

# Learning to Infer Kinematic Hierarchies for Novel Object Instances

Hameed Abdul-Rashid<sup>2</sup> Miles Freeman<sup>1</sup> Ben Abbatematto<sup>1</sup> George Konidaris<sup>1</sup> Daniel Ritchie<sup>1</sup>

<sup>1</sup>Brown University <sup>2</sup>University of Illinois at Urbana-Champaign

## Abstract—

Manipulating an articulated object requires perceiving its *kinematic hierarchy*: its parts, how each can move, and how those motions are coupled. Previous work has explored perception for kinematics, but none infers a complete kinematic hierarchy on never-before-seen object instances, without relying on a schema or template. We present a novel perception system that achieves this goal. Our system infers the moving parts of an object and the kinematic couplings that relate them. To infer parts, it uses a point cloud instance segmentation neural network and to infer kinematic hierarchies, it uses a graph neural network to predict the existence, direction, and type of edges (i.e. joints) that relate the inferred parts. We train these networks using simulated scans of synthetic 3D models. We evaluate our system on simulated scans of 3D objects, and we demonstrate a proof-of-concept use of our system to drive real-world robotic manipulation.

## I. INTRODUCTION

People frequently interact with articulated objects: opening doors, putting things away in drawers, adjusting the height of chairs, etc. To a human, such interactions are effortless, but the underlying process is complex: one must perceive the *kinematic hierarchy* of the object (i.e. which parts can move, how they move, and how those motions are coupled). Then, one must plan a sequence of actions to transform the object from its current kinematic pose into another, and finally, manipulate the object to execute that plan. The perception part of this process is particularly challenging, as objects with similar function can vary considerably in both their structure and their geometry. Cabinets, for example, can have differing numbers of drawers in different arrangements and many different doorknob shapes. To succeed at manipulating such objects in the wild, an agent must infer the kinematic hierarchies of never-before-seen object instances with these kinds of variations.

This problem has received considerable attention in computer vision and robotics literatures. Vision researchers have made progress on detecting atomic parts of 3D objects and predicting the presence and parameters of individual kinematic joints [29], [35], [34]. These approaches are trained on large datasets and thus have the potential to generalize to novel objects from a known category, however, they do not infer kinematic hierarchies. Roboticians have leveraged embodiment to solve the problem by providing the agent with human demonstrations or allowing it to interact with its environment [28], [13], [15]. Recent works infer kinematic hierarchies but assume a simple template (e.g. one door per cabinet) and thus cannot generalize to the structural variability that exists in real world objects [1], [19], [11].




Fig. 1. We propose a system for robot perception of the kinematic structure of never-before-seen object instances. The robot scans the object from multiple viewpoints, building a 3D point cloud representation of it. Our system then segments this point cloud into parts and infers the kinematic hierarchy that couples them, allowing for the robot to plan and execute articulated motions of the object.

In this paper, we propose a perception system that both infers kinematic hierarchies and also generalizes to novel object geometries and structures. Our system infers the moving parts of an object and the kinematic couplings between them. To infer parts, it uses a point cloud instance segmentation neural network. To infer kinematic hierarchies, it uses a graph neural network which learns to infer the existence, direction, and type of edges (i.e. joints) that relate the inferred parts. This system is trained on simulated scans of a large set of synthetic 3D models, enabling it to learn to generalize to new, never-before-seen object instances which may differ geometrically and structurally from the training data. We evaluate our system’s perception abilities on synthetic point clouds, and find that our system reliably detects moving parts and reconstructs kinematic hierarchies on synthetic data. We also demonstrate a proof-of-concept application of using the output of our perception system to drive manipulation executed by a real-world robot platform.

In summary, our contributions are:

- The first method for inferring kinematic hierarchies from never-before-seen instances of 3D objects, without reliance on a schema or template.
- A novel graph neural network approach for predicting hierarchies from part segments.

