Experimental Analysis of Human Control Strategies in Contact Manipulation Tasks

Ellen Klingbeil, Samir Menon, and Oussama Khatib

Stanford Robotics Lab, Department of Computer Science
Stanford University, Stanford CA 94305, USA,
ellenrk7@stanford.edu, {smenon,ok}@cs.stanford.edu

Abstract. Here, we present insights into human contact-control strategies by defining conditions to determine whether a human controls a contact state, empirically analyzing object-to-environment contact geometry data obtained from human demonstrations in a haptic simulation environment, and testing hypotheses about underlying human contact-control strategies. Using haptic demonstration data from eleven subjects who inserted non-convex objects into occluded holes, we tested the following human contact-control hypotheses: (*h1*) humans follow a task trajectory that tracks pre-planned contact-state waypoints organized in a *contact-state graph* (*contact-waypoint hypothesis*); (*h2*) humans traverse the contact-state graph, explicitly controlling some contact states or subsets of contact states, in addition to the pre-determined initial and final goal states (*controlled subgraph hypothesis*); (*h3*) humans use a control policy where the only controlled states are the starting state for the task and the goal state (*state policy hypothesis*). Notably, we found that humans tend to visit a select few contact states once they enter each state's vicinity in the graph, which is evidence against *h3*. Yet humans do not always visit said states (visit probability $< 40\%$), which is, in addition, evidence against *h1* provided different humans adopt similar strategies. We show that a classifier to determine when humans control their trajectories to visit specific contact states, when parameterized correctly, is invariant to graph aggregation operations across the false-positive to false-negative tradeoff spectrum. This indicates our results are robust given the data we obtained and suggests that efforts to characterize human motion should focus on *h2*.

Keywords: Haptics, contact-control, human skills analysis

1 Introduction and Related Work

Efforts to transfer human skills to robots are presently limited by the lack of models that characterize human contact-control strategies. Such models help identify control primitives and allow using haptic interfaces for human-robot teaching by demonstration. Unlike free-space rigid-body motions, which are readily described by spatial trajectories or continuous dynamical systems [1] controlled with potential fields [2] or proxies [3], motions in contact also require describing

discontinous contact-state dependent kinematic and force constraints [4–9]. Contact states (and their constraints) can be described as kinematic configurations where one or more vertices, edges, or planes of an object contacts another, and the connectivity between contact states can be be represented by a graph [10]. Graph traversals performed can thus describe a specific trajectory of states encountered while controlling an object in contact [10]: contact-control trajectories constitute a motion plan. However, graph size grows combinatorially with contact states [11, 12], and accumulated errors across graph traversals render such motion plans unachievable in real-world scenarios. Remarkably, humans overcome these problems and efficiently perform motor tasks with complex contact states.

Research in emulating human contact-control strategies with robots has involved identifying a walk in the contact-state graph betweeen two arbitrary states, along with engineering hybrid force-position controllers to transition between each of the states [13–15]. Such *"contact-state planning"* strategies successfully achieved contact-control for tasks with few contact states, for instance, placing a box in a corner [16] or inserting pegs in holes in two-dimensions [17]. This approach, however, does not readily accommodate efforts to map human demonstrations to robots because physical noise complicates identifying contact states in real world demonstrations [18], and humans rarely traverse the same contact graph trajectory even in noiseless virtual worlds. Moreover, even if humans repeatably seem to traverse the same trajectory, no method exists to characterize what subset of contact states in the graph walk that humans explicitly control. Contradicting the central assumption of contact-state planning, humans visit a large number of contact states as task complexity increases but visit only a small subset of the states reliably [19]. Moreover, the vast majority of contact states are visited for time periods of less than 40ms, negating the possibility of human feedback control due to physiological delay limits [20]. These results led us to the hypothesis that humans may only control a subset of contact states, using controllers that are invariant to other states.

Here, we present a framework to gain insight into human contact-control strategies by analyzing data from contact-state graph traversals for humans performing two multi-step object insertion tasks (L- and S-shaped object insertion) in a six-dimensional physics-based haptic simulation (Fig. 1; see [19] for details). Using this framework, we tested three hypotheses for human contact-control strategies: (*h1*) humans follow a task trajectory that tracks pre-planned (geometric) contact-state waypoints organized in a *contact-state graph*; (*h2*) humans traverse the contact-state graph, explicitly controlling some states or subsets of states, in addition to the pre-determined initial and final goal states; (*h3*) humans use a control policy where the only controlled states are the starting state for the task and the goal state (for instance, a state-action policy produced by a reinforcement learning algorithm where the state does not include the contact state). We discuss results that constitute evidence against hypotheses *h1* and *h3* and demonstrate that our analysis methods are invariant to varying levels of spatial smoothing introduced by aggregating contact-state graph nodes.

