acquired, the body schema is part of an agent's internal forward and inverse models and is used to plan and predict sensorimotor interactions. The robot morphology refers to the mechanical arrangement and spatial description of the body parts and the type and placement of sensors and effectors (*8*). We envision learning the robotic body schema as a series of learning stages that the robot must go through to discover core features of its body morphology. These features are part of the set

$$\mathcal{S} = \{N, \mathcal{G}, \mathbf{A}, \boldsymbol{\lambda}, \boldsymbol{\theta}\} \tag{1}$$

whose elements are the number of bodies N in the kinematic chain, the graph $\mathcal{G}$ that describes the sensorimotor interactions that result from the robot's embodiment, the adjacency matrix $\boldsymbol{A}$ that captures the topology—mechanical arrangement of the bodies and joints, its kinematic description $\boldsymbol{\lambda}$, and the inertial parameters of the links $\boldsymbol{\theta}$. Thus, a robot capable of acquiring knowledge of the elements of $\mathcal{S}$ is an agent that builds an understanding of its physical self. An overview of this process is depicted in fig. S1.

Most of the methods in the research literature to learn and adapt the body schema of a robot rely on a visual input source from which the robot can build an understanding of its geometry (*9*–*14*). In contrast, complementary works have explored alternatives to extract body knowledge using the robot's perceptual apparatus to find internal representations (*15*) and to find and encode sensorimotor maps (*16*). Information-theoretic approaches have also been applied to learn sensorimotor networks based on information flow (*17*), develop sensor sensoritopic maps according to Crutchfield's information metric (*18*), and extract knowledge about robot morphology from sensorimotor regularities (*19*). Other explored body models include Bayesian networks (*13*) and classical kinematic representations (*20*). In addition, there are methods closely connected to our present work

---

Chair of Robotics and Systems Intelligence, MIRMI—Munich Institute of Robotics and Machine Intelligence, Technical University of Munich, Georg-Brauchle-Ring 60-62, München 80992, Germany.
*Corresponding author. Email: fernando.diaz@tum.de

that discuss the advantages of continuous self-modeling from proprioception (*2*, *21*), use graph theory as a formalism to obtain the dynamic equations of tree-structure robots (*22*), apply screw formulations in a kinematic identification scheme (*23*), and search for the kinematic structure of a serial robot using inertial measurement units (IMUs) (*24*).

Although many works address one or more of the learning stages in fig. S1, to the best of our knowledge, there is a lack of a unifying scheme that breaches the gaps to define a synergistic integration of all the learning stages and to produce a fully characterized body schema from only knowledge about the sensorimotor signals. We focus on the first three stages of fig. S1 (highlighted in red) and relinquish the need for any exteroceptive information (for example, any form of visual or positional input), depending solely on proprioception to determine essential aspects of the robot body morphology. Subscribing to the belief that body structure lies concealed in the relationships among the proprioceptive signals, we leverage concepts from graph and information theory to seek a representation of said relationships that facilitates their study in connection to the physical laws that govern the body morphology to infer the robot's topology and characterize its kinematic description.

Embodiment and information structure

The embodiment of a robot determines its perceptual and behavioral capabilities and influences the emergence of sensorimotor regularities arising from body-environment interaction (*8*). Information-theoretic quantities have been used to capture these regularities as information flows between the various modalities in the sensorimotor system (*17*, *18*, *25*–*27*). To illustrate how embodiment influences information structure, consider a hypothetical robot equipped with a set of sensors, depicted in Fig. 1A. From an information-theoretic perspective, these sensors are information sources that share information with each other. The strength of these relations is represented graphically as edges connecting the sensors (light blue edges). As a result of the underlying embodiment of the robot, stronger connections between sensors exist (dark blue edges). These strong connections hint at the robot's body topology (Fig. 1B), a tree-like serial kinematic structure where the stronger exchange of information occurs between sensors of adjacent bodies. Succinctly, body structure becomes the hypothesis behind the strength of information sharing.

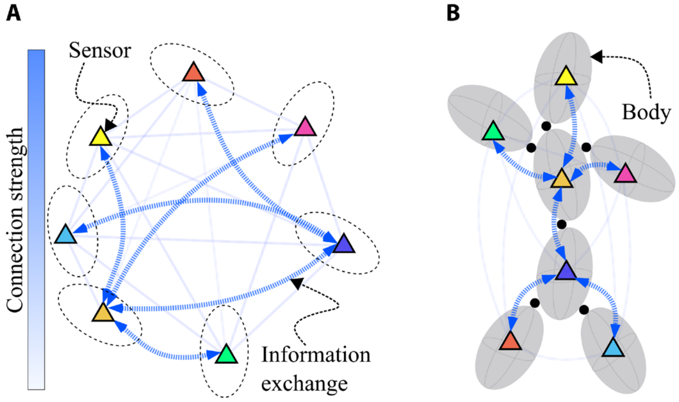

Fig. 1. Embodiment and information structure. (**A**) Sensors on a robot exchange information at different strength levels. (**B**) This strength is associated with the embodiment of the robot, a kinematic tree in this case.

Inferring connectivity

Network topology inference (NTI) involves discovering and visualizing relationships among system components, often represented as a network or graph (*28*); further details about graph-related terminology are provided in Supplementary Methods. Measurement-based approaches for estimating network topology rely on physically motivated and statistical models, including correlation, probabilistic graphical models, entropy, mutual information (MI), and transfer entropy (*28*). Some NTI methods fall under the framework of graph signal processing (GSP) (*29*), offering new analysis possibilities (*30*–*32*). Functional connectivity (FC) is an NTI method that uses information-theoretic metrics to assess dependencies based on probability distributions of observed signals (*33*). It can be categorized as undirected or directed, with the latter focused on exploring statistical causation from the data (*34*). In contrast, effective connectivity deals with dynamic influence among network elements under specific causal models (*35*).

Robot proprioception

The proprioception of an agent provides the ability to sense the position, orientation, movement, and force of its body parts (*36*), which is fundamental to learning about its body structure. In robotics, proprioception results from integrating various complementary sensors, for example, encoders, gyroscopes, accelerometers, and force/torque sensors (*37*). Apart from the typical measurements of the robot's joint position q, velocity $\dot{q}$, and joint torque τ, the use of IMUs to measure the acceleration $\dot{\boldsymbol{v}}$ and angular velocity $\boldsymbol{\omega}$ of the robot's bodies has had applications in state (*38*) and joint velocity and acceleration estimation (*39*, *40*), the control of mobile robots (*41*), and safe human-robot interaction (*42*).

The body morphology learning problem

Acknowledging that the negligible knowledge about the robot structure impedes the determination of the sensorimotor analytic relationships, we used NTI to study the FC among the robot's proprioceptive signals by defining the robot's π-graph, that is, a graph whose connectivity is associated with the amount of pairwise information sharing. Under the premise that embodiment and information structure are connected, we analyzed the π-graph and demonstrate how the robot's body topology lies encoded in its connectivity. We used the inferred topology and fundamental first-order principles from differential kinematics to determine the orientation and location of the robot joint axes (namely, its kinematic description). Our proposed body morphology inference method relies on the assumption of the type of structure and the awareness of the number and modality of the available proprioceptive signals but is bound to neither the knowledge of their hierarchical arrangement nor the location and orientation of their corresponding sensors. Our work differs from classical kinematic calibration, in that no previous kinematic model is considered, and from vision-centered kinematic estimation, because no motion-tracking systems are required. We applied our method to a manipulator with 7 degrees of freedom (DoF), an 18-DoF hexapod, and a 25-DoF humanoid, demonstrating that, if sufficient kinematic excitation is provided to all bodies, the robot morphology can successfully be recovered from the π-graph.

RESULTS

A general overview of our proposed framework to learn the robot morphology is shown in Fig. 2. In the context of our framework, we argue that the information provided by the tuple ($\dot{\boldsymbol{v}}, \boldsymbol{\omega}$) is essential to complement joint proprioception and reveal spatial properties of the robot's morphology. In particular, having a minimum of one IMU per body in the kinematic chain enables the identification of mechanical connections between bodies (*23*, *24*). Yet, precise knowledge of the location and orientation of the sensor coordinate frames is unnecessary. Therefore, our focus lies on robots with N joints equipped with sensors that record a time series

$$\boldsymbol{x}(t) = \left[\boldsymbol{q}^T \dot{\boldsymbol{q}}^T \boldsymbol{\tau}^T \boldsymbol{\omega}_1^T \dots \boldsymbol{\omega}_{N+1}^T \dot{\boldsymbol{v}}_N^T \dots \dot{\boldsymbol{v}}_{N+1}^T\right]^T \in \mathbb{R}^m \tag{2}$$

of proprioceptive signals, where $\boldsymbol{q}$, $\dot{\boldsymbol{q}}$, and $\boldsymbol{\tau} \in \mathbb{R}^N$ and $\dot{\boldsymbol{v}}$ and $\boldsymbol{\omega} \in \mathbb{R}^{3(N+1)}$, from which n samples are taken to define the data matrix

$$\boldsymbol{X} = [\boldsymbol{x}(t)\ \boldsymbol{x}(t+1)\ \dots\ \boldsymbol{x}(t+n)] \in \mathbb{R}^{m \times n} \tag{3}$$

With $m = 3\,N + 2[3(N + 1)]$. Under the fundamental premise that (assumption 1) the robot is a serial kinematic tree with revolute joints and that knowledge about the number and modalities of the signals $\boldsymbol{x}(t)$ is available, yet agnostic to their hierarchical arrangement, the body morphology learning problem consists of using $\boldsymbol{X}$ to infer the robot's body topology and associated kinematic description (location and orientation of its joints axes). To solve this problem, we present a body morphology inference method composed of a sequence of five subprocesses, depicted in Fig. 2, where (subprocess 1) proprioceptive signals $\boldsymbol{x}(t)$ are generated from motor babbling and (subprocess 2) collected in a replay

