# Safe and Robust Robot Learning from Demonstration through Conceptual Constraints

Carl L. Mueller
carl.mueller@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

Bradley Hayes
bradley.hayes@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

## ABSTRACT

This thesis summary presents research focused on incorporating high-level abstract behavioral requirements, called 'conceptual constraints', into the modeling processes of robot Learning from Demonstration (LfD) techniques. This idea is realized via an LfD algorithm called *Concept Constrained Learning from Demonstration*. This algorithm encodes motion planning constraints as temporally associated logical formulae of Boolean operators that enforce high-level constraints over portions of the robot's motion plan during learned skill execution. This results in more easily trained, more robust, and safer learned skills. Current work focuses on automating constraint discovery, introducing conceptual constraints into human-aware motion planning algorithms, and expanding upon trajectory alignment techniques for LfD. Future work will focus on how concept constrained algorithms and models are best incorporated into effective interfaces for end-users.

## CCS CONCEPTS

• **Computing methodologies** → **Learning from demonstrations**; *Robotic planning*; • **Computer systems organization** → *Robotic autonomy*; External interfaces for robotics.

## KEYWORDS

Learning from Demonstrations, Robotic Learning, Human-Robot Interaction

**ACM Reference Format:**
Carl L. Mueller and Bradley Hayes. 2020. Safe and Robust Robot Learning from Demonstration through Conceptual Constraints. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20 Companion), March 23–26, 2020, Cambridge, United Kingdom.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3371382.3377428

## 1 INTRODUCTION

Modern industrial robots predominantly exist in the realm of large-scale processes; those that are highly repetitive, precise, and relatively unchanging [4]. A blossoming niche of robotics research called Human-Robot Interaction focuses on robots designed or programmed to work with human counterparts [3] and may extend the benefits enjoyed by large-scale industrial automation to more dynamic small-scale industries. However, human-robot collaboration presents a number of challenges not often present in industrial settings: safety in shared workspaces, rapidly changing task requirements, decision-making, and adhering to human expectations of behavior. Recent advances in AI and robotics provide the means to overcome such challenges, inviting a new era of more capable, adaptable, and collaborative robots that expands automation into industries previously inaccessible to automation. As such, the thesis work summarized herein focuses on providing human users the means to easily train a collaborative robot to execute dynamic skills using robotic Learning from Demonstration (LfD) techniques while adhering to important behavioral restrictions.


ACM ISBN 978-1-4503-7057-8/20/03.
https://doi.org/10.1145/3371382.3377428

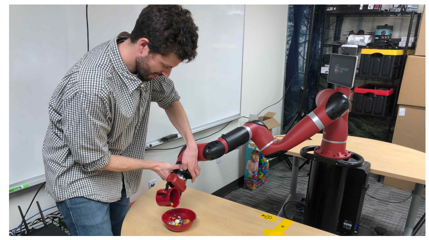

**Figure 1:** A user teaching a robot a task via kinesthetic learning.

## 2 PRIOR WORK

*Concept Constrained LfD.* The underlying motivation for this thesis research is the idea that incorporating abstract behavioral restrictions into robotic LfD methods might encourage an awareness of safety requirements and increase the learning efficiency of the system [11]. LfD comprises techniques that enable non-expert users with no programming knowledge to teach a robot how to perform a task [3]. Low-level data (e.g. configuration space) traditionally used by these techniques poorly captures important factors and abstract concepts essential to successful skill learning [5, 7, 13]. A robotic learning system must possess awareness of high-level considerations to truly 'learn' a skill. For example, when teaching a robot how to carry a cup of coffee, low-level demonstration data will not adequately inform the system that the cup *must remain upright if it contains liquid.*

An algorithm presented in [12] called Concept Constrained Learning from Demonstration (CC-LfD) introduces 'conceptual constraints' to represent abstract restrictions on the behavior of the robot (e.g. keeping a cup of coffee upright). These constraints are encoded as Boolean operators that evaluate whether a given environment state (e.g. c-space, relative distances, and sensor data) satisfies the high-level abstract idea it represents. By augmenting low-level robot state data with high-level abstract information a learned model will much more closely resemble the ground truth representation of a task or skill.