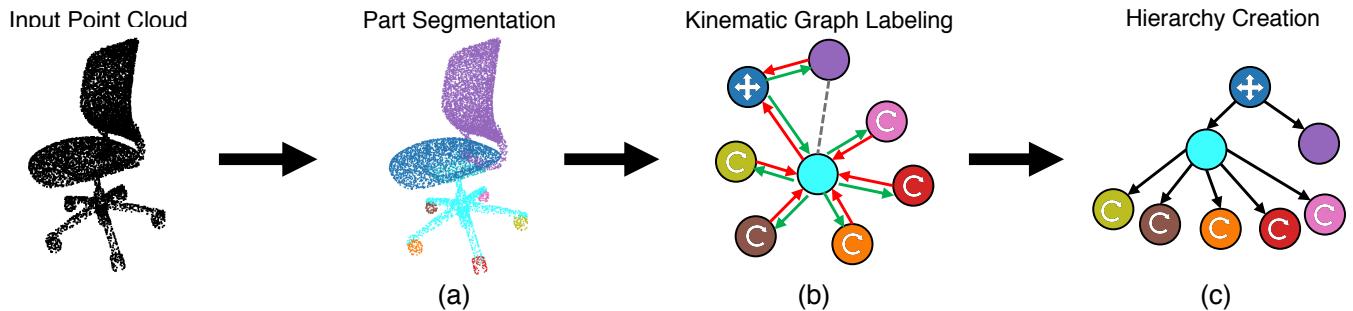


Fig. 2. Overview of our system. Given an input 3D point cloud representation of an object, our system (a) segments the point cloud into parts using an instance segmentation network, (b) constructs a part graph and uses a graph neural network to label its nodes and edges with kinematic properties, and (c) constructs a kinematic hierarchy from this graph which can be used for robot manipulation.

## II. RELATED WORK

### A. Approaches from the Robotics Community

Kinematic model and articulated pose estimation are well-studied problems in the robotics literature. Observation-based approaches [28], [24], [22], [6], [33] fit motion models to an observed sequence of part poses. Interactive approaches leverage the embodiment of the manipulator to generate part motion autonomously [13], [10], [2], [20], [16], [3] and to reason about the intended effects of actions [12]. More recently, learning-based approaches to category-level model and pose estimation [19], [1], [11], [36] have shown promise, but are thus-far limited to a-priori known connectivity graphs. Several methods estimate only articulated pose given known kinematic models and geometry [4], [7], [23].

### B. Approaches from the Computer Vision Community

Wang *et al.* train a model to jointly predict movable parts and their motion parameters from an input point cloud. [29]. Yi *et al.* take a pair of articulated 3D shapes of the same category but in different kinematic poses, and jointly infer moving parts and their kinematic motion parameters [35]. Yan *et al.* infer movable parts by hallucinating their motions over multiple time steps [34]. These approaches train deep neural networks which generalize to geometrically and structurally different input objects. However, none infer the structure of the object’s kinematic hierarchy: they treat each articulated part as it if moves in isolation. Our approach bears resemblance to RigNet, a system for inferring the skeletal joint structure of articulated humanoid and creature characters [32]. Like us, they also use a graph neural network for predicting graph topology information.

## III. APPROACH

In this section, we overview our method for kinematic mobility perception. Fig. 2 shows a schematic of our pipeline.

### A. Part Segmentation

The input to our system is an unlabeled point cloud depicting an articulated object; a robot can obtain such input by unprojecting and consolidating frames from an onboard depth sensor. Our system first segments this point cloud into parts, some of which are movable (Fig. 2a). This stage relies

upon an existing neural network architecture for point cloud instance segmentation (i.e. segmenting a point cloud without consistent part labels) [21].