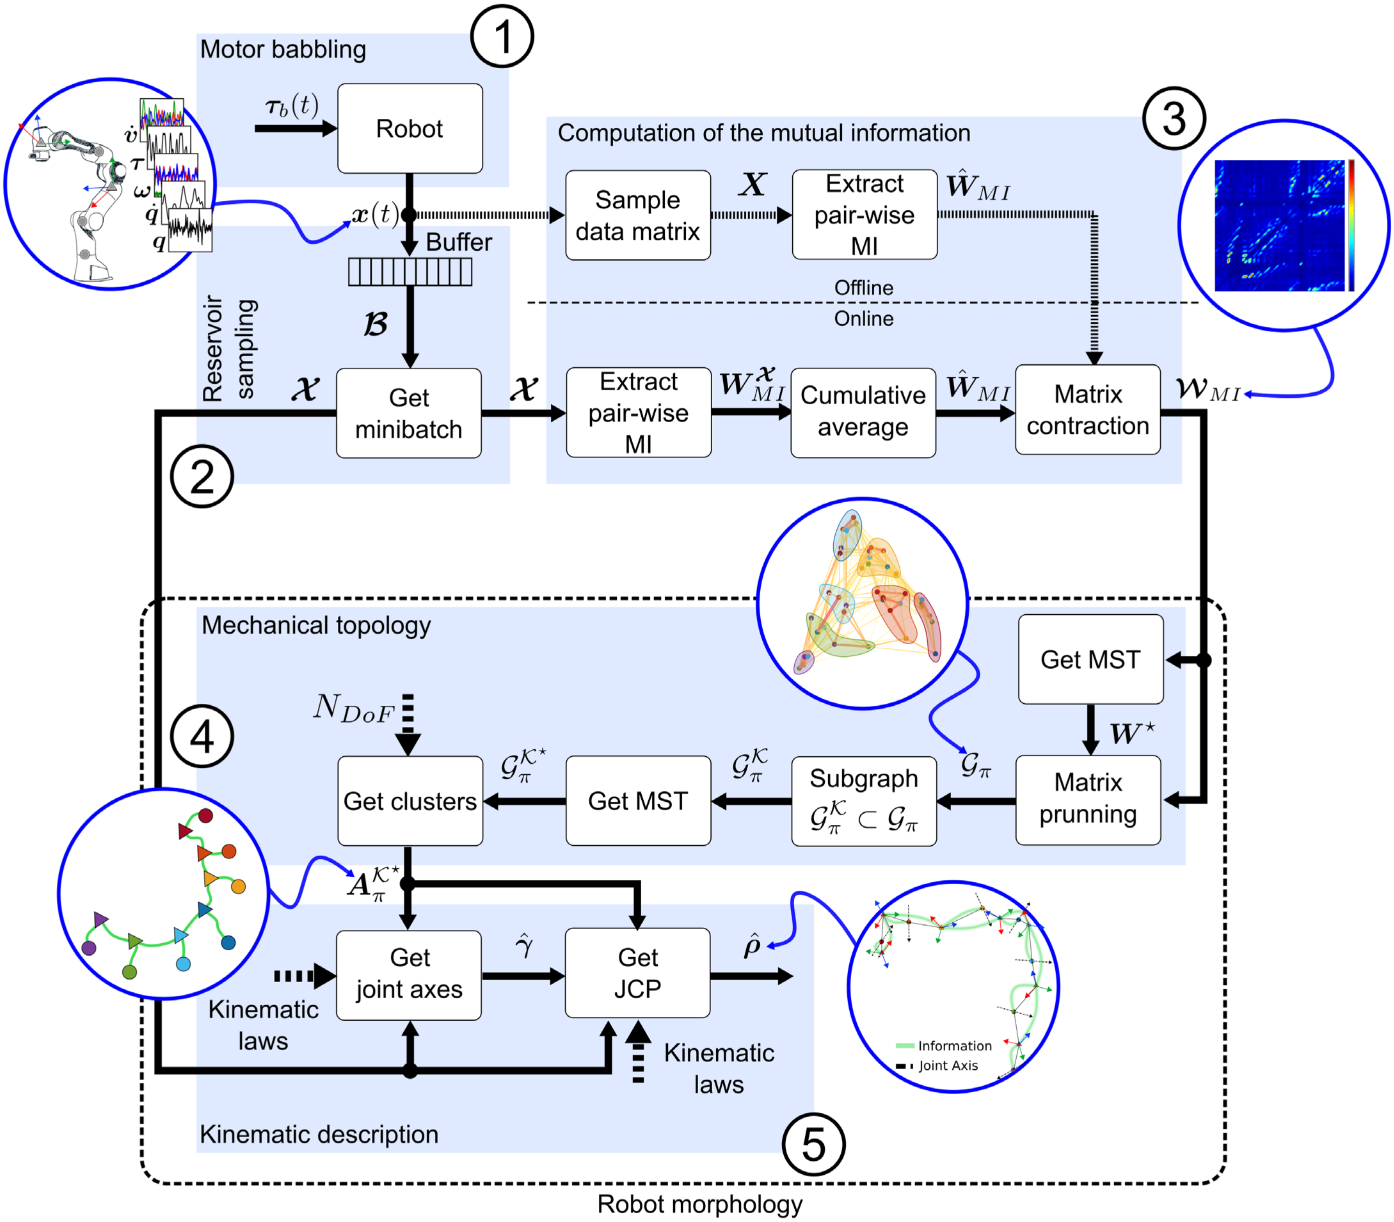

Fig. 2. Inferring robot body morphology from proprioception. Via motor babbling, proprioceptive signals with high variability are generated. The pairwise MI from the signals is used to infer a graph whose adjacency matrix leads to the robot's body topology and associated kinematic description.

buffer from which to extract the data matrix $\boldsymbol{X}$. We use these data to (subprocess 3) construct a relational matrix $\widehat{\boldsymbol{W}}_{\mathrm{MI}}$ that summarizes the proprioceptive pairwise information sharing. Consequently, (subprocess 4) $\widehat{\boldsymbol{W}}_{\mathrm{MI}}$ is used to define a graph $\mathcal{G}_\pi$ that reflects the relevant relationships among the proprioceptive signals and helps reveal subsets of them that relate to bodies and joints in the kinematic chain. Then, we extract the connectivity between these subsets and represent it as a binary matrix $\boldsymbol{A}_\pi^{\mathcal{K}*}$ that summarizes the body topology of the robot. Last, (subprocess 5) based on the inferred topology, the kinematic description of the robot is determined by learning a vector $\widehat{\boldsymbol{\lambda}}$ whose parameters express the location and orientation of the joint axes.

In the following, we review the results obtained from the application of our framework to robots of different morphology and with movable and static bases (see movie S1). Details about the experiments are provided in Supplementary Methods.

The morphology of a manipulator

We tested our method using a simulated seven-DoF Franka Emika research robot manipulator (*43*). To allow the learning of the kinematic parameters $\widehat{\boldsymbol{\lambda}}$ of links proximal to the base (as will be discussed later), motion to the manipulator's base was given by adding a virtual three-DoF joint connecting the base body frame to the world frame. The joint angle, velocity, torque, link acceleration, and angular velocity for each of the $N + 1$ bodies in the kinematic chain were measured, resulting in $m = 69$ signals. From these measurements, the data matrix $\boldsymbol{X}$ was constructed and used to compute the estimated MI adjacency matrix $\widehat{\boldsymbol{W}}_{\mathrm{MI}}$ (Fig. 3A). This led to the extraction of the π-graph $\mathcal{G}_\pi$, from which two subgraphs were identified, the kinematics subgraph $\mathcal{G}_\pi^{\mathcal{K}}$ (consisting only of kinematics-related signals) and its associated maximum spanning tree (MST) $\mathcal{G}_\pi^{\mathcal{K}*}$, shown together with its adjacency matrix in Fig. 3 (D and E), with the MST highlighted in green. The latter is associated with the mechanical topology of the robot and was used to infer it.

To assess whether other state-of-the-art NTI methods are suitable for robot topology inference, we used $\boldsymbol{X}$ to estimate weighted adjacency matrices using correlation- and GSP-based methods (see Materials and Methods). Figure 3B shows the comparison between the local structure of the MST subgraph $\mathcal{G}^{\mathcal{K}*}$ for the correlation-, GSP-, and MI-based

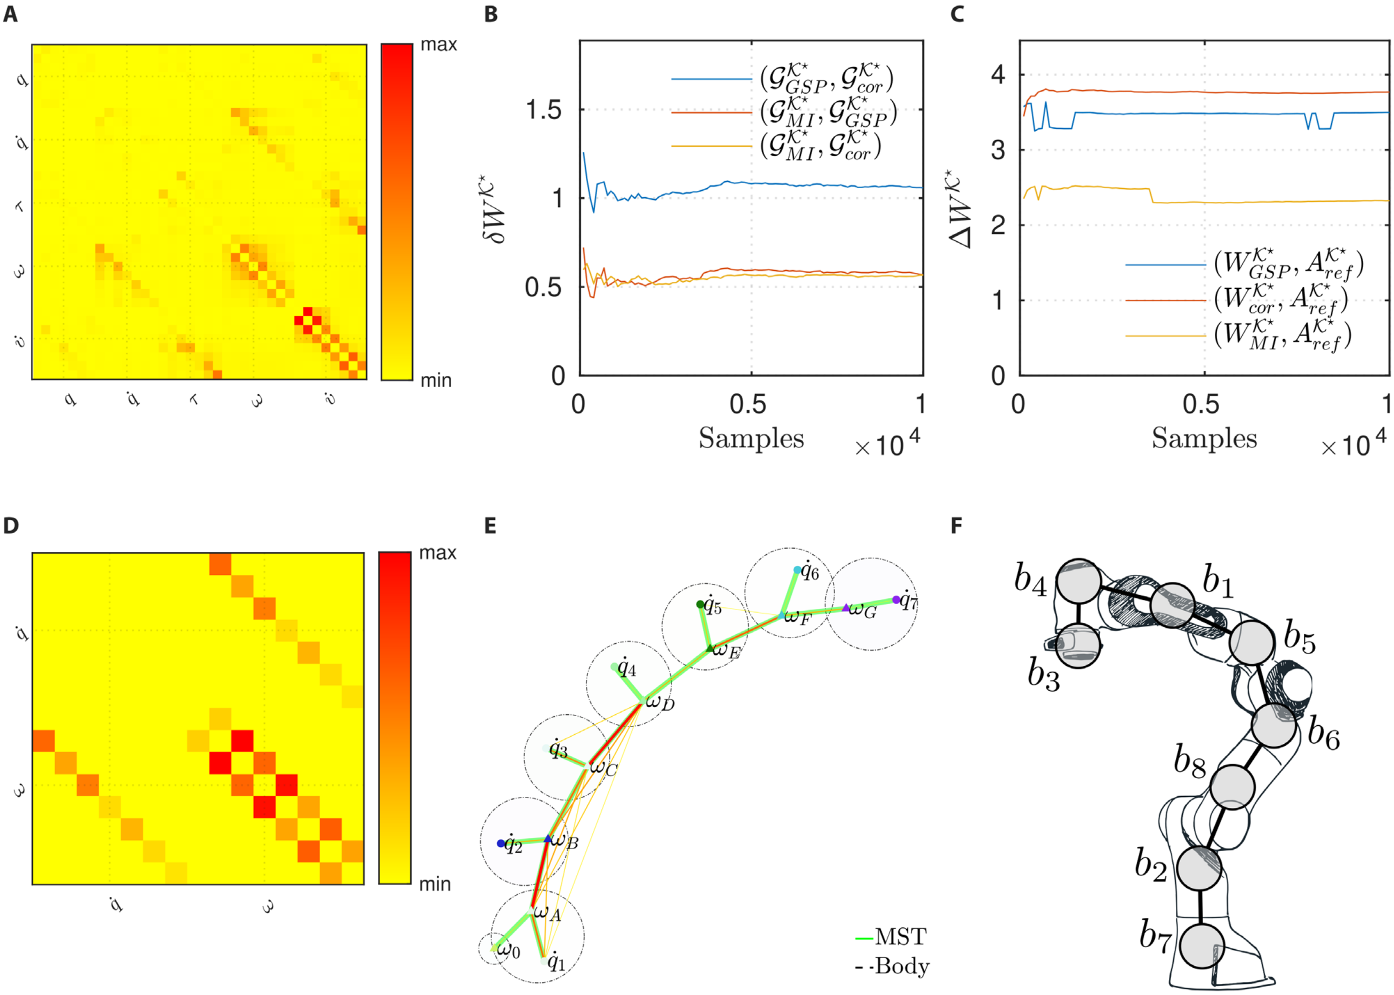

Fig. 3. Inferring body structure using the proprioceptive graph. (**A**) The MI matrix, (**B**) spectral and (**C**) Frobenius matrix distances of the kinematics subgraph, (**D**) the adjacency matrix of the subgraph $\mathcal{G}_\pi^{\mathcal{K}*}$, and (**E**) its graph representation. (**F**) The connectivity of the clusters corresponds to the topology of the robot.

graphs via the adjacency spectral distance (*44*) $\delta\boldsymbol{W}$. It can be seen that the number of samples seems to have a null effect on the similarity. $\mathcal{G}^{\mathcal{K}*}_{\mathrm{GSP}}$ and $\mathcal{G}^{\mathcal{K}*}_{\mathrm{MI}}$ appear to share local structural resemblance. In addition, to determine whether the subgraphs $\mathcal{G}^{\mathcal{K}*}_{\mathrm{MI}}$, $\mathcal{G}^{\mathcal{K}*}_{\mathrm{cor}}$, and $\mathcal{G}^{\mathcal{K}*}_{\mathrm{GSP}}$ are consistent with the mechanical connectivity of the robot, we looked at the Frobenius matrix distance (shown in Fig. 3C) between their adjacency matrices and the (binary) ground-truth adjacency matrix $\boldsymbol{A}^{\mathcal{K}*}_{\mathrm{ref}}$. This distance is more suitable to compare the pairwise affinities captured by the adjacency matrices. Evidently, the high dissimilarity of $\mathcal{G}^{\mathcal{K}*}_{\mathrm{cor}}$ shows that it is not suitable for topology inference. Moreover, although $\mathcal{G}^{\mathcal{K}*}_{\mathrm{GSP}}$ and $\mathcal{G}^{\mathcal{K}*}_{\mathrm{MI}}$ share a high spectral similarity, only the latter is sufficiently structurally close to $\boldsymbol{A}^{\mathcal{K}*}_{\mathrm{ref}}$ to justify its use to infer the mechanical topology of the robot.

