Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)
Journal Track

Abstraction in Data-Sparse Task Transfer (Extended Abstract)*

Tesca Fitzgerald[1†], **Ashok Goel**[2], **Andrea Thomaz**[3]
[1]Carnegie Mellon University
[2]Georgia Institute of Technology
[3]University of Texas at Austin
tesca@cmu.edu, goel@cc.gatech.edu, athomaz@ece.utexas.edu

Abstract

When a robot adapts a learned task for a novel environment, any changes to objects in the novel environment have an unknown effect on its task execution. For example, replacing an object in a pick-and-place task affects where the robot should target its actions, but does not necessarily affect the underlying action model. In contrast, replacing a tool that the robot will use to complete a task will effectively alter its end-effector pose with respect to the robot's base coordinate system, and thus the robot's motion must be replanned accordingly. In this abstract, summarizing our full article [Fitzgerald *et al.*, 2021a], we present our taxonomy of transfer problems based on the effect of environment changes on a robot's ability to complete a task. We also describe a knowledge representation called the Tiered Task Abstraction (TTA) and demonstrate its applicability to a variety of transfer problems within this taxonomy. Our experimental results indicate a trade-off between the generality and data requirements of a task representation, and reinforce the need for multiple transfer methods that operate at different levels of abstraction.

1 Introduction

To date, robots do not have the ability to adapt their behavior to the wide variety of novel objects, tasks, and interactions that are inherent to everyday life. For example, a robot tasked with restocking objects in a warehouse will need to adapt to visual variations in product branding, the unique grasping constraints of different objects, and the unpredictability of human co-workers in the robot's vicinity. As our expectations of robots' interactive and adaptive capacity grow, it will be increasingly important for them to reason about and respond to novel situations appropriately.

The high-dimensionality of robot perception and action presents a challenge when adapting a robot's task knowledge to accommodate novelty in its environment. While a robot can learn a "task model" from a teacher's demonstration of a task [Argall *et al.*, 2009; Chernova and Thomaz, 2014; Schaal, 2006; Pastor *et al.*, 2009], demonstrating these behaviors for all variations of that task would require tedious effort, and still would not adequately prepare a robot for every scenario it may encounter. This is due to the various dimensions along which the original ("source") and novel ("target") settings may differ, such as changes in objects, manipulation tools, constraints, and dynamics that the robot must adapt to. Demonstrations alone are insufficient to enable a robot to generalize to novel situations as they do not encode the relationship between these *task differences* and the corresponding modifications to the robot's task model. While some of these task differences have been addressed for reinforcement learning agents [Taylor and Stone, 2009], they do not address all task differences that a robot may encounter (e.g., tool changes) and typically rely on the robot receiving substantial training data in the target environment.

Rather than attempt to pre-train a robot for all task variations it will encounter, we posit that we can develop more capable and robust robots by assuming that they will inevitably encounter novelty that they are initially unprepared to address. We thus consider a data-sparse paradigm, with the goal of enabling a robot to reason over the information it needs to address whatever specific form of novelty it encounters. We particularly focus on the effect of changes in the robot's *environment* on task execution (rather than adapting to changes in other aspects of the task, such as goals or reward functions).

This abstract summarizes our work [Fitzgerald *et al.*, 2021a] characterizing the problem of task transfer in terms of similarity between the source and target environments. We summarize the Tiered Task Abstraction (TTA): a generalized task representation that implements these principles of task similarity and abstraction. We then demonstrate the TTA representation's effectiveness on a physical robot transferring two pick-and-place tasks to a variety of target environments. Finally, we conclude with a discussion of autonomous and collaborative methods for grounding an abstracted task representation in a new environment.