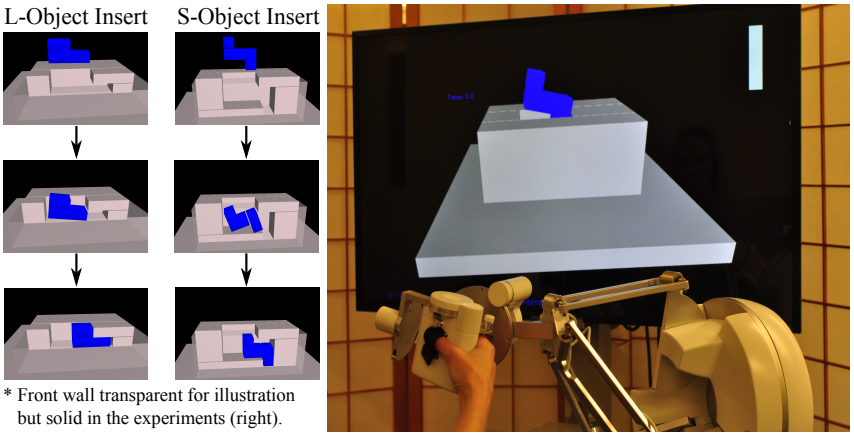

Fig. 1. Experimental conditions. Subjects inserted an L-object and S-object into corresponding holes using the Force Dimension Sigma-7 haptic device (right). Each motion consisted of a free-space *'Hold'*, a *'Move'* to the goal, and a final *'Maintain'* at the goal.

2 Technical Approach

To develop our framework for studying human contact-control strategies, we will begin by defining the contact-state graph. We then introduce a novel definition for a controlled graph state, which we will use in our analysis of human contact-control strategies.

2.1 Contact-State Graph

A contact state s is a configuration of an object with respect to its environment defined as $s \in (\{\mathbb{V}, \mathbb{E}, \mathbb{P}\}_{object} \times \{\mathbb{V}, \mathbb{E}, \mathbb{P}\}_{environment})$ where $\mathbb{V}$, $\mathbb{E}$, and $\mathbb{P}$ represent a vertex, edge, or plane [21].[1] A contact-state graph $\mathbb{S}$ is the set of contact states along with a transition function that takes as an input an arbitrary starting state and ending state and returns as an output either 0 (disconnected), or 1 (forward-connected). Bi-directional connections, in addition, return a 1 when the states are passed in inverse order.[2]

Empirical Contact Subgraph Since enumerating the states in a full contact-state graph is computationally prohibitive for anything beyond the simplest of objects [12], here, we decided to conduct our analyses on a reduced contact-state graph $\hat{\mathbb{S}}$ whose set of contact states and transitions are empirically obtained from

[1] In this paper, we ignore contact states consisting of $(\{\mathbb{V}\}_{object} \times \{\mathbb{V}\}_{environment})$ and $(\{\mathbb{E}\}_{object} \times \{\mathbb{E}\}_{environment})$ where the edges are parallel, since they are degenerate cases.

[2] In [10], the contact graph transition function was specifically defined to return a 1 if a given state can be reached by another without losing contact.

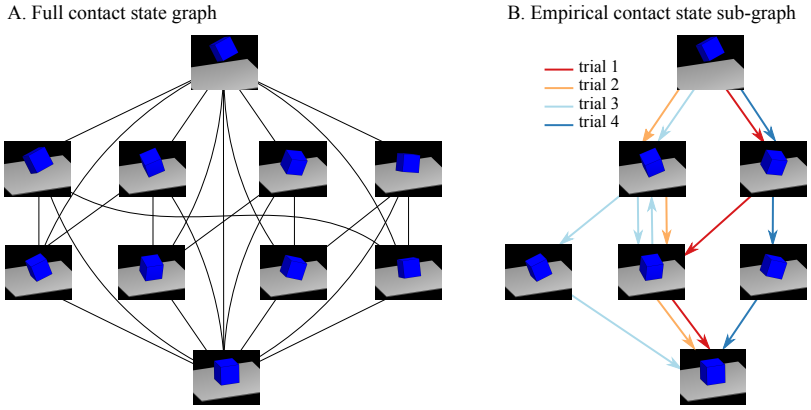

Fig. 2. Empirical contact subgraph. A. Full contact-state graph for the task of placing the bottom surface of a box on a plane. **B.** Empirical contact-state subgraph generated by observing contact-state visits from four human demonstrations of the placing task. The figures demonstrate how empirical graphs are smaller.