The manipulator's subgraph $\mathcal{G}^{\mathcal{K}*}_{\mathrm{MI}}$, hereafter referred as $\mathcal{G}^{\mathcal{K}*}_{\pi}$, reflects important information from its topology. Its edges capture potential relationships between a rigid body (represented by its angular velocity $\boldsymbol{\omega}$) and its driving joint (indicated by the joint velocity $\dot{q}$). They also hint at the connectivity between the bodies in the chain. Consequently, from the edges of $\mathcal{G}^{\mathcal{K}*}_{\pi}$, the body-joint tuples b_i

$$(\dot{q}_7,\boldsymbol{\omega}_G)\in b_3\ (\dot{q}_6,\boldsymbol{\omega}_F)\in b_4\ (\dot{q}_5,\boldsymbol{\omega}_E)\in b_1\ (\dot{q}_4,\boldsymbol{\omega}_D)\in b_5\ (\dot{q}_3,\boldsymbol{\omega}_C)\in b_6$$
$$(\dot{q}_2,\boldsymbol{\omega}_B)\in b_8\ (\dot{q}_1,\boldsymbol{\omega}_A)\in b_2\ (\boldsymbol{\omega}_0)\in \underbrace{b_7}_{\text{root link}} \qquad (4)$$

were identified (depicted as dashed circles in Fig. 3E). This leads to the conclusion that the body topology is potentially defined by the following body-to-body connectivity

$$b_7\sim b_2\sim b_8\sim b_6\sim b_5\sim b_1\sim b_4\sim b_3$$

which matches the topology of the manipulator (see Fig. 3F). From this topology—represented by the binary adjacency matrix $\boldsymbol{A}^{\mathcal{H}*}_{\pi}$—and that the only body in Eq. 4 that is not associated with a joint variable is the robot's base (body b_7 in this case), the direction of propagation of angular the velocity was deduced to be

$$\boldsymbol{\omega}_0\rightarrow\boldsymbol{\omega}_A\rightarrow\boldsymbol{\omega}_B\rightarrow\boldsymbol{\omega}_C\rightarrow\boldsymbol{\omega}_D\rightarrow\boldsymbol{\omega}_E\rightarrow\boldsymbol{\omega}_F\rightarrow\boldsymbol{\omega}_G\rightarrow\boldsymbol{\omega}_H \qquad (5)$$

With this information, the kinematic description of the robot was characterized by defining the sensor-to-sensor rotation matrices ${}^{\{P\}}\hat{\boldsymbol{R}}_{\{C\}}$ between an IMU located on a parent body b_p and on a child body b_c, as well as the orientation of the rotation axes ${}^{\{C\}}\hat{\boldsymbol{\zeta}}_C$ corresponding to the jth joint that connects the bodies, together with their location $({}^{\{P\}}\hat{r}_j, {}^{\{C\}}\boldsymbol{r}_j)$—namely, the elements of the kinematic parameter vector $\hat{\boldsymbol{\lambda}}$ defined in Eq. 17. We estimated the manipulator's kinematics solving the problems in Eqs. 13 and 16 via offline optimization ($\hat{\boldsymbol{\lambda}}_{\mathrm{off}}$) and online learning ($\hat{\boldsymbol{\lambda}}_{\mathrm{on}}$). The results are presented in Fig. 4, where we also included the (offline) estimated kinematics for the case when the base is fixed. The estimation accuracy, computed according to Eqs. 18 and 21, is shown in Fig. 4 (B and C). When the base is excited, offline and online estimation provide comparable results. However, there is a stark contrast between the results of a movable and a fixed base. In the latter, its lack of motion and the robot having only revolute joints result in insufficient information to estimate the kinematic parameters of the proximal links. This is illustrated in the red highlighted area in Fig. 4A, where the estimated sensor-to-sensor rotation matrices and the estimated joint-center-point (JCP) location vectors for the first two bodies were inaccurate, as indicated by the relatively large errors depicted in Fig. 4 (B and C) (for the fixed base case). The origin of this problem is addressed in Discussion.

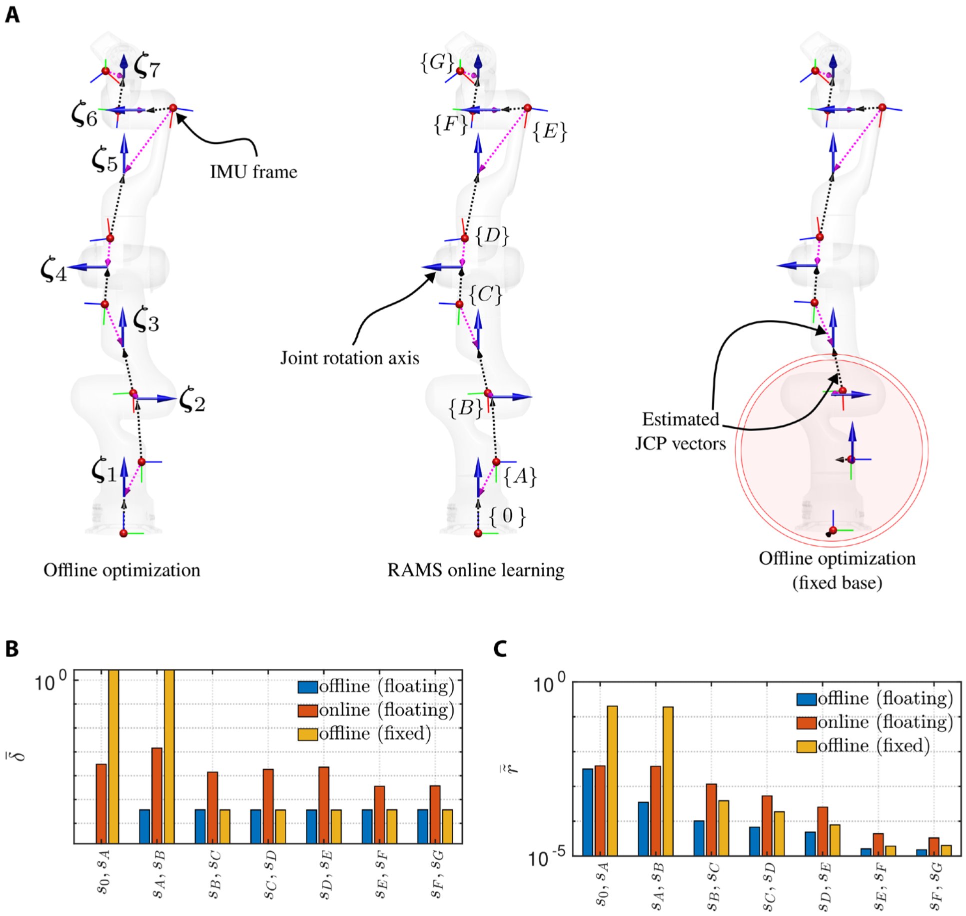

Fig. 4. The inferred kinematic structure of the Franka Emika manipulator. (**A**) The structure was obtained via offline optimization, RAMS online learning, and offline with a motionless base. The red spheres in the images represent the location of the IMUs, and the solid blue arrows indicate the orientation and position of the joint rotation axes, determined by the JCP vectors (dashed arrows). Comparisons between the real and learned (**B**) sensor-to-sensor rotation matrices and (**C**) location vectors.

Physical experiment

To confirm the results obtained in simulation, we also estimated the body structure using the real Franka Emika robot arm. The inferred morphology is shown in Fig. 5. The motion of the base, or lack thereof, does not seem to affect the relationships that the MI is able to determine for the bodies (excluding the base itself, naturally). This is observed in Fig. 5B, where $\mathcal{G}_{\pi}^{\mathcal{K}*}$ contains the correct body-to-body and body-to-joint relationships (see Fig. 5A). In contrast, a motionless base does influence the quality of the estimated parameter vector $\widehat{\boldsymbol{\lambda}}$—estimated offline in this case—as shown in Fig. 5 (C and D). With a fixed base, the proximal bodies lack proper excitation, a fact that is reflected in the incorrect location and orientation of the rotation axes corresponding to the proximal joints. Yet, when the spatial motion of the bodies is sufficient (as is the case for the distal bodies of the chain), the estimation is successful. This is observed in the relatively small errors associated to the distal links. In summary, the topology and kinematic estimation obtained in the physical experiment are a replica of those obtained in simulation and are consistent with the actual structure of the manipulator.

More complex body structures

Further testing confirmed that π-graphs serve to infer more complex tree-structure topologies. The first scenario contemplated a simulated PhantomX hexapod robot (18 DoF). The IMU signals are labeled from 0 to R and joint signals from 1 to 18. The generation and collection of proprioceptive signals resulted from motor babbling with the robot on the floor. Contacts occurred between the hexapod, the environment, and itself; joint limits were sometimes reached. However, our results show that these events did not hinder the estimation of the π-graph. Motion of the extremities needs to excite the base of the robot (through contacts with the ground) so that the observed base acceleration and angular velocity can be used to determine body-body pairings between the base and its adjacent links. Figure 6 depicts the hexapod's estimated topology and kinematic description. As shown in Fig. 6B, from the subgraph $\mathcal{G}_{\pi}^{\mathcal{K}*}$, the joint-to-body and body-to-body pairings are identifiable. Using the inferred $A_{\pi}^{\mathcal{K}*}$, the parameter vector $\widehat{\boldsymbol{\lambda}}$ was computed, leading to the morphology shown in Fig. 6C. To assess the quality of the estimated structure, Fig. 6 (D and E) shows the estimation errors for the sensor-to-sensor rotation matrices and position vectors for the offline and online learning cases. The results show accuracy on the order of 10^{-3} m for the JCP vectors and a difference rotation within the order of 10^{-2} radians.

A second experiment considered the 25-DoF humanoid Poppy. The mechanical arrangement of its trunk is such that the intersection of the joint axes of the two joints proximal to the pelvis and those of the joints that actuate the robot's chest (highlighted in yellow in Fig. 7A) results in virtual two-DoF joints. Figure 7B shows the inferred topology from the connectivity of $\mathcal{G}_{\pi}^{\mathcal{K}*}$. Albeit mostly coincident with the robot's topology (Fig. 7A), the experiments consistently resulted in the MI misidentifying some of the relationships in the torso, as shown by the two joint velocity signals associated with a single IMU (red areas in Fig. 7B). Such an effect seems to stem from the fact that the bodies in the circled areas operate as a unit driven by a two-DoF joint. This particular joint arrangement has an effect on the determination of the unique joint-body relationships, because the