2 Taxonomy: Categorizing Task Differences

We refer to a *task* as a sequence of object-oriented task steps, each consisting of their own action model and performed in series to achieve a goal. As an example, a *scooping* task may

*This abstract summarizes our full article published in Artificial Intelligence Journal.
†Contact Author

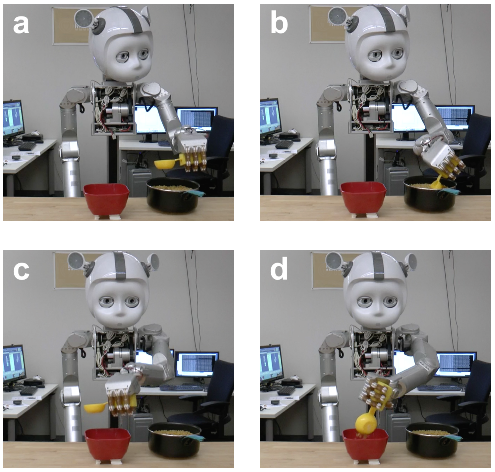

Figure 1: Steps comprising a *scooping* task

consist of several steps (Fig. 1): moving the scoop from the initial position at the robot's side to the pasta bowl, scooping the pasta, moving the scoop to the target bowl, and then pouring the scoop over the target bowl. This example suggests three key elements of a task representation: the robot's state with respect to its environment (e.g. objects), the action model comprising each task step, and the goal that is achieved by executing all action models. These correspond to the state space, action space, and goal/rewards commonly used to define a Markov Decision Process or other task-planning problem. In this article, we only address transfer problems that result from changes in the robot's state space, as opposed to changes in the robot's goal or action spaces. For a scooping task originally learned in the environment shown in Figure 2a, Figs. 2b-e illustrate a range of possible state space changes. We categorize these changes as follows:

- *Structural* changes in which the relationship between objects within the robot's environment is altered. E.g., Fig. 2b, where objects are moved around the scene.
- *Perceptual* changes in which the robot's environment appears different while remaining functionally the same. E.g., Fig. 2c containing new objects.
- *Functional* changes that affect the relationship between the robot and its environment. E.g., Fig. 2d, containing a new tool that constrains the robot's motion when in use.

These task differences illustrate a *spectrum* of similarity between the source and target; at one end of the spectrum, the source and target differ in ways that have a small effect on the robot's execution of the task, such as object configurations. At the other end of the spectrum, they contain more differences, until finally (as in Fig. 2e, where there is nothing for the robot to scoop), the target environment cannot be addressed via transfer. While we have highlighted discrete levels of similarity in this spectrum, we do not claim this to be an exhaustive categorization of transfer problems. In prior work, we have also explored the requirements for creativity in order to address more dissimilar environments [Fitzgerald *et al.*, 2017].

3 Representation: Tiered Task Abstraction (TTA)

Based on this similarity spectrum, we propose the Tiered Task Abstraction (TTA) representation to reflect the relationship between (1) changes in the state space and (2) their effect on the task transfer. Overall, the TTA representation consists of four elements: an action model $\mathbf{a}$, parameterization function $\mathbf{p}$, feature selector $\mathbf{f}$, and feature values E. These four elements represent a task demonstration as a series of action models as follows:

$$\mathbf{a_0}(\mathbf{p_0}(\mathbf{f_0}(E))), \ldots, \mathbf{a_m}(\mathbf{p_m}(\mathbf{f_m}(E))) \qquad (1)$$

where:

- $\mathbf{a_i}(\mathbf{p})$ is an **action model** (e.g., a Gaussian Mixture Model, Dynamic Movement Primitive, or neural network) that is parameterized by p and outputs a trajectory consisting of a series of poses for the robot to execute. In our implementation, we used a Dynamic Movement Primitive (DMP) to model each step of the task.
- $\mathbf{p_i}(\mathbf{f})$ is a **parameterization function** for the action model a. The parameters returned by this function serve as a task-relevant encoding of the input features f. In our implementation, we defined a parameterization function for each task model that represents the end-effector's position with respect to the nearest object in the robot's environment (e.g., defining scooping as being performed at the surface of bowl #1, and later, pouring as being performed 7cm above bowl #2).
- $\mathbf{f_i}(\mathbf{E})$ serves as a **feature selector** that returns a task-relevant feature representation of the robot's **input space** E. In our implementation, this consisted of unique, non-descriptive object IDs assigned to the segmented objects in the robot's environment.