CC-LfD enables users to temporally assign constraints during demonstration. These constraints are incorporated into a technique called Keyframe LfD [1, 2] where data points of temporally aligned demonstration trajectories are clustered into sequential groups across demonstrations. Distributions learned using the data within each cluster form *keyframe* models. These models are sequentially sampled for waypoints that the robot follows to perform a skill. In CC-LfD, constraints across demonstrations are combined during alignment to create disjunctive normal form formulae, which are assigned to temporally appropriate keyframes. Waypoints generated from the keyframe models pass through a rejection sampling filter where each point is evaluated using the keyframe's constraint formula. This ensures the robot follows constraint-compliant waypoints. As a consequence, robotic learning systems employing CC-LfD require far fewer demonstrations than standard Keyframe LfD to produce robust learned skills [12]. For example, in a cup pouring task, when given three demonstration trajectories that poorly perform the skill, one additional constraint-annotated demonstration resulted in a high level of objective skill performance. Without the constraint annotations, many more 'gold' demonstrations were required to achieve improvement in skill performance [12].

## 3 CURRENT WORK

*Autonomous Concept Constrained LfD.* CC-LfD enables users to communicate constraints over a demonstrated skill, but the burden that constraint assignment places on the user during demonstration can be prohibitive. While conceptual constraints can capture very high-level information, they also easily capture low-level constraints such as orientation and position constraints [11]. For a given skill, these types of constraints (such as object-object relative position) can be numerous thereby placing a significant cognitive load on a user as they must maintain awareness of all constraints. With this in mind, Autonomous Concept Constrained LfD (ACC-LfD) seeks to automate the currently entirely human-driven selection and annotation of constraints while retaining the benefits of CC-LfD. The motivation is to automate low-level constraint assignment, freeing the user to focus on assigning the most abstract conceptual constraints that might be very difficult or impossible to learn.

Inspired by the Transition State Clustering (TSC) algorithm [10], ACC-LfD uses a combination of Variational Gaussian Mixture Models (VGMM) to cluster keyframes and constraint-specific heuristics to parameterize conceptual constraints. Similar to the TSC algorithm, a VGMM clusters keyframes based on common information contained within the demonstration data. Using a conceptual constraint that restricts the orientation of a cup as an example, a VGMM might generate two clusters representative of a pouring task: an upright orientation cluster and a pouring orientation cluster. Heuristics for this orientation constraint could be the 'average' orientation and angle of deviation from that average, calculated using the data points within each cluster. This average and the angle of deviation would thus populate the constraint's parameters to evaluate keyframe sample points for the orientation constraint.

*Concept Constrained Motion Planning.* Both CC-LfD and ACC-LfD do not consider conceptual constraints when relying on motion planning algorithms to traverse between waypoints. This forces both algorithms to require more tightly spaced keyframing than necessary in order to avoid constraint violation during intermediate poses. One way to address this problem is by incorporating conceptual constraints into existing offline and online motion planning algorithms. A key challenge of this work is that constrained planning often must occur in a higher-dimensional space than conventional fast configuration-space planning allows. Offline sampling based motion planners (e.g. RRT*, KPIECE) [9, 15] require abstract constraints to have a geometric representation in this space, or be integrated into a cost function that evaluates the generated local plans for constraint compliance. Local plans are the small incremental movements of joints that the robot conducts during the execution of a chosen automated motion plan. Similarly, this cost function could be used to scale the first and second order vector fields generated over the state space employed by Dynamic Motion Primitive algorithms, such as the end-effector space [8].

*Robot Trajectory Alignment.* Successful keyframe-based LfD techniques depend upon adequate alignment of demonstration trajectory data. Many keyframe techniques rely on a prominent time-series alignment algorithm called Dynamic Time Warping (DTW) [14]. This algorithm produces an alignment of indices and overall distance/cost between two time series. For the purposes of Keyframe LfD, this alignment step provides a crucial linkage between the data points in one demonstration trajectory to the data points in all the other trajectories. Violating the underlying assumption that clustered data used to model a keyframe represents the same aspect of a skill may result in a poorly representative model. Much of the LfD literature does not consider the pros and cons of different approaches to DTW [1, 2, 6, 16]: alignment on different state spaces other than configuration or end-effector space; combining DTW with other clustering approaches such as k-nearest neighbors techniques or mixture model techniques. Current work of this thesis is exploring how these considerations might ultimately produce better representative keyframe models and how they might enable learned models to account for multi-directional demonstration paths that accomplish the same task.