human demonstrations of the given task. For example, if the contact-state sequence $[s_i, s_j]$ is visited in a given human trial, the states s_i and s_j are added to the graph if $s_i \notin \hat{\mathbb{S}}$ or $s_j \notin \hat{\mathbb{S}}$, along with adding the transition $t_i j$. In contrast with the difficulty in even enumerating states in the full contact-state graph (according to the method of [10]), our empirical contact-state subgraphs generated from human demonstration data are compact and tractable to analyze (Fig. 2). This makes them an attractive tool to study human contact-control.

Graph State Aggregation Despite moving to empirical subgraphs, the number of states obtained is usually still large. Moreover, humans do not traverse states uniformly and may tend to visit a few states more often than others. We note that nodes with few visits can be a consequence of stochastic human-simulation interactions. Since any analysis of human contact-control should not be affected by this stochasticity, any results should be invariant to perturbations in state. To achieve this, we decided to delocalize the notion of a node by clustering it with neighboring nodes. We did this by introducing an aggregation operator for contact-state graphs and explored its effect on the graph topology. Our operator, the K-level neighbor aggregator, groups together nodes which are within K edges from the cluster central node (Fig. 3). The aggregation operation can also be viewed as creating clusters around every state. As such, graph aggregation leads to a smoother distribution of visits across state clusters (compared to states alone). Note that this method of clustering results in every state belonging to more than a single cluster.

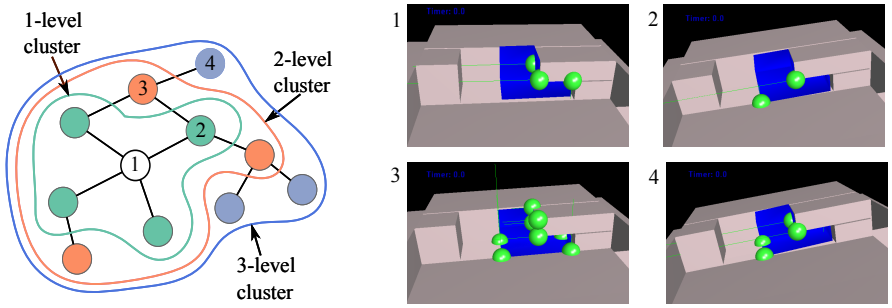

Fig. 3. Contact-graph state aggregation. *K-level neighbor aggregation* groups nodes within K edges from a central node. The example is a subgraph from the actual empirical contact-state graph from human demonstrations for the L-object insertion task.

2.2 Controlled Graph State

We assume that humans can delineate different contact states and, in the process of doing so, can obtain a (potentially noisy) estimate of an error between the two. In particular, we assume that given two contact states, humans can compute a distance metric between them. If a human can not distinguish between the two states, the metric is zero. In contrast, if they can distinguish between the two states, the metric is greater than zero. Since we can not presently identify this specific metric, we will now develop a definition for a controlled state that is contingent purely upon the metric's existence.

We define a *controlled state* as a state $s_i \in \hat{\mathbb{S}}$ that has the following properties: 1) a human's estimate of the error between s_i and $s_{\neq i}$ has an upper bound; 2) a human must be able to correct said error and reduce it below a desired distance threshold ϵ; 3) Given a controlled state s_i and surrounding nodes within a distance ϵ ($\epsilon_{max} > \epsilon > 0$), there exists a graph distance metric δ ($\delta > \epsilon$) such that the likelihood of a trajectory passing through ϵ given that it passed through δ is greater than $\frac{1}{number of humans}$. The ratio $\frac{\delta}{\epsilon}$ serves as a confidence measure for whether the state is controlled.