This definition of each element in the TTA representation is intentionally conceptual, as the specific implementation, input, and output of each of these functions is dependent on the robot and the domain in which it is deployed. **The defining characteristic of the TTA representation is that each element is parameterized by the next.** By omitting one or more elements from the task representation, the resulting representation is one that is *abstracted*. In doing so, TTA enables a task to be represented at the level of abstraction that is common to both the source and target environments. Figure 2 defines three abstractions of this representation.

Once a representation is abstracted, it must be *grounded* in the target environment in order to produce an output that is executable by the robot. In an embodied system, such as a robot, grounding refers to parameterizing a representation based on perception of the physical world. A representation is *grounded* in a target environment when all of its elements (action model, parameterization function, feature selector, and feature values) are present and defined based on information derived in the target environment (either by perception or interaction in the target environment). If the representation cannot be fully grounded in a target environment, the robot may need to re-learn the task within the context of the target environment. In the context our experiments, we manually provide the information needed to ground each abstraction.

	Non-Abstracted	Abstraction 1	Abstraction 2	Abstraction 3	Dissimilar
Retained Knowledge	Action Models Parameterization Functions Feature Selectors Feature Values	Action Models Parameterization Functions Feature Selectors	Action Models Parameterization Functions	Action Models	None
Grounded Knowledge	None	Feature Values	Feature Selectors Feature Values	Parameterization Functions Feature Selectors Feature Values	Action Models Parameterization Functions Feature Selectors Feature Values

Figure 2: The grounding requirements for each level of abstraction

4 Case Study: Transferring a Task At Multiple Abstraction Levels

We evaluate whether the TTA representation reflects the relationship between task differences and the resulting data requirements to enable task transfer. To do this, we demonstrate two tasks (*table-setting* and *scooping*) on the Curi robot's 7-DOF arm (Fig. 1), model the task demonstrations separately using the TTA representation, and then apply abstractions of each representation to three task variations:

1. *Displaced-Object variations*: Contain the same objects as in the original demonstration, but displaced as shown in Fig. 2b (for the scooping task) and Fig. 4b (for the table-setting task).
2. *Replaced-Object variations*: Contain objects that are different than those used in the original demonstration. Additional "distractor" objects are also provided that are irrelevant to completing the task. An example is shown in Fig. 2c (scooping) and Fig. 4c (table-setting).
3. *New-Constraint variations*: Contain the same objects as in the replaced-object environment, but with an additional constraint. In the table-setting task, the cup and utensil are jointly rotated 45-90 degrees away from the robot (Fig. 4d). In the scooping task, the scoop is replaced with one containing a longer handle (Fig. 2d).

We represent each task model at the *three* levels of abstraction shown in Figure 2 and manually provide the information needed to ground each abstracted representation in the target environment (see [Fitzgerald *et al.*, 2021a] for more details). We then evaluate each (grounded) abstraction on *ten* target environment variations in each of the *three* variation categories, resulting in a total of *90* transfer evaluations for each of the two tasks. Figure 3 provides the success rate of each abstraction level when grounded and applied to variations of the table-setting and scooping environments. Abstraction 1 succeeded consistently on *only* the displaced-objects environments. Abstraction 2 resulted in consistent performance in the displaced- and replaced-objects environments, and additionally, succeeded in a few of the new-constraint variations of the table-setting task. Finally, these results also indicate that Abstraction 3 succeeded consistently across all three categories of target environments.