### B. Kinematic Graph Labeling

Given a part-segmented point cloud, our system next infers the articulation properties of those parts: the type of articulated motion they support (if any) and how the motions of different parts are coupled. To do this, it converts the segmented point cloud into a graph and phrases the problem as one of *graph labeling*: nodes in the graph represent parts, and edges between them represent potential kinematic connections (e.g. joints) between parts. A graph neural network predicts whether each node is static, rotating, or translating, as well as whether each edge should exist (i.e. whether two parts should be kinematically connected), the probability of each node being the root of the tree, and the direction of kinematic dependency (Fig. 2b).

### C. Hierarchy Creation

Given the labeled graph, we construct a kinematic hierarchy by converting the bidirectional graph into a n-ary tree. Xu *et al.* address this problem by constructing a pairwise matrix where each entry represents the negative log probability of the corresponding parts possessing an edge connection [32]. The pairwise matrix is traversed from a predicted root node and a n-ary tree is extracted via a Minimum Spanning Tree (MST) algorithm [25]. We adopt this method to convert our predicted bidirectional graphs to an n-ary tree.(Fig. 2c).

## IV. PART SEGMENTATION

To segment point clouds into parts, we use a neural network based on that of PartNet [21]. This architecture uses a PointNet++ network [27] to compute per-point embeddings, each of which is then fed to a semantic segmentation branch (assigns semantic labels to each point) and an instance segmentation branch (assign each point to one of  $N$  possible instance masks; we use  $N = 24$ ). As our system does not assume parts are semantically labeled, we could remove the semantic segmentation branch. However, PartNet showed that semantic segmentation helps improve the quality of




Fig. 3. Segmented point clouds, their initial overcomplete part graphs, and their ground truth part graphs (as determined by a human labeler).

instance segmentation. Thus, we treat the semantic segmentation branch as a “motion type prediction” branch (static, rotating, translating, rotating and translating)—labels which we do assume of our data.

We train the segmentation network on 3D objects from the PartNet-Mobility dataset [30]. For augmentation, we use multiple poses for each moving part of each object, as described in more detail in Section VI. To simulate a robot’s onboard depth sensor, we render point clouds from these objects using a software simulator of the Kinect sensor [18]. We train the network using the Adam optimizer [14] with learning rate  $10^{-4}$ . Training was run for at most 8 hours per category.

## V. KINEMATIC GRAPH ANNOTATION

Here we describe how our system takes a segmented point cloud, converts it to a graph, and labels this graph with the information required to construct a kinematic hierarchy.

### A. Point Cloud to Graph Conversion

The kinematic hierarchy for a segmented object is a directed tree: tree nodes correspond to parts, nodes are labeled with a motion type (rotation and/or translation) and edge directions indicate kinematic couplings (i.e. edge  $A \rightarrow B$  means part  $B$  moves with part  $A$ ). To infer this tree from a segmented point cloud, our system first constructs an overcomplete, undirected, unlabeled graph over the parts, and then prunes edges, determines edge directions, and labels nodes using a graph neural network. Our system

constructs the initial graph by adding an undirected edge between any two parts whose Euclidean distance is below a small threshold. Fig. 3 shows some example segmented point clouds and their initial graphs.

### B. Graph Labeling Network

Given a graph constructed by the above procedure, we use a graph neural network to label its nodes and edges with the information required to construct a kinematic hierarchy.

**Input features:** Every node corresponds to one part, i.e. one subset of points from the object point cloud. We encode these point subsets into per-node vectors  $\mathbf{x}$  using a PointNet [26].

**Graph network:** To make per-node and per-edge predictions, our system must understand the context of each node or edge within the object. To satisfy this goal, we use a graph convolutional network to convert node features  $\mathbf{x}$  into context-aware features  $\mathbf{y}$  (Fig. 4). Our network passes all nodes through three GraphSAGE convolution layers [9], which compute new node features based on learned aggregations of neighbor node features. Each convolution layer is followed by a pooling and unpooling layer based on learned edge collapses [8]. In early experiments, we found this pooling + unpooling scheme to perform better than architectures that used only convolution. Finally, the outputs of all convolution layers and all pooling layers are concatenated to form a multi-scale feature vector  $\mathbf{y}$  for each node.