## 4 FUTURE WORK

While concept constrained algorithms and motion planning might provide effective means to inject abstract information into LfD methods, the process of injection must be considered. Designing more sophisticated interfaces that provide adaptable and efficient means of communicating constraints constitutes the final research effort of this thesis work prior to the completion of a dissertation. This future research will investigate methods to visualize constraints via augmented reality that are intuitive to the user. Such visual interfaces might enable a user to edit the parameterization of the encoded constraints. For example, the user might edit what is considered an allowed orientation for a cup carrying task. These interfaces will be evaluated for efficacy with human-subjects studies that explore both the objective performance increases in robotic skill execution and the subjective burden placed upon human users.

## ACKNOWLEDGMENTS

This work was funded in part by NSF Award #1830686.

## REFERENCES


[1] Baris Akgun, Maya Cakmak, Karl Jiang, and Andrea L Thomaz. 2012. Keyframe-based learning from demonstration. *International Journal of Social Robotics* 4, 4 (2012), 343–355.

[2] Baris Akgun, Maya Cakmak, Jae Wook Yoo, and Andrea Lockerd Thomaz. 2012. Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 391–398.

[3] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems* 57, 5 (2009), 469–483.

[4] Mohd Aiman Kamarul Bahrin, Mohd Fauzi Othman, NH Nor Azli, and Muhamad Farihin Talib. 2016. Industry 4.0: A review on industrial automation and robotic. *Jurnal Teknologi* 78, 6-13 (2016), 137–143.

[5] Sonia Chernova and Andrea L Thomaz. 2012. Special Issue on Robot Learning from Demonstration.

[6] Tesca Fitzgerald, Ashok K Goel, and Andrea L Thomaz. 2014. Representing skill demonstrations for adaptation and transfer. In *2014 AAAI Fall Symposium Series*.

[7] Michael A Goodrich, Alan C Schultz, et al. 2008. Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction* 1, 3 (2008), 203–275.

[8] Auke Jan Ijspeert, Jun Nakanishi, Heiko Hoffmann, Peter Pastor, and Stefan Schaal. 2013. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation* 25, 2 (2013), 328–373.

[9] Sertac Karaman and Emilio Frazzoli. 2010. Incremental sampling-based algorithms for optimal motion planning. *arXiv preprint arXiv:1005.0416* (2010).

[10] Sanjay Krishnan, Animesh Garg, Sachin Patil, Colin Lea, Gregory Hager, Pieter Abbeel, and Ken Goldberg. 2018. Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning. In *Robotics Research*. Springer, 91–110.

[11] Andrey Kurenkov, Baris Akgun, and Andrea L Thomaz. 2015. An evaluation of GUI and kinesthetic teaching methods for constrained-keyframe skills. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3608–3613.

[12] Carl Mueller, Jeff Venicx, and Bradley Hayes. 2018. Robust Robot Learning from Demonstration and Skill Repair Using Conceptual Constraints. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6029–6036.

[13] Claudia Pérez-D'Arpino and Julie A Shah. 2017. C-learn: Learning geometric constraints from demonstrations for multi-step manipulation in shared autonomy. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4058–4065.

[14] Hiroaki Sakoe and Seibi Chiba. 1978. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing* 26, 1 (1978), 43–49.

[15] Ioan A Şucan and Lydia E Kavraki. 2011. A sampling-based tree planner for systems with complex dynamics. *IEEE Transactions on Robotics* 28, 1 (2011), 116–131.

[16] Aleksandar Vakanski, Iraj Mantegh, Andrew Irish, and Farrokh Janabi-Sharifi. 2012. Trajectory learning for robot programming by demonstration using hidden Markov model and dynamic time warping. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 42, 4 (2012), 1039–1052.