2.3 Human Contact-Control Strategy

A contact manipulation event requires a control strategy to realize a sequence of contact states between two objects, starting in some contact state (usually free space), passing through many intermediate states, and ending in some final state (usually a desired stable configuration). A human contact-control strategy maps the current state of a manipulated object, $q(t), \dot{q}(t)$, the final contact state s_f, and measured interaction forces $f(t)$, to an applied control force $f_c(t)$. Applying a control strategy results in a traversal of the contact-state graph. A human control strategy is thus the mapping, $(s_f, f(t), q(t), \dot{q}(t)) \rightarrow f_c(t + \delta\, t)$. The present contact state is a deterministic function of the object configuration, $s(t) = state(q(t))$. However, this function may not be known to a human.

Here, we focus on the role of contact-state control in the model for human contact strategies, with the ultimate goal being to completely identify the human's control strategy.

We make the following assumptions in our subsequent analysis of human contact-control strategies: 1) After an initial training period, the change in human contact-control strategies is small; 2) Different humans have similar contact-control strategies when at contact states that are visited more times than the third quartile; 3) It is feasible to study human contact-control strategies with physics simulations that provide realistic visual and haptic feedback; 4) The control force response time, δt, is empirically assigned to the median human neural signal transmission plus response return times [20].

Having developed semantics that allow us to describe contact-control, we proceeded to test the following three hypotheses for human contact-control:

(*h1*) **Contact-waypoint hypothesis**. The human contact-control strategy involves tracking pre-planned contact-state waypoints, resulting in graph traversals that tend to be deterministic. As a consequence, this hypothesis would predict that each human has a set of strategies to draw upon while traversing the graph, and that the number of possible unique traversals would grow linearly with the number of subjects. Moreover, this hypothesis would also predict that each contact-waypoint (a controlled state) would be visited roughly the same number of times by humans while they traverse the contact graph. As such, we can reject this hypothesis if the traversal counts of the majority of controlled states differ by at least $\mathcal{O}(n_{subjects})$ for a small multiplicative constant. Alternate evidence would be the lack of repeatability in human contact trajectories across demonstrations.

(*h2*) **Controlled subgraph hypothesis**. Humans explicitly control some subsets of states in the graph encountered during a traversal, in addition to the pre-determined initial and final goal states. If this hypothesis is true, we expect clusters of controlled graph states to emerge in the analysis.

(*h3*) **State policy hypothesis**. Humans use a control policy where the only controlled states are the starting state for the task and the goal state. This hypothesis can be rejected if there exists another controlled state besides the initial and final.

Note that testing our hypotheses requires empirical data from a task of sufficient complexity, which led us to analyze data from the L- and S-object tasks instead of the box-on-plane task. The complexity of an experiment to reject the contact-waypoint hypothesis, for instance, should ideally limit the ability of humans to emulate other humans' contact-state trajectories even if they have the identical contact-control strategy. That is, whether or not they have an identical controller, it should be likely that they will end up exploring different parts of the contact-state graph.

2.4 Analyzing Controlled States

Performing the hypotheses tests requires a method to analyze controlled states in the human demonstration data. If a state s_i is controlled, it is likely to be

visited frequently across graph traversals from the human trials. Additionally, if a region in a graph is entered, it is likely that the graph trajectory will also pass within a small distance ϵ of the state if it is controlled. In this paper, ϵ was chosen to be zero, as identifying appropriate distance metrics between contact states is non-trivial and will be explored in future work. To discover possibly controlled states in the human graph traversals, a cluster of states C_i is formed around each state s_i (which we will refer to as the central node of the cluster) using the K-level neighbor aggregation operator. For each state s_i, we count the number of times n_{C_i} where a human graph traversal entered cluster C_i (i.e. a state $s_j \in C_i$ is visited after previously visiting a state $s_k \notin C_i$). If s_i is also visited when C_i is entered, the number of visits n_{s_i} to the state s_i is incremented by 1. The ratio n_{s_i}/n_{C_i} indicates the likelihood of passing through the state s_i whenever the surrounding region of the graph was entered across all the human graph traversals.

3 Experiments

To test our hypotheses for human contact-control strategies, we developed a haptic simulation environment to observe humans performing two insertion-type contact tasks [19]. Our haptic simulation environment uses a right-handed *Force Dimension Sigma 7* haptic interface [22], a dynamic proxy controller, and the *Open Dynamics Engine* for physics [23]. While creating stable simulations is challenging [24–29], we achieved stable contact manipulation by requiring geometric shapes to be composed of sets of cuboids and resolved instabilities by pruning contact state transitions that happened at a frequency associated with solver instabilities. We avoided using massless proxy methods since ignoring inertial effects can introduce a bias in studies of human motion. The tasks we studied were selected to have sufficient complexity to introduce a large number of contact states (see Fig. 1). The experiments consisted of having 11 subjects [3] insert L- and S-shaped objects into corresponding holes.