**Node and edge attribute prediction:** The per-node features  $\mathbf{y}$  are fed to multi-layer perceptrons (MLPs) for predicting node attributes. For edge attributes, the  $\mathbf{y}$ ’s for both edge endpoints are concatenated and fed to the MLP.

### C. Training

We train the graph labeling network on synthetic 3D models from PartNet-Mobility [30]. We first pre-train the node feature PointNet by training it as the encoder in an autoencoder framework; this network trains for 500 epochs using Nesterov SGD with a learning rate of  $10^{-3}$ . The graph labeling network trains for 30 epochs using Adam with a learning rate of  $10^{-3}$ . Training both networks takes 1.5 - 3 hours on a machine with a NVIDIA RTX 2080 Ti, 6-core Intel i7 CPU, and 32GB of RAM

## VI. RESULTS & EVALUATION

Here we evaluate our method’s ability to infer kinematic hierarchies via controlled experiments and a demonstration on a real-world robot. For experiments on synthetic PartNet-Mobility objects, we use the same train/test splits introduced by Mo *et al.* [21], restricted to models with valid kinematic trees. We augment the training and test datasets by uniformly sampling articulated poses within each part’s ground truth range of motion. We train on 249 storage furniture models, 22 lamp models, and 53 chair models; we test on 69, 8, and 12 models, respectively. We use 18 pose augmentations for training and testing.

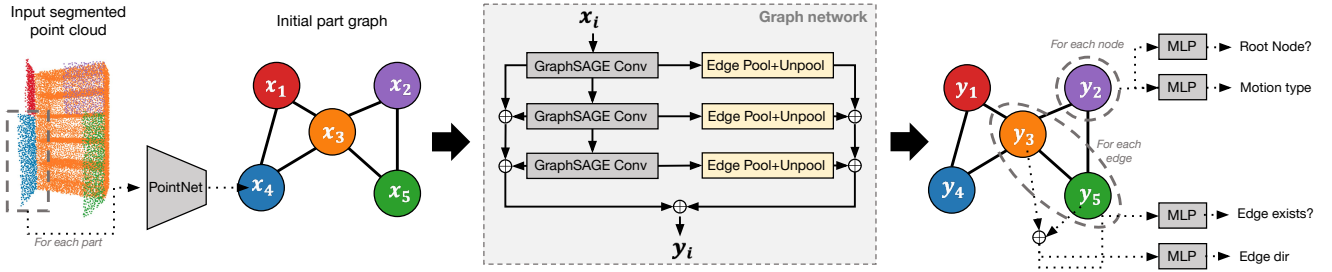


Fig. 4. Neural network architecture for kinematic graph labeling. A PointNet [26] converts each part point cloud into an initial node feature  $\mathbf{x}$ . These features go through graph convolutional + pooling layers to produce features  $\mathbf{y}$  which are passed to MLPs for predicting node and edge attributes.

TABLE I  
PART SEGMENTATION ON PARTNET-MOBILITY (MAP @ 0.5 IoU)

Category	Clean	Noisy
<i>Storage Furniture</i>	0.922	0.907
<i>Lamp</i>	0.695	0.593
<i>Chair</i>	0.824	0.811

### A. Part Segmentation

Table I summarizes the performance of our part segmentation method on synthetic 3D objects. We use an established instance segmentation metric [21]: we compute the intersection over union (IoU) of each predicted part segment with its closest ground truth part segment, classify it as a “correct” prediction if the IoU is over 0.5, and then compute the mean average precision of these classifications across all parts in all test set objects. We compute this metric under two conditions: (1) *Clean* point clouds sampled directly from the surface of PartNet-Mobility 3D meshes, (2) *Noisy* point clouds generated by simulated Kinect scans of PartNet-Mobility models as described in Section IV. Fig. 5 shows some qualitative examples. The network reliably segments parts in clean point clouds, and this performance degrades slightly in the noisy depth scans. The lamp category is especially challenging due to the small number of models present in the training data.