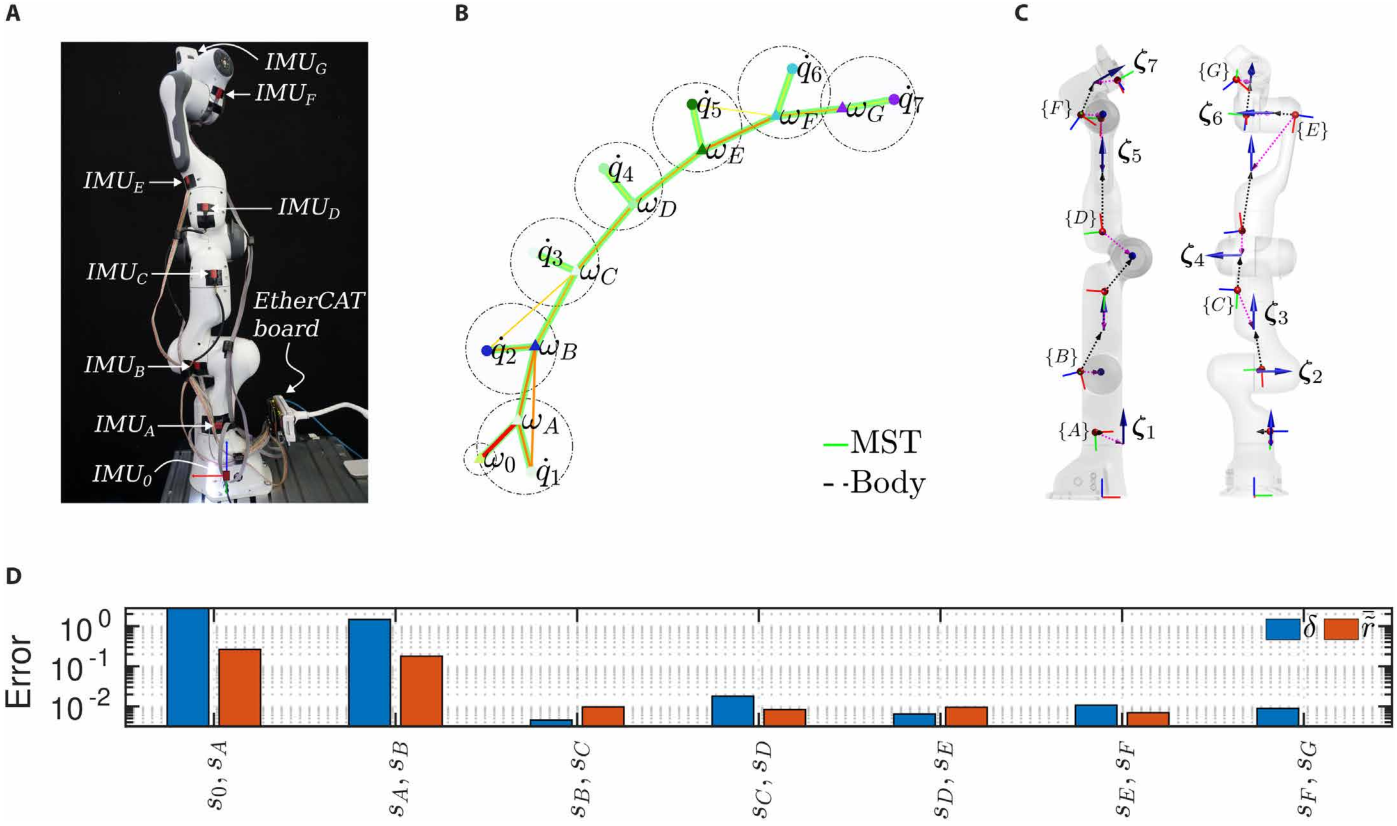

Fig. 5. Experimental results from the robot manipulator. (**A**) The robot arm equipped with IMUs, (**B**) its topology expressed by the $\mathcal{G}_{\pi}^{\mathcal{K}}$ graph, (**C**) the offline-learned kinematic description (fixed base), and (**D**) errors corresponding to the sensor-to-sensor rotation matrices and joint axis location vectors.

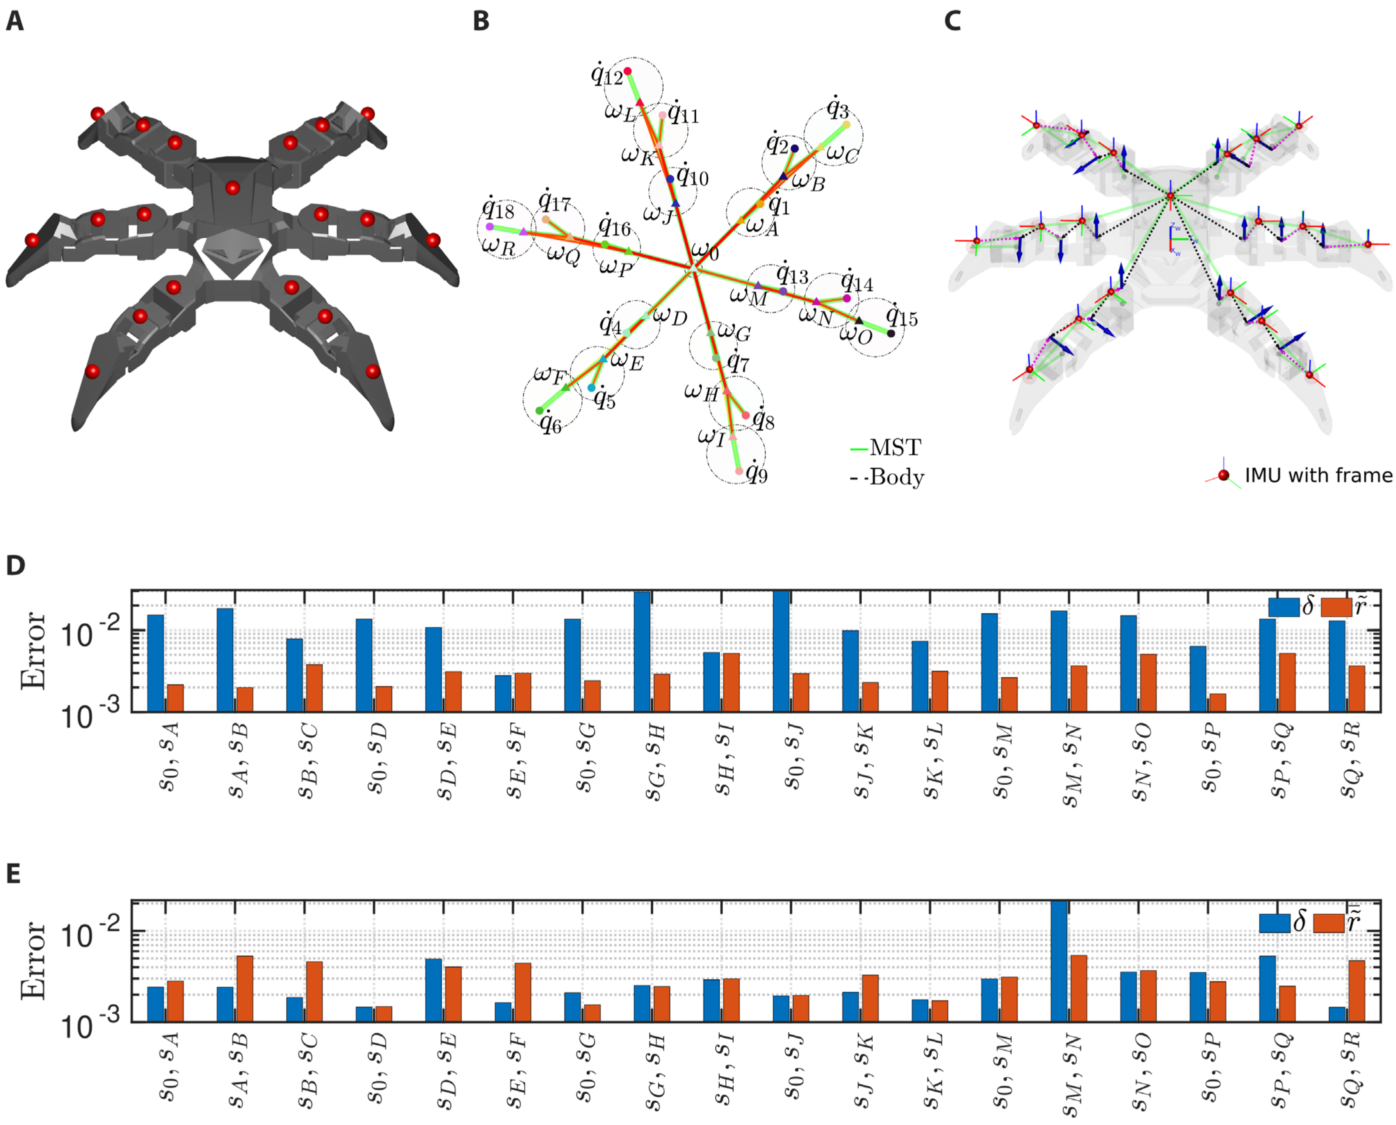

Fig. 6. The hexapod robot. (**A**) The PhantomX hexapod robot with its IMUs (red spheres), (**B**) the kinematics graph $\mathcal{G}_\pi^{\mathcal{K}}$, (**C**) the learned kinematic structure, and errors in the estimated body structure via (**D**) offline learning and (**E**) online learning (blue: difference rotation errors $\bar{\delta}$ [rad], orange: sensor-to-sensor vector errors $\bar{\tilde{r}}[m]$).

pairs of bodies corresponding to IMUs (s_K, s_L) and (s_N, s_O) are ill-defined. In addition, $\mathcal{G}_\pi^{\mathcal{K}*}$ shows bodies other than the root link (the robot pelvis) without driving joints, which presents a problem when using the expressions in Eqs. 10 and 12, because the corresponding joint variables are missing. One practical way to find these missing joint variables is to refer to the additional connections in the superseding graph $\mathcal{G}_\pi^{\mathcal{K}}$ to identify other potential driving joint candidates. Figure S9 shows the MI values between the trunk angular and joint velocities. Because one of those relationships ($\boldsymbol{\omega}_M$, $\dot{q}_{13}$) was unambiguously identified, the joint velocity to use in Eq. 12 was determined using a straightforward elimination process. Last, after resolving the inconsistencies in the joint-to-body pairings, the kinematic parameter vector $\hat{\boldsymbol{\lambda}}$ for the humanoid robot was estimated following the offline approach. Figure 7C shows the resulting kinematic description. According to the errors depicted in Fig. 7 (D and E), it matches very closely that of the actual robot.

DISCUSSION

Our results support that two essential aspects of the robot's body morphology, topology and kinematic description, can be learned from its sensorimotor proprioceptive signals through the analysis of π-graphs. Furthermore, they show that the MI between these signals can effectively reflect the embodiment of the robot. The experiments also confirm that, for body structure discovery, MI-based graph connectivity inference is superior to other state-of-the-art NTI methods. In addition, we demonstrated that from the connectivity of the kinematic proprioceptive subgraph $\mathcal{G}_\pi^{\mathcal{K}}$, the body topology of the robot could be extracted and used to deduce the correct dependency relationships in the differential kinematic laws, avoiding unnecessary tests when learning the kinematic description. The results showed—except for the fixed-base case—that exteroception (vision in particular) is optional to estimate fundamental aspects of the body morphology, and we can prescind from it and still achieve a

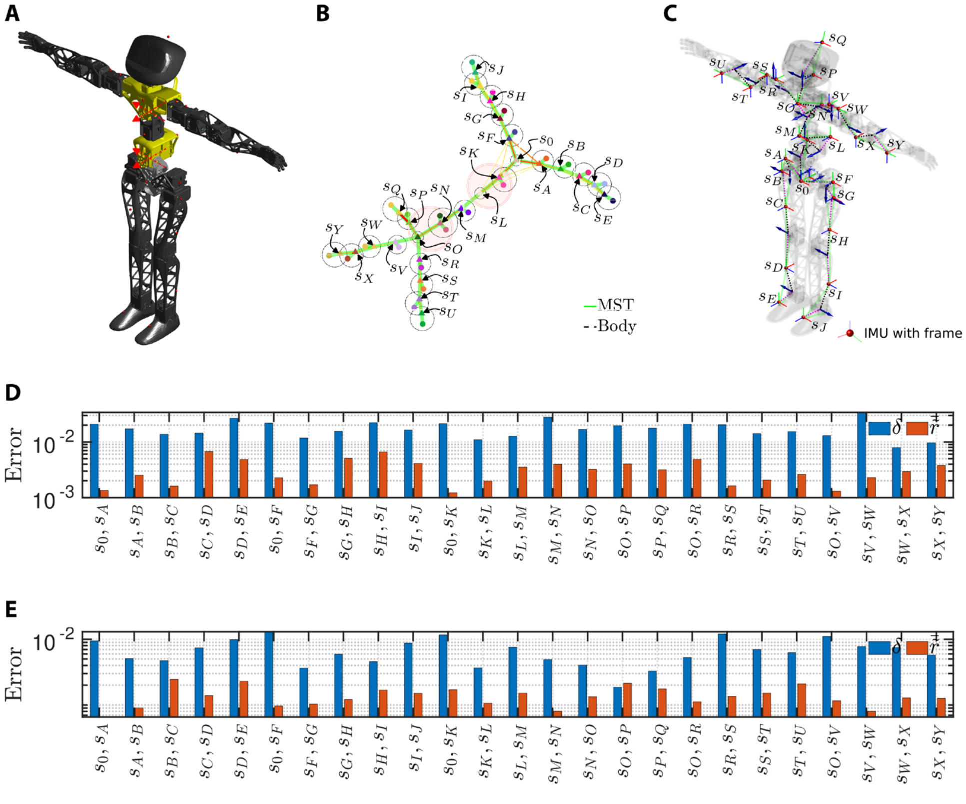

Fig. 7. The humanoid robot. (**A**) The Poppy humanoid robot with its IMUs (red spheres), (**B**) the proprioceptive kinematics graph $\mathcal{G}_\pi^{\mathcal{K}}$ highlighting the MST that shows the robot topology, (**C**) the learned morphology, and (**D**) offline and (**E**) online errors in the estimated structure (blue: $\bar{\delta}$ [rad], orange: $\bar{r}[m]$).

reasonably accurate morphological representation. Nevertheless, notwithstanding the promising results of the presented method, certain limiting aspects deserve special attention.