Subjects were pre-trained over 20 trials to apply control forces and moments within the device's limits, which were indicated with a slider bar on the screen (removed later for data-trials). During the training period, the time taken was displayed, and the longest time taken from the last ten training trials was set as the upper bound. Subjects then performed 50 data trials for each task. The environment pose was randomized to a new configuration for each trial. A success or failure message was displayed in text if the peak force or maximum time was exceeded. Subjects were not provided with any additional instructions on how to perform the task except to attempt to stay below the maximum force limits and perform the task in the allotted time.

[3] **Human Subjects:** Healthy, right-handed subjects with no history of motor disorders: 20m:6'0", 28m:5'9", 31f:5'4", 20m:6'0", 19m:6'0", 20m:5'7", 29m:5'11", 21f:5'2", 32m:5'11", 30m:5'8", 29m:5'8". Informed consent obtained in advance on a protocol approved by Stanford University's Institutional Review Board (IRB).

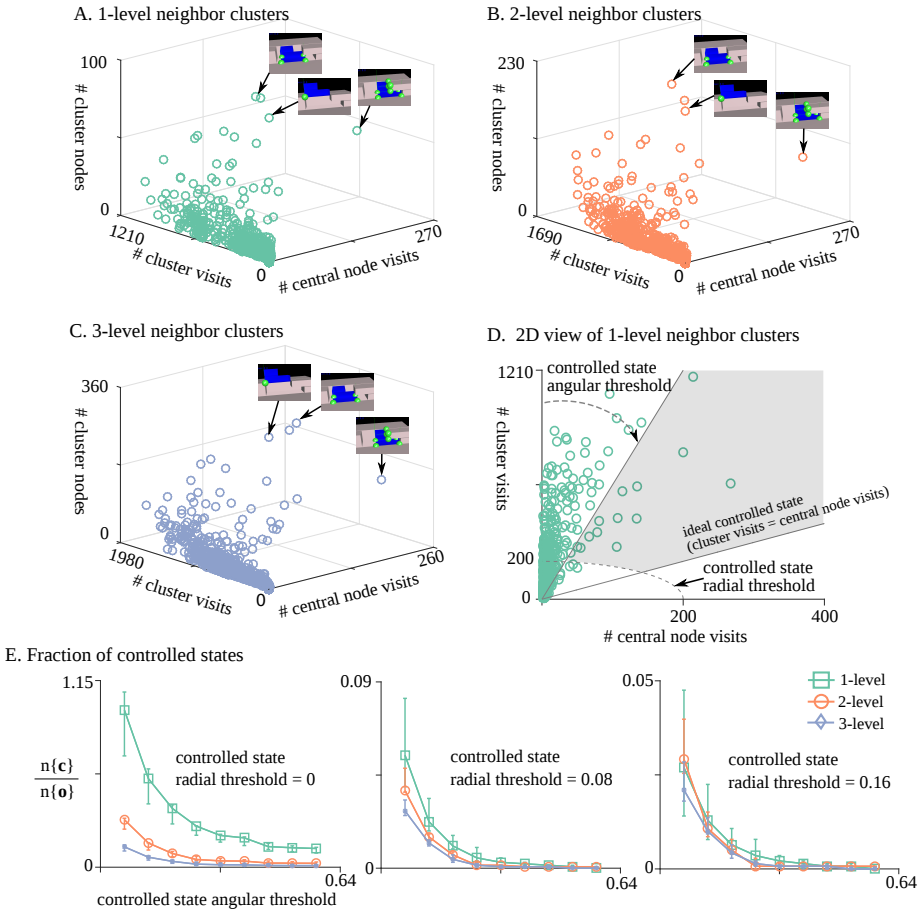

Fig. 4. Controlled states in L-Object trajectories. A, B, C. Clustering neighboring nodes revealed that only a few contact states were visited regularly whenever a human trajectory entered the state's cluster, indicating they were *controlled* ($\{c\}$). **D.** Possibly controlled states can be observed by varying the controlled state angular threshold and the controlled state radial threshold. The ideal controlled state is one where the state is visited every time its surrounding cluster is entered. **E.** There exists a value for the radial threshold where the fraction of states visited regularly was invariant to the level of clustering across varying values of angular controlled threshold used to classify a state as controlled or other ($\{o\}$).