### B. Kinematic Structure Prediction

We next evaluate our graph neural network module’s ability to infer the structure of an object’s kinematic hierarchy. To assess this performance, we use four metrics:

- $\mathbf{E}_{\text{type}}$ : % of nodes whose motion type is mis-predicted.
- $\mathbf{E}_{\text{exist}}$ : % of edges whose existence is mis-predicted.
- $\mathbf{E}_{\text{dir}}$ : % of edges whose direction is mis-predicted.
- $\mathbf{E}_{\text{root}}$ : % of models whose root node is mis-predicted.
- **Tree F1**: the F1 score of the predicted kinematic hierarchy with respect to the ground truth one. Precision is computed via top-down traversal of the predicted tree, counting the fraction of nodes and edges which match their counterparts in the ground truth tree. Recall computation uses the same procedure, with the roles of the predicted and ground-truth trees reversed.

Table II summarizes the performance of our network on Clean and Noisy synthetic data. As expected, the GNN eval-

uated on clean data with ground truth instance segmentation preforms well on all classes. Even in the case where  $\mathbf{E}_{\text{dir}}$  is relatively high for the lamp and chair class the negative log probability pairwise matrix complemented with MST is able to construct trees with a negligible amount of error. On noisy data, performance again degrades slightly. Fig. 6 shows some qualitative examples.

TABLE II  
KINEMATIC STRUCTURE ERROR FOR DIFFERENT DATA TYPES

Category	Data Type	$\mathbf{E}_{\text{type}} \downarrow$	$\mathbf{E}_{\text{exist}} \downarrow$	$\mathbf{E}_{\text{dir}} \downarrow$	$\mathbf{E}_{\text{root}} \downarrow$	<b>Tree F1</b> $\uparrow$
<i>Storage Furniture</i>	Clean	1.16	1.2	2.22	0	99.73
	Noisy	0.39	2.03	29.04	0	99.02
<i>Lamp</i>	Clean	0.47	4.63	31.11	0.23	98.24
	Noisy	0	0.98	17.24	0.92	97.82
<i>Chair</i>	Clean	0.08	0.1	39.64	0	99.77
	Noisy	0.24	0.99	41.73	0	98.67

### C. Robot Demo

Finally, we demonstrate that the model enables a Kinova Movo robot to manipulate novel object instances. A storage furniture object was densely scanned using the robot’s Kinect sensor by driving the robot around the object. Point-clouds were merged using RTAB-Map [17], which performed SLAM using the robot’s odometry, base laser scans, and visual data. The object was manually segmented from the scene and assigned a frame of reference. Our model successfully performed part segmentation and kinematic hierarchy inference; predictions are visualized in Figure 1. Motion model predictions, i.e. the axes of rotation, were generated using the motion type predictions from the estimated kinematic tree and a heuristic: the axes of rotation were determined from the intersection of part bounding boxes and oriented in the world vertical direction.

Using the predicted tree and motion models, the agent is able to plan and execute interactions with the object. We fixed a grasp on the relevant object part, and ran a simple controller which moves the end-effector along the predicted direction of motion toward a goal configuration. Fig. 7 shows frames from a video of the robot scanning the cabinet and closing one of its doors.




Fig. 5. Examples of part segmentation for both Clean and Noisy point clouds.




Fig. 6. Examples of predicted kinematic hierarchies for both Clean and Noisy point clouds. All examples shown use ground truth part segmentations.

## VII. FUTURE WORK

Training our neural networks requires collections of part-segmented, kinematically-annotated 3D models. Such data is not widely available, as it requires nontrivial human annotation effort. Recent computer vision research has demonstrated the possibility of segmenting 3D models with limited or even no human supervision [5], [37]. The development of similar techniques for identifying articulated motion within 3D shapes would allow our methods to be applied to any object type for which unsegmented 3D models are available [31].