Robot structure and type of joints

Our method infers the morphology of robots defined by a serial kinematic tree structure. For other mechanical arrangements, like those of parallel robots, no guarantees exist about the successful inference of the topology and associated kinematic description. Furthermore, we only considered revolute joints with one DoF. As seen in the humanoid case, joints with multiple DoF can complicate the inference of the robot's topology via MI and require additional search routines and decision criteria to determine the correct kinematic relationships.

The influence of measurement noise

We artificially distorted the measurements for all the simulated experiments by injecting zero-mean additive white Gaussian noise with a signal-to-noise ratio (SNR) of 20 dB. This noise was comparable to that observed from the gyroscope and accelerometer measurements in the physical robot experiment. Tests with various SNR levels suggested that sensor noise did not affect the estimation of the robot topology, as shown in fig. S10. However, as expected, noise alters the estimated kinematic description of the robot because the estimation quality naturally depends on the quality of the measurements. Sensor noise influences the estimated orientation of the IMU coordinate frames and joint rotation axes and the location of the estimated JCPs. Efficient data fusion and filtering techniques can improve the estimation of the kinematic parameter vector $\widehat{\boldsymbol{\lambda}}$. For instance, multisensor data fusion techniques have been shown to decrease the noise power of accelerometer readings, thereby enhancing the SNR (*45*).

Motion and MI

Incremental learning of the body morphology implies that the robot needs to determine the connectivity of its π-graph via joint torque commands that generate motions to increase the MI between pairs of relevant proprioceptive signals. Potentially adaptable methods to achieve this have been discussed recently (*46*, *47*). Unexpectedly, our observations suggested that the π-graph can be effectively inferred without using elaborate methods to generate exploratory motions and relying only on motor babbling. We suspect that this is related to the phenomenon that the MI of signals resulting from random motions reflects the relationships associated with the body morphology. In contrast, structured motions (regular patterns such as those corresponding to locomotion or reaching motions) would obscure these relationships and reflect the information sharing due to the motion policy itself. Last, handling joint limits and managing potential self-collisions in physical robots, especially when no model is available, is still an open problem limited not only to babbling strategies.

Sampling efficiency and the accuracy of $\mathcal{G}_\pi^{\mathcal{K}*}$

The π-graph, the kinematics subgraph $\mathcal{G}_\pi^{\mathcal{K}}$, and its associated MST $\mathcal{G}_\pi^{\mathcal{K}*}$ are extracted only after the information content of $\widehat{W}_{\mathrm{MI}}$ stabilizes (see fig. S7). As observed in numerous simulation experiments, a large percentage of the edges of $\mathcal{G}_\pi^{\mathcal{K}*}$ match the actual topology early in the data collection process (see fig. S8). Furthermore, our analysis suggests that after about 300,000 samples, the structure of the MST does not change anymore. However, these results depend on the variability of the data samples used to compute $\widehat{W}_{\mathrm{MI}}$. Consequently, the sensibility of the convergence and accuracy of the identified graph to the size and variability of the replay buffer $\mathcal{B}$ and its associated mini-batch size $N_{\mathcal{X}}$ are yet to be characterized. Furthermore, it is not immediately possible to determine whether $\mathcal{G}_\pi^{\mathcal{K}*}$ describes the correct robot topology. This graph guides the learning of $\widehat{\boldsymbol{\lambda}}$ but is also validated by it while solving the optimization problems given in Eqs. 13 and 16. If there are body-to-body or body-to-joint relationships in $\mathcal{G}_\pi^{\mathcal{K}*}$ that do not find kinematic validation, an alternative candidate edge needs to be drawn from $\mathcal{G}_\pi^{\mathcal{K}}$ (recall that $\mathcal{G}_\pi^{\mathcal{K}*} \subset \mathcal{G}_\pi^{\mathcal{K}} \subset \mathcal{G}_\pi$), and the kinematic validation needs to be executed again.

Observability of the parameters

As shown in Results, under our proposed scheme, tree-like robots with fixed bases and proximal bodies whose acceleration and angular velocity are either zero or do not span $\mathbb{R}^3$ will lead to kinematic parameters that are not observable. The rationale behind this can be understood by looking at Eqs. 10, 12, and 14. It is certainly possible to estimate the joint axis orientation in the frame of the corresponding IMU from its measured angular velocity, namely, vector ${}^{\{C\}}\hat{\boldsymbol{\zeta}}_C$ in Eq. 10. However, as illustrated in fig. S6, when it comes to the robot's proximal joints to a fixed base, there are limitations on the estimation of the sensor-to-sensor rotation matrices in Eq. 12 and the joint axis position vectors in Eq. 14, which determine the estimated JCP. For example, consider the simple one-DoF robot in fig. S6. Because its base is entirely immobile, its attached IMU cannot provide information neither to find ${}^{\{P\}}\hat{\boldsymbol{r}}_j$ nor to estimate the rotation matrix ${}^{\{C\}}\hat{\boldsymbol{R}}_{\{P\}}$. Similarly, the IMU on the first link, having limited DoFs because of the system structure, can only rotate about the first joint axis $\boldsymbol{\zeta}$. From its measurements, the radial vector $\boldsymbol{a}$ to the axis can be estimated, but they are insufficient to estimate the location vector ${}^{\{C\}}\hat{\boldsymbol{r}}_j$. This is precisely the reason for the errors reported in Fig. 4 (B and C) (for the fixed base case) and Fig. 5D. This lack of excitation affects only robots with a fixed base and does not influence mobile robots as long as enough motion is produced.

Comparison with end-to-end learning

One final area that deserves consideration is the use of end-to-end learning techniques, neural networks in particular, to achieve similar results to those presented here. One work that illustrates this philosophy (*24*) retrieves structural kinematic information from the variations in the measurements from IMUs randomly located on the body of serial manipulators. The measurements are fed through a convolutional neural network (CNN) trained on samples from a platform that generates serial manipulators of different kinematic properties. This work exemplifies some of the caveats that using neural networks would have. On the one hand, supervised learning techniques would necessitate a pool of morphologically characterized kinematic trees from which to learn; on the other, these would need to be varied enough to capitalize on the interpolation power of supervised techniques. For instance, CNNs were trained with more than 100,000 samples of only serial manipulators and achieved a success rate of 80% when identifying the structure. Considering tree-like structures would demand a far more extensive training set. Take, for example, m bodies. According to Cayley's formula (*48*), there are m^{m-2} possible tree-like robots that could be formed, each of which could have infinitely many kinematic variations. In contrast, topology inference based on MI does not rely on such a massive training dataset.

In summary, our work shows how from complementary proprioceptive signals, we have learned $\mathcal{G}_\pi$, $\boldsymbol{A}_\pi^{\mathcal{K}*}$, and $\hat{\boldsymbol{\lambda}}$, which are analogous to the elements $\mathcal{G}$, $\boldsymbol{A}$, and $\boldsymbol{\lambda}$ of the body schema set in Eq. 1. We also showed that these elements, provided sufficient excitation, can be readily learned online. Ultimately, our framework infers and can continuously adapt the kinematic description of the robot from proprioception only. Such a capability is essential to maintain forward and inverse kinematic models for different motion planning, control, and interaction strategies. Last, robots capable of autonomously and incrementally building an understanding of their morphology can monitor the state of their dynamics, adapt the representation of their body, and react to changes to it, for instance, for self-calibration and fault detection and isolation.

MATERIALS AND METHODS

Signal collection via motor babbling

This corresponds to subprocess 1 in Fig. 2. Empirical evidence showed that the robot's motions must result in proprioceptive signals with enough excitation to estimate a meaningful MI matrix. Conventional methods for designing excitation trajectories—see, for example, (*49*)—can only be used with body knowledge. In contrast, typical examples of body-agnostic exploration trajectories involve maximizing the entropy of the control policy (*50*). In our simulated experiments, we used a simple motor babbling strategy that proved effective at highlighting the morphology-relevant relationships via MI. In particular, each joint received a periodic torque command τ_j defined by

$$\tau_j(t) = \tau_{\max}[\boldsymbol{A}_1 \sin(\omega_0 t) + \boldsymbol{A}_2 \sin(2\omega_0 t) + \boldsymbol{A}_3 \sin(4\omega_0 t)] \qquad (6)$$

with the base frequency $\omega_0 = 2\pi/T$ and $\tau_{\max}$ being the maximum commanded torque. The amplitudes $\boldsymbol{A}_i$ were changed at random every $T = 3$ s and adapted to ensure that $|\boldsymbol{A}_1| + |\boldsymbol{A}_2| + |\boldsymbol{A}_3| = 1$.

Data storage

Maintaining a memory of previously seen data samples is done in subprocess 2 of Fig. 2. The number of data points n required to compute a reliable estimate of the MI between a pair of time series is large, because it needs to capture the probability distribution of the signals. If no limitations exist regarding memory, time, or computation power, the most straightforward strategy is the offline computation of the pairwise MI values after having sampled and stored a data matrix $\boldsymbol{X}$ with many points n. However, the computational demand increases rapidly, because it depends on the number of data points and the number of pairwise MI evaluations required to build the MI adjacency matrix $\widehat{\boldsymbol{W}}_{\text{MI}}$, which grows exponentially with the number of sensor signals. To ameliorate the computational load and to allow the online estimation of the MI matrix from streaming sensorimotor signals (and consequently, the incremental estimation of the body morphology), we introduce a replay buffer $\mathcal{B}$ that can host $N_{\mathcal{B}} = 10{,}000$ points. Data storage follows a reservoir sampling strategy to capture approximately uniformly distributed samples from the data stream history. At any given point during learning, a minibatch $\boldsymbol{\mathcal{X}}$ of $N_{\mathcal{X}} = 100$ samples can be sampled at random from $\mathcal{B}$. Using the minibatches $\boldsymbol{\mathcal{X}}$, we can incrementally approximate an $\widehat{\boldsymbol{W}}_{\text{MI}}$ comparable to that obtained from a large data matrix $\boldsymbol{X}$.

Proprioceptive information sharing

Subprocess 3 in Fig. 2 involves the computation of the pairwise information sharing among the sensorimotor signals. On the basis of the connection between embodiment and information structure, our hypothesis is that properties of the body morphology can be extracted by studying the information sharing among the proprioceptive signals $\boldsymbol{x}(t)$. For this purpose, information theory has provided model-free metrics that can capture both linear and nonlinear relationships between signals. Of particular interest to us is the MI, because it has been used in different contexts to represent and analyze the relationships between variables (*51*). The MI between two signals $I\,(X, Y)$ is understood as the amount by which a measurement of a random variable Y reduces the uncertainty of the random variable X (*52*); see Supplementary Methods for further details.