4 Results

Figure 4(A-C) shows, for each s_i, the number of states in the cluster, the number of visits to the cluster, and the number of visits to state s_i (the central node of the cluster) for the L-object insertion task accumulated over all of the human subject trials. Observe that for large values of central node visits and number of cluster nodes the density of states is low, and thus the number of controlled

states is likely low relative to the total number of visited states. Additionally, the relative locations of the likely controlled states (Fig. 4(A-C)) varied little across the three cluster groupings. Fig. 5(A-D) shows similar results on a different task (the S-Object insertion).

Interestingly, there is a similar distribution of states as the cluster size is varied. This can be seen more clearly by varying two parameters – the controlled state angular threshold (i.e. the ratio of central node visits to cluster visits) and the controlled state radial threshold (Fig. 4(D)) – and observing how the ratio of controlled states ($\{c\}$) to other states ($\{o\}$) changes. The radial parameter removes states that are likely visited due to noise or random exploration. States closer to the ideal controlled line indicate a state that is controlled across all human subjects. The grey region in (Fig. 4(D)) indicates the states that would be considered as controlled for the given value of the controlled state angular and radial thresholds.

Figure 4(E) shows the ratio of controlled states to other as the controlled state angular and radial thresholds are varied for the different levels of clustering. The angular threshold is normalized with respect to the ideal controlled state (i.e. a state which is visited every time its surrounding cluster is entered), and the radial threshold is normalized with respect to the maximum radius of any state across the three levels of clustering. The error bars represent $95th$ percentile confidence intervals obtained by bootstrapping the ratio of controlled states to other across the given values of controlled state angular and radial thresholds [30]. Note that there exists a value for the controlled state radial threshold where the data points for each of the $1, 2, 3-level$ clustering on Fig. 4(E) closely coincide across varying values of the controlled states angular threshold. Thus, the fraction of states visited regularly (possibly controlled states) was invariant to the level of clustering. Similar results are observed on the S-object insertion task (Fig. 5(E)).

5 Experimental Insights

In summary, our goal was to construct a framework to help map human contact-control demonstrations to a notion of the underlying control strategy. To do so, we developed a definition for a controlled state and a classification method to help categorize whether specific contact states encountered by humans were controlled or not. An open problem in contact control is how to overcome the combinatorial growth in complexity associated with manipulation in contact for non-convex objects. We earlier found that it is unlikely that humans develop control strategies that explicitly rely on contact trajectory planning [19]. In addition, we found that it is also unlikely that humans simply perform random walks using a state-policy. Specifically, our results posit that humans tend to reliably visit a few specific contact states if they enter the state's vicinity in the graph, which is evidence against *h3*, yet do not always visit said states (visit probability $< 40\%$), which is evidence against *h1* if different humans adopt similar strategies. The remaining possibility within our set of contact-control hypotheses, that humans control to a select subset of states while performing a directed walk along the

