

# Abstract Constraints for Safe and Robust Robot Learning from Demonstration

**Carl L. Mueller**

carl.mueller@colorado.edu  
University of Colorado Boulder  
1111 Engineering Drive  
Boulder, Colorado 80309

## Abstract

My thesis research incorporates high-level abstract behavioral requirements, called ‘conceptual constraints’, into the modeling processes of robot Learning from Demonstration (LfD) techniques. My most recent work introduces an LfD algorithm called *Concept Constrained Learning from Demonstration*. This algorithm encodes motion planning constraints as temporal 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. Future work will incorporate conceptual constraints into human-aware motion planning algorithms. Additionally, my research will investigate how these concept constrained algorithms and models are best incorporated into effective interfaces for end-users.

**Introduction** Whether they are articulated arms in automotive factories or Cartesian platforms of industrial drug manufacturing, most robots exist in the realm of large-scale industrial processes; those that are highly repetitive, precise, and relatively unchanging (Bahrin et al. 2016). A blossoming niche of robotics research called Human-Robot Interaction focuses on robots designed or programmed to work with human counterparts (Argall et al. 2009). Such robots have the potential to expand 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, perhaps most challenging, adhering to human expectations of behavior. Overcoming such challenges will invite a new era of more capable, adaptable, and collaborative robotics that revolutionize a wide range of industries previously inaccessible to automation. Recent advances in AI and robotics have provided the necessary foundations to effect transformative change. As such, my thesis work focuses on providing human users the means to easily train a collaborative robot to execute dynamic skills while adhering to important behavioral restrictions.

## Concept Constrained Learning from Demonstration

The first two years of my research have been motivated



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

by the idea that incorporating abstract behavioral restrictions into robotic learning methods might precipitate safety awareness and increase the learning efficiency of the system. My first paper, accepted to the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems, realizes this approach by introducing an algorithm called Concept Constrained Learning from Demonstration (CC-LfD) (Mueller, Venicx, and Hayes 2018). The core of CC-LfD draws from an area of robotics research known as Learning from Demonstration (LfD), comprising a set of techniques enabling non-expert users with no programming knowledge to teach a robot how to perform a task. Traditionally, these techniques utilize teleoperation, kinesthetic learning, or imitation learning, methods of demonstration that record low-level data such as end-effector position and robot configuration (Goodrich, Schultz, and others 2008).

Low-level data has limited information bandwidth for capturing important factors and abstract concepts essential to successful skill learning and execution (Chernova and Thomaz 2012). CC-LfD introduces ‘conceptual constraints’ to represent abstract restrictions on the behavior of the robot (e.g. keeping a pitcher upright until over a cup). These constraints are encoded as Boolean operators that evaluate whether a given low-level environment state satisfies the high-level abstract idea it represents. The motive is to augment low-level robot state data with high-level abstract information such that the learned model much more closely resembles the ground truth representation of a task or skill.

CC-LfD enables users to dictate when and where conceptual constraints must hold true during the demonstration of a task. These constraints are incorporated into a technique called Keyframe LfD (Akgun et al. 2012) where the data points of temporally aligned demonstration trajec-

tories are clustered into sequential groups across demonstrations. Through statistical modeling of the clusters into *keyframes*, these clusters can be used to generate waypoints that the robot follows sequentially to perform a skill. In CC-LfD, waypoints generated from the keyframe models pass through a rejection sampling filter where each point is evaluated with the constraints' Boolean operators assigned to the given keyframe. This ensures the robot moves through a sequence of constraint-compliant waypoints. As a consequence, we were able to show that robotic learning systems employing CC-LfD require far fewer demonstrations than standard Keyframe LfD to produce robust learned skills.

I am currently working on extending CC-LfD toward Autonomous Concept Constrained LfD (ACC-LfD), automating the currently entirely human-driven selection and annotation of constraints while retaining the benefits of CC-LfD. This work serves to substantially reduce the burden that constraint assignment places on the user during demonstration. Inspired by the Transition State Clustering (TSC) algorithm (Krishnan et al. 2018), 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.

**Constrained Compliant Robot Motion Planning** Both CC-LfD and ACC-LfD suffer from a common issue in that they do not consider conceptual constraints when relying on offline 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. To address this problem, I will be focusing on incorporating conceptual constraints into existing online and offline motion planning algorithms (to be submitted to AAAI 2021). 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. Thus, each abstract constraint must either have a geometric representation in this space, or must utilize 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 (Ijspeert et al. 2013).

**Natural Language and Augmented Reality Interfaces** 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 my degree (2021 – 2022 timeframe) and the final piece and capability required for my PhD work's proposed LfD system. Natural language interfaces can provide an intuitive and high-information bandwidth mechanism for users to dictate conceptual constraints. At present, CC-LfD uses a very simplistic natural language interface, consisting mainly of keyword recognition or one word commands. A major challenge inhibiting the use of more sophisticated language interfaces is the ambiguity naturally present in all languages. When a user says, "keep a safe distance from any electronics in the environment.", the robotic system must parameterize the 'safe distance' constraint and rely on object classification to identify where the electronics are in the workspace.

One solution could be to present the user with an understandable visualization of the encoded constraints to provide assurance that the learning system correctly encoded communicated constraints. My research will investigate methods to visualize constraints via augmented reality that are intuitive to the user as a supplement to the natural language interface. This interface will also 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.

## References

- Akgun, B.; Cakmak, M.; Jiang, K.; and Thomaz, A. L. 2012. Keyframe-based learning from demonstration. *International Journal of Social Robotics* 4(4):343–355.
- Argall, B. D.; Chernova, S.; Veloso, M.; and Browning, B. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems* 57(5):469–483.
- Bahrin, M. A. K.; Othman, M. F.; Azli, N. N.; and Talib, M. F. 2016. Industry 4.0: A review on industrial automation and robotic. *Jurnal Teknologi* 78(6-13):137–143.
- Chernova, S., and Thomaz, A. L. 2012. Special issue on robot learning from demonstration.
- Goodrich, M. A.; Schultz, A. C.; et al. 2008. Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction* 1(3):203–275.
- Ijspeert, A. J.; Nakanishi, J.; Hoffmann, H.; Pastor, P.; and Schaal, S. 2013. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation* 25(2):328–373.
- Krishnan, S.; Garg, A.; Patil, S.; Lea, C.; Hager, G.; Abbeel, P.; and Goldberg, K. 2018. Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning. In *Robotics Research*. Springer. 91–110.
- Mueller, C.; Venicx, J.; and Hayes, B. 2018. Robust robot learning from demonstration and skill repair using conceptual constraints. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 6029–6036. IEEE.