By computing the pairwise MI between the $\{x_i\}_{i=1}^{m}$ proprioceptive signals, an information-based relational matrix $\widehat{\boldsymbol{W}}_{\text{MI}} \in \mathbb{R}^{m \times m}$ can be constructed, where its (i, j) entries are given by

$$(\widehat{\boldsymbol{W}}_{\text{MI}})_{i,j} = I(x_i, x_j) \tag{7}$$

Notice that we use the "hat" notation to emphasize that this matrix is an estimate and there is no actual ground truth. Last, to simplify the analysis of $\widehat{\boldsymbol{W}}_{\text{MI}}$, we account for the three-dimensional nature of the IMUs' signals and their relation to other body and joint signals by introducing a contraction preprocessing step $\boldsymbol{\mathcal{W}}_{\text{MI}} = C(\widehat{\boldsymbol{W}}_{\text{MI}})$ in which the blocks from $\widehat{\boldsymbol{W}}_{\text{MI}}$ that correspond to acceleration and angular velocities are contracted into a single scalar by taking their Frobenius norm. Complementary details about the contraction process are provided in Supplementary Methods and fig. S3.

Mechanical connectivity and MI

Subprocess 4 in Fig. 2 takes as input the contracted MI matrix $\boldsymbol{\mathcal{W}}_{\text{MI}}$ to infer the robot topology. Our problem is centered around the study of the FC among the proprioceptive signals x_i, implying that we wish to find a symmetric adjacency matrix $\boldsymbol{W}$. The already found contracted MI matrix $\boldsymbol{\mathcal{W}}_{\text{MI}}$ serves this exact purpose; it is a symmetric adjacency matrix that encodes the relationships between the $\{x_i\}_{i=1}^{m}$ proprioceptive signals by means of their information sharing, namely, building the MI matrix is a form of information-based NTI. Considering that $\boldsymbol{\mathcal{W}}_{\text{MI}}$ is a dense nonnegative matrix, making its associated graph $\mathcal{G}_{\text{MI}}$ a complete graph, we argue that there exists a subgraph $\mathcal{G}_\pi \subset \mathcal{G}_{\text{MI}}$ (with adjacency matrix $\boldsymbol{W}_\pi$), which we denote as the proprioceptive information graph (π-graph), that meets the following assumption (assumption 2) that the π-graph is a sparse graph whose edges represent the most relevant connection among the $\{x_i\}_{i=1}^{m}$ proprioceptive signals.

To obtain $\boldsymbol{W}_\pi$ from $\boldsymbol{\mathcal{W}}_{\text{MI}}$, a threshold is needed to prune irrelevant connections from $\mathcal{G}_{\text{MI}}$. Typically, this is done by manually determining such a threshold. In our study, we look for the graph's MST, denoted $\boldsymbol{W}^*$ and defined as the set of edges that connects all nodes and whose edge weight sum is maximal. Its lowest value is used as a pruning threshold to define a sparse matrix $\boldsymbol{W}_\pi$; in other words

$$(\boldsymbol{W}_\pi)_{i,j} = \max(\min(\boldsymbol{W}^*), (\boldsymbol{\mathcal{W}}_{\text{MI}})_{i,j})$$

$\boldsymbol{W}_\pi$ defines the proprioceptive information graph $\mathcal{G}_\pi$, whose edges express the largest information sharing observed among the sensorimotor signals.

The found $\mathcal{G}_\pi$ is used to explore further properties of the robot's body by studying the subgraph $\mathcal{G}_\pi^{\mathcal{K}} \subset \mathcal{G}_\pi$ whose set of nodes corresponds to signals from angular velocity $\boldsymbol{\omega}$ and joint angular velocity $\dot{q}$. Consider a simple serial kinematic chain (see fig. S2A) consisting of three rigid bodies connected by two joints. Each rigid body has an IMU, denoted by s_{l+1}, s_l, and s_{l-1}. Because we rely on the assumption that the strongest relationships existing among the proprioceptive signals reflect the morphology of the robot, we hypothesize that, on average, the MI I_A between the signals of adjacent bodies is larger than the MI I_{NA} of signals from nonadjacent bodies. This implies that

$$\mathbb{E}[I(s_{l+1}, s_l)] > \mathbb{E}[I(s_{l+1}, s_{l-1})] \tag{8}$$

where the operator $\mathbb{E}[\,\cdot\,]$ denotes expectation. Therefore, considering expression in Eq. 8 and assumption 1 leads to the assumption (assumption 3) that the mechanical connectivity between the bodies in the kinematic chain (robot topology) could be inferred from the pairwise MI among the kinematic-related proprioceptive signals in $\mathcal{G}_\pi$.

The set of edges in $\mathcal{G}_\pi^{\mathcal{K}}$ represents the large MI values $I(\boldsymbol{\omega}_j, \dot{q}_k)$ between the kth joint angular velocity $\dot{q}_k$ and the jth body angular velocity $\boldsymbol{\omega}_j$ (see fig. S2B). This reflects the fact that, in a serial kinematic chain, the motion of a body depends primarily on its driving joint. Therefore, the entries $(W_\pi^{\mathcal{K}})_{j,k} = I(\boldsymbol{\omega}_j, \dot{q}_k)$ define tuples $(\dot{q}_k, \boldsymbol{\omega}_j)$ that indicate likely joint-to-body pairings. To determine the tuples that may correspond to actual physical connections, we identified the MST of $\boldsymbol{W}_\pi^{\mathcal{K}}$ (which we denote $\boldsymbol{W}_\pi^{\mathcal{K}*}$) and identified $N + 1$ clusters in $\mathcal{G}_\pi^{\mathcal{K}*}$ on the basis of the edges between body and joint velocities. Ideally, each of the identified $\{b_i\}_1^{N+1}$ clusters represents a body in the kinematic chain, defined as the tuple $(\dot{q}_k, \boldsymbol{\omega}_k) \in b_i$. Similarly, the subset of edges associated with the entries of $(\boldsymbol{W}_\pi^{\mathcal{K}*})_{i,j} = I(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j)$ linking pairs of body angular velocity measurements in $\mathcal{G}_\pi^{\mathcal{K}*}$ hints at body-to-body couplings, namely, the relationships between bodies in the kinematic chain (also exemplified in fig. S2B). Thus, if a body b_i is represented by a tuple consisting of its corresponding body angular velocity and the associated joint angular velocity $(\boldsymbol{\omega}_i, \dot{q}_t)$, then, on the basis of $I(\boldsymbol{\omega}_i, \boldsymbol{\omega}_j)$, it can be concluded that the bodies b_i and b_j are connected by a joint. This is expressed as $b_i \sim b_j$. Note that according to the tree-like assumption, these connections not only are one to one but also can be one to many, that is, a body connected to several bodies. From this discussion, it can be concluded that $\boldsymbol{W}_\pi^{\mathcal{K}*}$ reflects the mechanical topology of the robot. An example illustrating this process is provided in Supplementary Methods in fig. S4.

Body kinematic description

Subprocess 5 in Fig. 2 finds a kinematic description congruent with the topology in $\boldsymbol{A}_\pi^{\mathcal{K}*}$. The actual existence of a physical connection $b_P \sim b_C$ between a parent body b_P and a child body b_C is validated by evaluating the difference

$$\tilde{\boldsymbol{\omega}}_C = \boldsymbol{\omega}_C - \hat{\boldsymbol{\omega}}_C(\boldsymbol{\omega}_P, q_C, \dot{q}_C; \hat{\boldsymbol{\gamma}}) \leq \epsilon \tag{9}$$

between the measured angular velocity of the child body $\boldsymbol{\omega}_C$ and its estimated counterpart $\hat{\boldsymbol{\omega}}_C$ given by the differential kinematics expression

$$\hat{\boldsymbol{\omega}}_C(\boldsymbol{\omega}_P, q_C, \dot{q}_C; \hat{\boldsymbol{\gamma}}) = \underbrace{{}^{\{P\}}\hat{\boldsymbol{R}}_{\{C\}}^{T}(q_C; {}^{\{C\}}\hat{\boldsymbol{\xi}}_P, \hat{\phi}, {}^{\{C\}}\hat{\boldsymbol{\zeta}}_C)}_{1}{}^{\{P\}}\boldsymbol{\omega}_P + \dot{q}_C \underbrace{{}^{\{C\}}\hat{\boldsymbol{\zeta}}_C}_{2} \tag{10}$$

with

$$\hat{\boldsymbol{\gamma}} = [{}^{\{P\}}\hat{\boldsymbol{\xi}}_C^{T}\, \hat{\phi}\, {}^{\{C\}}\hat{\boldsymbol{\zeta}}_C^{T}]^{T} \tag{11}$$

Equation 10 models the propagation of angular velocity from body b_P to body b_C under the assumption that their connecting joint is revolute. The term "1" represents the estimated rotation matrix ${}^{\{P\}}\hat{\boldsymbol{R}}_{\{C\}}^{T}$ corresponding to the orientation of the IMU's coordinate system $\{P\}$ on body b_P relative to the coordinate system $\{C\}$ of the IMU on body b_C. This matrix is expressed in axis-angle form as

$$^{\{P\}}\hat{\boldsymbol{R}}_{\{C\}}(q_c;\hat{\boldsymbol{\gamma}}) = (\exp^{[^{\{P\}}\hat{\xi}_C]\hat{\phi}})(\exp^{[^{\{C\}}\hat{\zeta}_C]q_c}) \tag{12}$$

with the terms $\hat{\boldsymbol{\xi}} \in \mathbb{R}^3$ and $\hat{\boldsymbol{\zeta}} \in \mathbb{R}^3$ being unitary rotation axes and [•] denoting a vector expressed as a skew-symmetric matrix. Similarly, “2” corresponds to the estimated rotation axis of the joint that connects the two bodies, expressed in the coordinate system $\{C\}$ of the child body’s IMU. Because Eq. 10 assumes a direction of propagation of angular velocity for the body-body relationship $b_P \sim b_C$, which we denote as $\boldsymbol{\omega}_P \rightarrow \boldsymbol{\omega}_C$, this assumed directionality is considered correct if the difference given in Eq. 9 is below a threshold value ϵ, where ϵ accounts for the fact that noise in the IMU measurements hinders perfect equality. The optimal values for the estimated parameter vector $\hat{\boldsymbol{\gamma}}$ are found by solving a regression problem minimizing the mean squared error of Eq. 9 for a subset of samples $\mathcal{B} \subset X$ with cardinality $|\mathcal{B}| \leq n$, that is,

$$\begin{aligned} \min_{\hat{\boldsymbol{\gamma}}} \quad & J_{\tilde{\boldsymbol{\omega}}_C} = \frac{1}{2|\mathcal{B}|}\sum\nolimits_{k\in\mathcal{B}} \|\boldsymbol{\omega}_C(k) - \hat{\boldsymbol{\omega}}_C(\boldsymbol{\omega}_P(k), q_C(k), \dot{q}_C(k)); \hat{\boldsymbol{\gamma}}\|^2 \\ \text{subject to} \quad & \hat{\boldsymbol{\gamma}} \in \Gamma \end{aligned} \tag{13}$$

where Γ is the constraints set, namely, $\|^{\{P\}}\boldsymbol{\xi}_C\| = 1$, $\|^{\{C\}}\boldsymbol{\zeta}_C\| = 1$ and $\phi \in [0, \pi)$. Without knowledge about the robot’s mechanical topology, testing these parent-child relationships represents a combinatorial challenge in which the problem in Eq. 13 would need to be solved for every combination of body and joint angular velocities in Eq. 9. However, as a result of the analysis of the π-graph, only the entries of the connectivity matrix $\boldsymbol{A}_\pi^{\mathcal{K}*}$ need to be evaluated, because each entry relates two triplets of signals $(\boldsymbol{\omega}_p, q_p, \dot{q}_p) \ni b_{\text{pa}} \sim (\boldsymbol{\omega}_C, q_C, \dot{q}_C) \ni b_C$. Nonetheless, two possibilities exist for each body-body relationship: $\boldsymbol{\omega}_C \rightarrow \boldsymbol{\omega}_P$ and $\boldsymbol{\omega}_C \leftarrow \boldsymbol{\omega}_P$. Therefore, two solutions to the problem in Eq. 13 considering both cases are required to determine the actual directionality by evaluating $\min\left(J_{\tilde{\boldsymbol{\omega}}_C}, J_{\tilde{\boldsymbol{\omega}}_P}\right)$. A sufficient condition to deduce the direction of propagation is to use the adjacency matrix $\boldsymbol{A}_\pi^{\mathcal{K}*}$ and choose as the root of the kinematic chain the one body that lacks an association to a joint angular velocity, because this corresponds to the base link of the robot. Consequently, as a result of the tree-structure assumption, the direction of propagation of angular velocity can be directly inferred, and further testing is unnecessary. For example, we use the hypothetical robot in fig. S4. The root node corresponds to body b_2, implying that the root angular velocity is $\boldsymbol{\omega}_0$, and the velocity propagation directions are shown in fig. S4D.