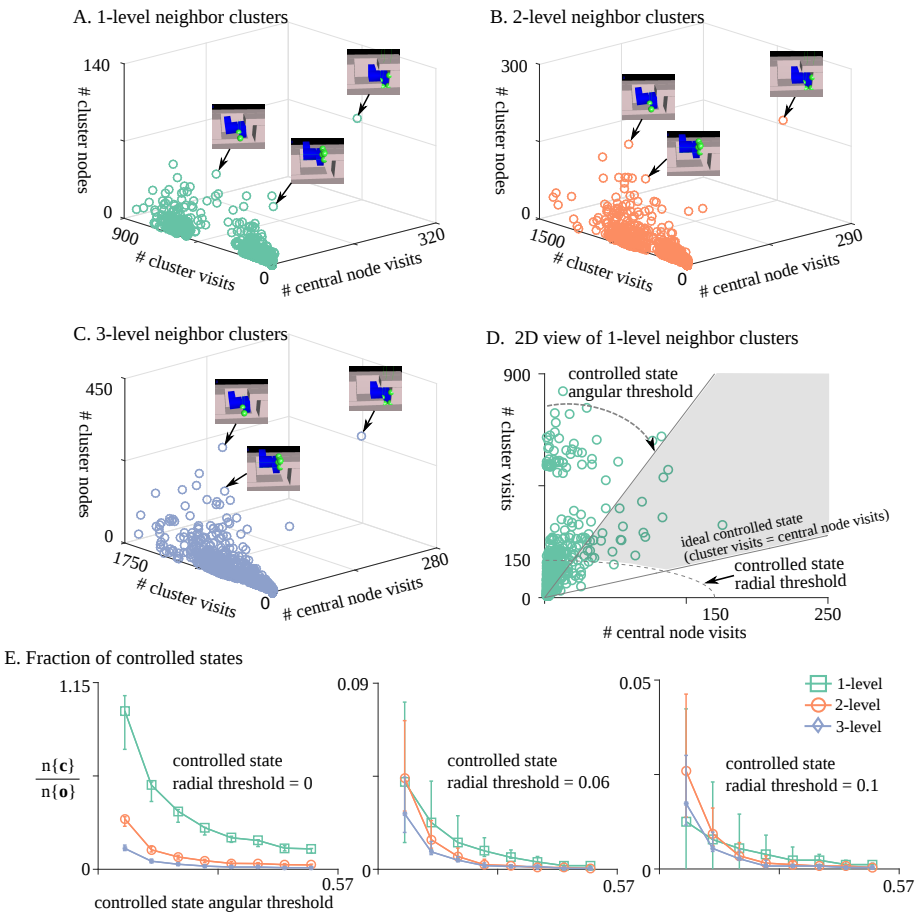

Fig. 5. Controlled states in S-Object trajectories. A, B, C. Similar to the L-Object, clustering neighboring nodes revealed that only a few contact states were visited regularly whenever a human trajectory entered the state's cluster, indicating they were *controlled* ($\{c\}$). **D.** Possibly controlled states can be observed by varying the controlled state angular threshold and the controlled state radial threshold. The ideal controlled state is one where the state is visited every time its surrounding cluster is entered. **E.** There exists a value for the radial threshold where the fraction of states visited regularly was invariant to the level of clustering across varying values of angular controlled threshold used to classify a state as controlled or other.

contact graph, is consistent with our results and should be considered a standing hypothesis for human contact control. Moreover, our haptic demonstration framework and controlled-state classification method can help empirically determine the controlled states while spanning the spectrum of false-positives and false-negatives. These results are invariant to graph aggregation, which serves as a control for stochastic effects.

Our results serve as a first step towards systematically studying human contact-control using empirically testable hypotheses. Future research directions include exploring inter-human variations in contact control, testing for invariance to other metrics besides graph aggregation, analyzing the effect of spatial distance across contact states, and modeling the error distribution while trying to control to a specific contact state.

Acknowledgements

We thank Keegan Go for his assistance with developing the haptic simulation environment. The project was supported by National Science Foundation National Robotics Initiative grant (IIS-1427396, O. Khatib and R. Bajcsy) and a grant from the SAIL-Toyota Center for AI Research at Stanford (O. Khatib).

References

1. Schaal, S., Peters, J., Nakanishi, J., Ijspeert, A.: Learning movement primitives. In: Robotics Research. The Eleventh International Symposium, Springer (2005) 561–572
2. Khatib, O.: The potential field approach and operational space formulation in robot control. Adaptive and Learning Systems: Theory and Applications (1986) 367–377
3. Ruspini, D., Khatib, O.: Haptic display for human interaction with virtual dynamic environments. Journal of Robotic Systems **18**(12) (2001) 769–783
4. Mason, M.T.: Compliance and force control for computer controlled manipulators. IEEE Transactions on Systems, Man, and Cybernetics **11**(6) (1981) 418–432
5. Whitney, D.E.: Historical perspective and state of the art in robot force control. The International Journal of Robotics Research **6**(1) (1987) 3–14
6. Hogan, N.: Stable execution of contact tasks using impedance control. In: Robotics and Automation. Proceedings. 1987 IEEE International Conference on. Volume 4., IEEE (1987) 1047–1054
7. Featherstone, R., Thiebaut, S.S., Khatib, O.: A general contact model for dynamically-decoupled force/motion control. ICRA **4** (May 1999) 3281–3286
8. Park, J., Khatib, O.: A haptic teleoperation approach based on contact force control. The International Journal of Robotics Research **25**(5-6) (June 2006) 575–591
9. Wang, D., Zhang, X., Zhang, Y., Xiao, J.: Configuration-based optimization for six degree-of-freedom haptic rendering for fine manipulation. In: Robotics and Automation (ICRA), 2011 IEEE International Conference on. (2011) 906–912
10. Ji, X., Xiao, J.: Planning motions compliant to complex contact states. The International Journal of Robotics Research **20**(6) (2001) 446–465
11. Xiao, J., Ji, X.: Automatic generation of high-level contact state space. The International Journal of Robotics Research **20**(7) (2001) 584–606
12. Kwak, S.J., Chung, S.Y., Hasegawa, T.: Generating a contact state graph of polyhedral objects for robotic application. In: Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on. (2010) 4522–4527