Our approach treats segmentation and graph labeling as separate sub-problems, but one can argue that they are

coupled: parts determines what motions are possible, and whether a motion is possible determines whether a proposed part is a good one. Accordingly, the Shape2Motion system found benefits in jointly segmenting and predict motion parameters for parts [29]. It is possible that jointly segmenting and predicting kinematic hierarchy could confer similar benefits.

We focused on inferring the structure of an object’s kinematic hierarchy. Our system could be combined with one that focuses on kinematic motion parameter prediction [29], [34], [35] to produce a complete mobility perception system.




Fig. 7. Frames from a video demonstrating a robot scanning and then interacting with an articulated object.

### VIII. CONCLUSION

We presented a perception system that infers kinematic hierarchies for never-before-seen object instances. Our system infers the moving parts of an object and the kinematic couplings between them. To infer parts, it uses a point cloud segmentation neural network. To infer kinematic couplings, it uses a graph neural network to predict the existence, direction, and type of edges (i.e. joints) that relate the inferred parts. We train these networks using simulated scans of synthetic 3D models. In experiments, our system inferred accurate kinematic hierarchies for simulated scans of 3D objects and scans of real-world objects gathered by a mobile robot.

### ACKNOWLEDGEMENTS

This research was supported by NSF CAREER Award 1844960 to Konidaris, and by the ONR under the PERISCOPE MURI Contract N00014-17-1-2699. Disclosure: George Konidaris is the Chief Robotist of Realtime Robotics, a robotics company that produces a specialized motion planning processor.

### REFERENCES

- [1] Ben Abbatematteo, Stefanie Tellex, and George Konidaris. Learning to generalize kinematic models to novel objects. In *Proceedings of the Third Conference on Robot Learning*, 2019.
- [2] Patrick R Barragán, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Interactive bayesian identification of kinematic mechanisms. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2013–2020. IEEE, 2014.
- [3] Manuel Baum, Matthew Bernstein, Roberto Martin-Martin, Sebastian Höfer, Johannes Kulick, Marc Toussaint, Alex Kacelnik, and Oliver Brock. Opening a lockbox through physical exploration. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 461–467, Nov 2017.
- [4] Jonathan Brookshire and Seth Teller. Articulated pose estimation using tangent space approximations. *The International Journal of Robotics Research*, 35(1-3):5–29, 2016.
- [5] Zhiqin Chen, Kangxue Yin, Matthew Fisher, Siddhartha Chaudhuri, and Hao Zhang. Bae-net: Branched autoencoder for shape co-segmentation. *Proceedings of International Conference on Computer Vision (ICCV)*, 2019.