At this point, the rotation axes’ location is still unknown. To completely characterize the robot’s kinematic description, they need to be translated to the actual location of the physical joint in the mechanical assembly. For this purpose, we introduce the notion of JCP, defined as the point $\boldsymbol{r}_j$ along the joint axis $\boldsymbol{\zeta}$ that is almost coincidental with the geometric center of the joint mechanical assembly. Many kinematically equivalent mechanisms can have a different JCP (see fig. S5A). Without exteroceptive information (e.g., touch or vision), it is not possible to unambiguously identify the actual JCP. However, an estimate of the JCP can be realized by finding the translation of the axes relative to the associated IMU coordinate frames $\{P\}$ and $\{C\}$ of the connected bodies. Note that the velocity and acceleration of any point $\hat{\boldsymbol{r}}_j$ located on the actual joint rotation axis $\boldsymbol{\zeta}$ must be the same in frames $\{P\}$ and $\{C\}$ (*53*) (see Eq. 14a). This means that we can express the axis translation by defining this point relative to the IMU frames, that is, by defining two position vectors $^{\{P\}}\hat{\boldsymbol{r}}_j$ and $^{\{C\}}\hat{\boldsymbol{r}}_j$. These vectors are estimated by using the expression for the propagation of acceleration in a kinematic tree, given in Eqs. 14b and 14c.

$$^{\{P\}}\hat{\dot{\boldsymbol{v}}}_j = {}^{\{P\}}\hat{\boldsymbol{R}}_{\{C\}}{}^{\{C\}}\hat{\dot{\boldsymbol{v}}}_j \tag{14a}$$

$$^{\{C\}}\hat{\dot{\boldsymbol{v}}}_j = {}^{\{C\}}\dot{\boldsymbol{v}}_C - ([^{\{C\}}\boldsymbol{\omega}_C][^{\{C\}}\boldsymbol{\omega}_C] + [^{\{C\}}\hat{\dot{\boldsymbol{\omega}}}_C])^{\{C\}}\hat{\boldsymbol{r}}_j \tag{14b}$$

$$^{\{P\}}\hat{\dot{\boldsymbol{v}}}_j = {}^{\{P\}}\dot{\boldsymbol{v}}_P - ([^{\{P\}}\boldsymbol{\omega}_P][^{\{P\}}\boldsymbol{\omega}_P] + [^{\{P\}}\hat{\dot{\boldsymbol{\omega}}}_P])^{\{P\}}\hat{\boldsymbol{r}}_j \tag{14c}$$

However, because any point $\hat{\boldsymbol{r}}_j$ satisfying Eq. 14 represents a valid translation, we introduce a constraint and define the estimated JCP as follows: The estimated JCP of the jth joint is a point $\hat{\boldsymbol{r}}_j$ along the joint axis $\boldsymbol{\zeta}$ that is a valid solution to Eq. 14 and whose location vectors $^{\{P\}}\hat{\boldsymbol{r}}_j$ and $^{\{C\}}\hat{\boldsymbol{r}}_j$ have minimal norms.

The estimated JCP is depicted in fig. S5B. It is worth mentioning that $^{\{P\}}\hat{\boldsymbol{r}}_j$ and $^{\{C\}}\hat{\boldsymbol{r}}_j$ together with the unit joint axis $^{\{P\}}\hat{\zeta}_{\{P\}}$ are analogous to the Plücker coordinates of the line defining the joint axis relative to the IMU coordinate frames.

Because the actual angular acceleration $\dot{\boldsymbol{\omega}}$ in Eq. 14 is not measured, numerical differentiation of the IMU-measured angular velocity (denoted $\hat{\dot{\boldsymbol{\omega}}}$) is used instead. Determining the estimated JCP implies minimizing the difference between Eq. 14; in other words

$$\tilde{\hat{\dot{\boldsymbol{v}}}}_j = {}^{\{P\}}\hat{\boldsymbol{R}}_{\{C\}}(q_C;\hat{\boldsymbol{\gamma}})^{\{C\}}\hat{\dot{\boldsymbol{v}}}_j - {}^{\{P\}}\hat{\dot{\boldsymbol{v}}}_j \leq \epsilon \tag{15}$$

Unlike the work in (*53*), which uses Eq. 15 based only on the difference of the norms of the accelerations in Eq. 14, we leverage the found sensor-to-sensor rotation matrices, given in Eq. 12, to have component-wise equality. Therefore, to determine $\boldsymbol{\rho}_j = [^{\{P\}}\hat{\boldsymbol{r}}_j^T, {}^{\{C\}}\hat{\boldsymbol{r}}_j^T]^T$, we solve

$$\min_{\boldsymbol{\rho}_j} \boldsymbol{J}_{\tilde{\dot{\boldsymbol{v}}}_j} = \frac{1}{2|\mathcal{B}|}\sum_{k\in\mathcal{B}}(\tilde{\hat{\dot{\boldsymbol{v}}}}_j(k))^2 + \beta(\boldsymbol{\rho}_j^T, \boldsymbol{\rho}_j) \tag{16}$$

where $\beta > 0$ enforces the minimum-norm constraint for the estimated JCP vectors.

In summary, the problem in Eq. 13 is solved first to find the orientation of the sensor frames, which are then used to solve Eq. 16 to find the location of the JCPs. The complete set of parameters describing the orientation of the IMU frames and the JCPs (and, by extension, the location of the IMU sensors) is collected in the vector

$$\hat{\boldsymbol{\lambda}}_j = \left[{}^{\{P\}}\hat{\boldsymbol{\xi}}_C^T \hat{\phi}\, {}^{\{C\}}\hat{\boldsymbol{\zeta}}_C^T\, {}^{\{P\}}\hat{\boldsymbol{r}}_j^T\, {}^{\{C\}}\hat{\boldsymbol{r}}_j^T\right]^T \tag{17}$$

Alternatively, both problems in Eqs. 13 and 16 can be solved in parallel using online learning via mini-batch stochastic gradient descent with the constraint set Γ enforced by adding penalizing terms to the cost function (for example, a barrier function). Nevertheless, because the axis vectors belong to the spherical manifold of dimension two

($\boldsymbol{\xi}, \boldsymbol{\zeta} \in \mathcal{S}^2$), we can use the Riemannian AMS (RAMS) gradient descent method with experience replay (*54*) to compute the updates to each $\widehat{\boldsymbol{\Upsilon}}_j$ without explicitly considering Γ. The method requires a buffer $\mathcal{B}$ to draw mini-batches $\mathcal{X}$ of samples at random to generate the Riemannian updates to $\widehat{\boldsymbol{\Upsilon}}_j$. Such a buffer is already available in our framework. In practice, for our experiments, we used offline optimization to learn a vector $\widehat{\boldsymbol{\lambda}}_{\text{off}}$ using standard interior-point constrained optimization with a batch of 1000 samples taken at random from the data matrix $\boldsymbol{X}$. Alternatively, we used online learning to find a vector $\widehat{\boldsymbol{\lambda}}_{\text{on}}$ via RAMS gradient descent leveraging the buffer $\mathcal{B}$ and mini-batches $\mathcal{X}$ of 50 samples.

Alternatives for NTI

To validate the use of MI as an NTI method that leads to the correct inference of the robot's mechanical topology, we considered two alternative NTI methods. The first technique (*18*) upgrades standard correlation-based NTI by searching for an inverse covariance matrix with Laplacian (to find valid adjacency matrices $\boldsymbol{W}_{\text{cor}}$) and structural constraints (requiring a sparse matrix to reduce the graph edge density). In addition, we consider a method (*55*) that learns a relational matrix $\boldsymbol{W}_{\text{GSP}}$ without prior structural information considering the signals in $\boldsymbol{x}(t)$ as graph signals (*28*). $\boldsymbol{W}_{\text{GSP}}$ is derived under the assumption that the signals on the graph change smoothly between connected nodes. We used the Graph Laplacian Learning (GLL) package (*56*) to calculate $\boldsymbol{W}_{\text{cor}}$, with regularization parameter $\gamma = 0.07$ and using a matrix $\boldsymbol{W}_0$ as connectivity prior with zero diagonal elements and ones elsewhere (denoting lack of structural knowledge). Likewise, the Graph Signal Processing Toolbox (GSPBOX) from (*57*) was used to compute $\boldsymbol{W}_{\text{GSP}}$ using $\alpha = 0.6$, $\beta = 1$ (parameters that control the edge weight magnitude and the sparsity of $\boldsymbol{W}_{\text{GSP}}$, respectively) and normalizing the required pairwise distance matrix $\boldsymbol{Z}$ between [0, 1].

Kinematic description metrics

Let $\boldsymbol{R}\,(\cdot)$ and $\widehat{\boldsymbol{R}}(\cdot\,;\widehat{\boldsymbol{\lambda}})$ be the actual sensor-to-sensor rotation matrix and its corresponding estimate from Eq. 12. To evaluate the distance between the estimates $\widehat{\boldsymbol{R}}$ and the ground truth $\boldsymbol{R}$, we use Eq. 18 to compute the angle $\delta_{j,k}$—measured in radians—of the difference rotation (for the kth sample of the jth joint) given by $\boldsymbol{R}\widehat{\boldsymbol{R}}^T$

$$\delta_{j,k} = \arccos\left(\frac{tr(\boldsymbol{R}(q_j(k))\widehat{\boldsymbol{R}}^T(q_j(k);\widehat{\lambda})) - 1}{2}\right) \tag{18a}$$

$$\bar{\delta}_j = \frac{1}{|\,\mathcal{B}\,|}\sum_{k\in\mathcal{B}}\delta_{j,k} \tag{18b}$$

Equation 18 needs to be evaluated for each of the N sensor-sensor pairs. A similar approach is considered for the JCP location vectors. We first define the configuration-dependent vector

$${}^{\{P\}}\widehat{\boldsymbol{r}}_{\overrightarrow{pc},j}(k) = -{}^{\{P\}}\widehat{\boldsymbol{r}}_j(k) + {}^{\{P\}}\widehat{\boldsymbol{R}}_{\{C\}}(\cdot)\,{}^{\{C\}}\widehat{\boldsymbol{r}}_j(k) \tag{19}$$

that expresses the estimated location vector $\boldsymbol{r}\{_{\overrightarrow{pc},j}(k)$ of the IMU s_c relative to the IMU s_p when their connecting joint is at an angle $q_c(k)$. Thus, there is a pairwise comparison for each of the j joints in the tree, that is,

$${}^{\{P\}}\tilde{\boldsymbol{r}}_j(k) = \|\,{}^{\{P\}}\widehat{\boldsymbol{r}}_{\overrightarrow{pc},j}(k) - {}^{\{P\}}\tilde{\boldsymbol{r}}_{\overrightarrow{pc},j}(k)\,\| \tag{20}$$

Consequently, the sensor location error per joint—in meters—for a set of samples $\mathcal{B}$ is simply the average of Eq. 20 over the samples expressed as

$$\bar{\tilde{r}}_j = \frac{1}{|\,\mathcal{B}\,|}\sum_{k\in\mathcal{B}}{}^{\{P\}}\tilde{r}_j(k) \tag{21}$$

Supplementary Materials

This PDF file includes:
Methods
Results
Figs. S1 to S10
References (*58–62*)

Other Supplementary Material for this manuscript includes the following:
Movie S1