13. Meeussen, W., Staffetti, E., Bruyninckx, H., Xiao, J., De Schutter, J.: Integration of planning and execution in force controlled compliant motion. Robotics and Autonomous Systems **56**(5) (2008) 437–450
14. Meeussen, W., Rutgeerts, J., Gadeyne, K., Bruyninckx, H., De Schutter, J.: Contact-state segmentation using particle filters for programming by human demonstration in compliant-motion tasks. Robotics, IEEE Transactions on **23**(2) (2007) 218–231
15. Skubic, M., Volz, R.A.: Acquiring robust, force-based assembly skills from human demonstration. IEEE Transactions on Robotics and Automation **16**(6) (Dec 2000) 772–781
16. Bruyninckx, H., De Schutter, J.: Specification of force-controlled actions in the ldquo;task frame formalism rdquo;-a synthesis. Robotics and Automation, IEEE Transactions on **12**(4) (1996) 581–589
17. Onda, H., e.a.: Assembly motion teaching system using position/force simulator-generating control program. In: Intelligent Robots and Systems, Proceedings of the IEEE/RSJ International Conference on. Volume 2. (Sep 1997) 938–945 vol.2
18. Gadeyne, K., Lefebvre, T., Bruyninckx, H.: Bayesian hybrid model-state estimation applied to simultaneous contact formation recognition and geometrical parameter estimation. The International Journal of Robotics Research **24**(8) (2005) 615–630
19. Klingbeil, E., Menon, S., Go, K.C., Khatib, O.: Using haptics to probe human contact control strategies for six degree-of-freedom tasks. In: Haptics Symposium (HAPTICS), IEEE (2014) 93–95
20. Kandel, E.R., Schwartz, J.H., Jessell, T.M.: Principles of Neural Science. Volume 4. New York: McGraw-Hill, Health Professions Division (2000)
21. Xiao, J.: Automatic determination of topological contacts in the presence of sensing uncertainties. In: Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on. (May 1993) 65–70 vol.1
22. Tobergte, A., Helmer, P., Hagn, U., Rouiller, P., Thielmann, S., Grange, S., Albu-Schaffer, A., Conti, F., Hirzinger, G.: The sigma.7 haptic interface for mirosurge: A new bi-manual surgical console. In: Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on. (2011) 3023–3030
23. Smith, R.: Open dynamics engine (2010)
24. Tan, H.Z., Srinivasan, M.A., Eberman, B., Cheng, B.: Human factors for the design of force-reflecting haptic interfaces. Dynamic Systems and Control **55**(1) (1994) 353–359
25. Ruspini, D., Khatib, O.: A framework for multi-contact multi-body dynamic simulation and haptic display. In: Intelligent Robots and Systems, 2000.(IROS 2000). Proceedings. 2000 IEEE/RSJ International Conference on. Volume 2., IEEE (2000) 1322–1327
26. McLaughlin, M.L., Hespanha, J.P., Sukhatme, G.S.: Touch in virtual environments. Prentice Hall (2002)
27. Salisbury, J.K., Conti, F., Barbagli, F.: Haptic rendering: Introductory concepts. IEEE Computer Graphics and Applications **24**(2) (March 2004) 24–32
28. McNeely, W.A., Puterbaugh, K.D., Troy, J.J.: Six degree-of-freedom haptic rendering using voxel sampling. In: ACM SIGGRAPH 2005 Courses, ACM (2005) 42
29. Kuchenbecker, K., Fiene, J., Niemeyer, G.: Improving contact realism through event-based haptic feedback. Visualization and Computer Graphics, IEEE Transactions on **12**(2) (2006) 219–230
30. Efron, B., Efron, B.: The jackknife, the bootstrap and other resampling plans. Volume 38. SIAM (1982)