- [6] Andrea F. Daniele, Thomas M. Howard, and Matthew R. Walter. Learning articulated object models from language and vision. 2017.
- [7] Karthik Desingh, Shiyang Lu, Anthony Opipari, and Odest Chadwicke Jenkins. Efficient nonparametric belief propagation for pose estimation and manipulation of articulated objects. *Science Robotics*, 4(30), 2019.
- [8] Frederik Diehl. Edge contraction pooling for graph neural networks. *arXiv preprint arXiv:1905.10990*, 2019.
- [9] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017.
- [10] Karol Hausman, Scott Niekum, Sarah Osentoski, and Gaurav S Sukhatme. Active articulation model estimation through interactive perception. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 3305–3312. IEEE, 2015.
- [11] Ajinkya Jain, Rudolf Lioutikov, and Scott Niekum. Screwnet: Category-independent articulation model estimation from depth images using screw theory. *arXiv preprint arXiv:2008.10518*, 2020.
- [12] Ajinkya Jain and Scott Niekum. Learning hybrid object kinematics for efficient hierarchical planning under uncertainty. In *IROS*, 2019.
- [13] Dov Katz. *Interactive perception of articulated objects for autonomous manipulation*. University of Massachusetts Amherst, 2011.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [15] Oliver Kroemer, Scott Niekum, and George Dimitri Konidaris. A review of robot learning for manipulation: Challenges, representations, and algorithms. *Journal of machine learning research*, 22(30), 2021.
- [16] Johannes Kulick, Stefan Otte, and Marc Toussaint. Active exploration of joint dependency structures. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 2598–2604. IEEE, 2015.
- [17] Mathieu Labbé and François Michaud. Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation. *Journal of Field Robotics*, 36(2):416–446, 2019.
- [18] Michael Landau. Kinect Infrared (IR) and Depth Image Simulator. MATLAB Central, 2017.
- [19] Xiaolong Li, He Wang, Li Yi, Leonidas Guibas, A Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *CVPR*, 2020.
- [20] R. Martín Martín and O. Brock. Online interactive perception of articulated objects with multi-level recursive estimation based on task-specific priors. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2494–2501, Sept 2014.
- [21] Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [22] Scott Niekum, Sarah Osentoski, Christopher G Atkeson, and Andrew G Barto. Online bayesian changepoint detection for articulated motion models. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 1468–1475. IEEE, 2015.
- [23] Jana Pavlasek, Stanley Lewis, Karthik Desingh, and Odest Chadwicke Jenkins. Parts-based articulated object localization in clutter using belief propagation. *arXiv preprint arXiv:2008.02881*, 2020.
- [24] Sudeep Pillai, Matthew R Walter, and Seth Teller. Learning articulated motions from visual demonstration. In *Robotics: Science and Systems*, 2015.
- [25] Robert Clay Prim. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401, 1957.
- [26] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- [27] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Neural Information Processing Systems*, 2017.
- [28] Jürgen Sturm, Cyrill Stachniss, and Wolfram Burgard. A probabilistic framework for learning kinematic models of articulated objects. *Journal of Artificial Intelligence Research*, 41:477–526, 2011.
- [29] Xiaogang Wang, Bin Zhou, Yahao Shi, Xiaowu Chen, Qinqing Zhao, and Kai Xu. Shape2motion: Joint analysis of motion parts and

- attributes from 3d shapes. *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [30] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [31] Xianghao Xu, David Charatan, Sonia Raychaudhuri, Hanxiao Jiang, Mae Heitmann, Vladimir Kim, Siddhartha Chaudhuri, Manolis Savva, Angel X Chang, and Daniel Ritchie. Motion annotation programs: A scalable approach to annotating kinematic articulations in large 3d shape collections. In *2020 International Conference on 3D Vision (3DV)*, pages 613–622. IEEE, 2020.
- [32] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. Rignet: Neural rigging for articulated characters. *ACM Trans. on Graphics*, 39, 2020.
- [33] Jingyu Yan and Marc Pollefeys. Automatic kinematic chain building from feature trajectories of articulated objects. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 1, pages 712–719. IEEE, 2006.
- [34] Zihao Yan, Ruizhen Hu, Xingguang Yan, Luanmin Chen, Oliver van Kaick, Hao Zhang, and Hui Huang. Rpm-net: Recurrent prediction of motion and parts from point cloud. *ACM Transactions on Graphics (Proceedings of SIGGRAPH ASIA 2019)*, 38(6):240:1–240:15, 2019.
- [35] Li Yi, Haibin Huang, Difan Liu, Evangelos Kalogerakis, Hao Su, and Leonidas Guibas. Deep part induction from articulated object pairs. *Transactions on Graphics*, 37(6), December 2018.
- [36] Vicky Zeng, Timothy E Lee, Jacky Liang, and Oliver Kroemer. Visual Identification of Articulated Object Parts. *arXiv*, 2020.
- [37] Chenyang Zhu, Kai Xu, Siddhartha Chaudhuri, Li Yi, Leonidas J. Guibas, and Hao Zhang. AdaCoSeg: Adaptive shape co-segmentation with group consistency loss. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2020.