REFERENCES AND NOTES

1. D. Pierce, B. Kuipers, Map learning with uninterpreted sensors and effectors. *Artif. Intell.* **92**, 169–227 (1997).
2. J. Bongard, V. Zykov, H. Lipson, Resilient machines through continuous self-modeling. *Science* **314**, 1118–1121 (2006).
3. M. Hoffmann, *Biologically Inspired Robot Body Models and Self-Calibration* (Springer, 2021), pp. 1–14.
4. P. Morasso, What is the use of the body schema for humanoid robots? *Int. J. Mach. Conscious.* **5**, 75–94 (2013).
5. D. De Santis, V. Mohan, P. Morasso, J. Zenzeri, Do humanoid robots need a body schema?, in *Biologically Inspired Cognitive Architectures 2012: Proceedings of the Third Annual Meeting of the BICA Society* (Springer, 2013), pp. 109–115.
6. P. Morasso, M. Casadio, V. Mohan, F. Rea, J. Zenzeri, Revisiting the body-schema concept in the context of whole-body postural-focal dynamics. *Front. Hum. Neurosci.* **9**, 83 (2015).
7. M. Hoffmann, H. Marques, A. Arieta, H. Sumioka, M. Lungarella, R. Pfeifer, Body schema in robotics: A review. *IEEE Trans Auton Ment Dev* **2**, 304–324 (2010).
8. R. Pfeifer, J. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press, 2006).
9. M. Mathew, R. Sapra, S. Majumder, A learning based approach to self modeling robots, in *2014 International Conference on Control, Instrumentation, Communication and Computational Technologies (ICCICCT)* (IEEE, 2014), pp. 758–762.
10. J. W. Hart, B. Scassellati, Robotic self-models inspired by human development, in *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence* (AAAI, 2010).
11. M. Hersch, E. Sauser, A. Billard, Online learning of the body schema. *Int. J. Humanoid Robot.* **5**, 161–181 (2008).
12. R. Martinez-Cantin, M. Lopes, L. Montesano, Body schema acquisition through active learning, in *2010 IEEE International Conference on Robotics and Automation* (IEEE, 2010), pp. 1860–1866.
13. J. Sturm, C. Plagemann, W. Burgard, in *Towards Service Robots for Everyday Environments* (Springer, 2012), pp. 131–161.
14. B. Chen, R. Kwiatkowski, C. Vondrick, H. Lipson, Fully body visual self-modeling of robot morphologies. *Sci. Robot.* **7**, eabn1944 (2022).
15. G. Schillaci, V. V. Hafner, B. Lara, Exploration behaviors, body representations, and simulation processes for the development of cognition in artificial agents. *Front. Robot. AI* **3**, 39 (2016).
16. L. Montesano, M. Lopes, A. Bernardino, J. Santos-Victor, Learning object affordances: From sensory–motor coordination to imitation. *IEEE Trans. Robot.* **24**, 15–26 (2008).
17. M. Lungarella, O. Sporns, Mapping information flow in sensorimotor networks. *PLOS Comput. Biol.* **2**, e144 (2006).
18. L. A. Olsson, C. L. Nehaniv, D. Polani, From unknown sensors and actuators to actions grounded in sensorimotor perceptions. *Connect. Sci.* **18**, 121–144 (2006).
19. N. Schmidt, "Sensorimotor contingencies in artificial intelligence and robotics," thesis, University of Zurich (2015).
20. P. Lanillos, E. Dean-Leon, G. Cheng, Enactive self: A study of engineering perspectives to obtain the sensorimotor self through enaction, in *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)* (IEEE, 2017), pp. 72–78.
21. R. Kwiatkowski, H. Lipson, Task-agnostic self-modeling machines. *Sci. Robot.* **4**, eaau9354 (2019).

22. J. A. Escalera, F. J. Abu-Dakka, M. Abderrahim, Symbolic geometric modelling of tree-structure robotic mechanisms using lie groups and graph theory. *Asian J. Control* **20**, 2088–2100 (2018).
23. N. D'Amore, C. Ciarleglio, D. L. Akin, Imu-based manipulator kinematic identification, in *2015 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2015), pp. 1437–1441.
24. M. M. Aref, J. Mattila, Deep learning of robotic manipulator structures by convolutional neural network, in *2018 Ninth International Conference on Intelligent Control and Information Processing (ICICIP)* (IEEE, 2018), pp. 236–242.
25. N. M. Schmidt, M. Hoffmann, K. Nakajima, R. Pfeifer, Bootstrapping perception using information theory: Case studies in a quadruped robot running on different grounds. *ACS* **16**, 1250078 (2013).
26. D. Polani, M. Möller, *Information Theory and Statistical Learning* (Springer, 2009), pp. 289–308.
27. T. Bossomaier, L. Barnett, M. Harré, J. T. Lizier, *An Introduction to Transfer Entropy* (Springer International Publishing, 2016).
28. X. Dong, D. Thanou, M. Rabbat, P. Frossard, Learning graphs from data: A signal representation perspective. *IEEE Signal Process. Mag.* **36**, 44–63 (2019).
29. L. Stanković, M. Daković, E. Sejdić, *Vertex-Frequency Analysis of Graph Signals* (Springer, 2019), pp. 3–108.
30. W. Zhang, J. Chien, J. Yong, R. Kuang, Network-based machine learning and graph theory algorithms for precision oncology. *NPJ Precis. Oncol.* **1**, 25 (2017).
31. W. Karwowski, F. Vasheghani Farahani, N. Lighthall, Application of graph theory for identifying connectivity patterns in human brain networks: A systematic review. *Front. Neurosci.* **13**, 585 (2019).
32. O. Sporns, Graph theory methods: Applications in brain networks. *Dialogues Clin. Neurosci.* **20**, 111–121 (2018).
33. K. J. Friston, Functional and effective connectivity: A review. *Brain Connect.* **1**, 13–36 (2011).
34. A. M. Bastos, J.-M. Schoffelen, A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Front. Syst. Neurosci.* **9**, 175 (2016).
35. H.-J. Park, K. Friston, Structural and functional brain networks: From connections to cognition. *Science* **342**, (2013).
36. U. Proske, S. C. Gandevia, The proprioceptive senses: Their roles in signaling body shape, body position and movement, and muscle force. *Physiol. Rev.* **92**, 1651–1697 (2012).
37. R. Siegwart, I. R. Nourbakhsh, D. Scaramuzza, *Introduction to Autonomous Mobile Robots* (MIT Press, 2011).
38. A. De Luca, D. Schroder, M. Thummel, An acceleration-based state observer for robot manipulators with elastic joints, in *Proceedings 2007 IEEE International Conference on Robotics and Automation* (IEEE, 2007), pp. 3817–3823.
39. S. A. B. Birjandi, J. Kühn, S. Haddadin, Joint velocity and acceleration estimation in serial chain rigid body and flexible joint manipulators, in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2019), pp. 7503–7509.
40. N. Rotella, S. Mason, S. Schaal, L. Righetti, Inertial sensor-based humanoid joint state estimation, in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2016), pp. 1825–1831.
41. S. Dutta, et al., Stability analysis of humanoid robots with gyro sensors subjected to external push forces, in *2019 2nd International Symposium on Devices, Circuits and Systems (ISDCS)* (IEEE, 2019), pp. 1–4.
42. M. Hamad, A. Kurdas, N. Mansfeld, S. Abdolshah, S. Haddadin, Modularize-and-Conquer: A generalized impact dynamics and safe precollision control framework for floating-base tree-like robots. *IEEE Trans. Robot.* **39**, 3200–3221 (2023).
43. S. Haddadin, S. Parusel, L. Johannsmeier, S. Golz, S. Gabl, F. Walch, M. Sabaghian, C. Jahne, L. Hausperger, S. Haddadin, The Franka Emika Robot: A reference platform for robotics research and education. *IEEE Robot. Autom. Mag.* **29**, 46–64 (2022).
44. P. Wills, F. G. Meyer, Metrics for graph comparison: A practitioner's guide. *PLOS ONE* **15**, e0228728 (2020).
45. S. A. Baradaran Birjandi, E. Pozo Fortunić, S. Haddadin, Evaluation of robot manipulator link velocity and acceleration observer. *IFAC-PapersOnLine* **56**, 292–299 (2023).
46. G. Martius, R. Der, N. Ay, Information driven self-organization of complex robotic behaviors. *PLOS ONE* **8**, e63400 (2013).
47. F. Gama, M. Shcherban, M. Rolf, M. Hoffmann, Goal-directed tactile exploration for body model learning through self-touch on a humanoid robot. *IEEE Trans. Cogn. Develop. Syst.* **15**, 419–433 (2021).
48. A. Shukla, A short proof of Cayley's tree formula. *Am. Math. Mon.* **125**, 65–68 (2018).
49. K.-J. Park, Fourier-based optimal excitation trajectories for the dynamic identification of robots. *Robotica* **24**, 625–633 (2006).
50. T. Haarnoja, A. Zhou, P. Abbeel, S. Levine, Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, in *International Conference on Machine Learning* (PMLR, 2018), pp. 1861–1870.
51. R. Steuer, J. Kurths, C. O. Daub, J. Weise, J. Selbig, The mutual information: Detecting and evaluating dependencies between variables. *Bioinformatics* **18**, S231–S240 (2002).
52. T. M. Cover, *Elements of Information Theory* (John Wiley & Sons, 1999).
53. T. Seel, J. Raisch, T. Schauer, IMU-based joint angle measurement for gait analysis. *Sensors* **14**, 6891–6909 (2014).
54. F. Diaz Ledezma, S. Haddadin, RIL: Riemannian incremental learning of the inertial properties of the robot body schema, in *2021 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2021), pp. 9354–9360.
55. V. Kalofolias, How to learn a graph from smooth signals, in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics* (PMLR, 2016), pp. 920–929.
56. H. E. Egilmez, Graph Laplacian Learning (GLL) package v2.1 (2021); https://github.com/STAC-USC/Graph_Learning.
57. N. Perraudin, J. Paratte, D. Shuman, L. Martin, V. Kalofolias, P. Vandergheynst, D. K. Hammond, GSPBOX: A toolbox for signal processing on graphs. arXiv:1408.5781 [cs.IT] (2014).
58. G. Mateos, S. Segarra, A. G. Marques, A. Ribeiro, Connecting the dots: Identifying network structure via graph signal processing. *IEEE Signal Process. Mag.* **36**, 16–43 (2019).
59. J. Walters-Williams, Y. Li, Estimation of mutual information: A survey, in *International Conference on Rough Sets and Knowledge Technology* (Springer, 2009), pp. 389–396.
60. J. T. Lizier, JIDT: An information-theoretic toolkit for studying the dynamics of complex systems. *Front. Robot. AI* **1**, 11 (2014).
61. F. Pedregosa, Scikit-learn: Machine learning in python. *J. Mach Learn Res.* **12**, 2825 (2011).
62. H. Peng, Mutual information computation; www.mathworks.com/matlabcentral/fileexchange/14888-mutual-information-computation.

Acknowledgments: We thank F. Wu from the Munich Institute of Robotics and Machine Intelligence for his feedback and support throughout the research process. **Funding:** We acknowledge the funding of this work by the Alfried Krupp von Bohlen und Halbach Foundation, the Lighthouse Initiative Geriatronics by StMWi Bayern (Project X, grant no. 5140951), the Lighthouse Initiative KI.FABRIK Bayern by StMWi Bayern, Forschungsund Entwicklungsprojekt, grant no. DIK0249, and Phase 1: Aufbau Infrastruktur, as well as the Federal Ministry of Education and Research of the Federal Republic of Germany (BMBF), by funding the project AI.D under project number 16ME0539K. **Author contributions:** The concepts were developed by F.D.L. and S.H. F.D.L. implemented and conducted all of the experiments and analyzed the data. F.D.L. and S.H. interpreted the results. F.D.L. and S.H. conceptualized the research. F.D.L. wrote the manuscript. S.H. revised and edited the manuscript. All of the authors read the paper. **Competing interests:** S.H. was a shareholder of Franka Emika GmbH, the manufacturer of the used robot manipulator. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the main manuscript or the Supplementary Materials. The datasets generated and analyzed in the current study are available at https://doi.org/10.5061/dryad.h44j0zpsf. Requests for additional materials should be addressed to S.H.

Submitted 19 February 2023
Accepted 14 November 2023
Published 13 December 2023
10.1126/scirobotics.adh0972

Science Robotics

Machine learning–driven self-discovery of the robot body morphology

Fernando Díaz Ledezma and Sami Haddadin

Sci. Robot. **8** (85), eadh0972. DOI: 10.1126/scirobotics.adh0972

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.