5 Discussion: Trade-Off Between Generality and Data-Efficiency

These results suggest that Abstraction 3 provides the most consistently successful results across the full range of transfer problems we tested, with Abstractions 2 and 1 each addressing fewer transfer problems, respectively. However, the more that the task representation is abstracted, the more data is required to ground that abstraction in the target environment, indicating a **trade-off between the generality of a task representation, and the amount of additional information required to ground that task representation in the target environment**. For a robot that operates in human environments, we aim for the robot to be able to ground its own task abstractions using continued interaction with a human teacher during task transfer. The aim of interactive grounding is to perform task transfer in a manner that (i) is partially autonomous (the robot interacts with a human teacher and may receive additional instruction, but does not require a full re-demonstration of the task), (ii) enables collaboration with the human teacher so that the robot may infer information about the task in the target environment, (iii) results in parameterization functions and/or action models that can ground an abstracted task representation, and (iv) grounds the TTA representation such that a trajectory can be executed in the target environment.

In prior work, we have defined two methods for interactive grounding, each targeting a different category of transfer problems. We have presented Mapping by Demonstra-

	Displaced Objects	Replaced Objects	New Task Constraints
Abstraction 1 (Least-abstracted)	100%	0%	0%
Abstraction 2	100%	100%	0%
Abstraction 3 (Most-abstracted)	100%	90%	80%

(a) Scooping Task

	Displaced Objects	Replaced Objects	New Task Constraints
Abstraction 1 (Least-abstracted)	100%	0%	10%
Abstraction 2	90%	100%	40%
Abstraction 3 (Most-abstracted)	100%	100%	70%

(b) Table-Setting Task

Figure 3: Success rates for each abstraction level when applied to two tasks

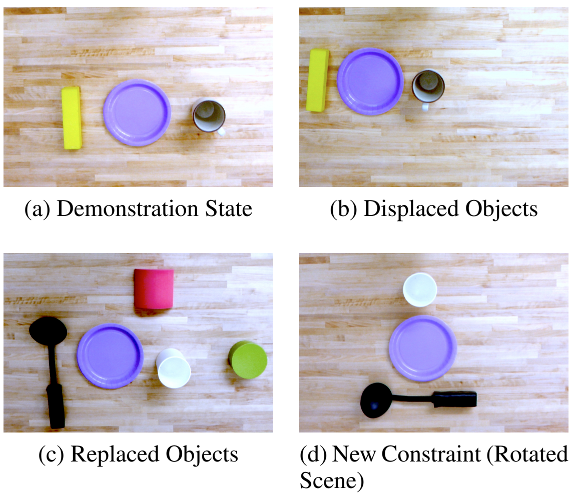

(a) Demonstration State (b) Displaced Objects

(c) Replaced Objects (d) New Constraint (Rotated Scene)

Figure 4: Variants of the *Table-Setting* Task Environment

tion [Fitzgerald *et al.*, 2018], in which a robot utilizes a targeted method of interaction with a human teacher (indicating the next object the robot should use) in order to infer an *object mapping* between objects in the source and target environments. This method enables the robot to use limited assistance with the first part of the task to infer which objects should be used to complete the remainder of the task autonomously. The resulting object mapping enables the robot to ground Abstraction 2 and thus address transfer problems with replaced objects.

We have also presented a method for grounding Abstraction 3 for transfer problems in which object replacements alter the manipulation constraints of the task. To address this category of transfer problems, a different interaction mode is needed to ground the relationship between (1) the new object and (2) the trajectory adaptations necessary to use the new object. We have employed *corrections* to record and model constrained points in the robot's motion [Fitzgerald *et al.*, 2019]. Furthermore, we presented a method for modeling the new constraints afforded by the tool within the context of the corrected task, and demonstrated that the learned model can also be reused on other tasks that provide a similar context for that tool (e.g. in the tool surfaces used to execute the task) [Fitzgerald *et al.*, 2021b].

Future work can leverage additional, existing methods for interactive robot learning for grounding task abstractions, such as preferences [Sadigh *et al.*, 2017], critiques [Cui and Niekum, 2018], and dialogue [Scheutz *et al.*, 2017]. The selection of interaction type affects the quantity and specificity of the data that is derived from interpreting the teacher's resulting feedback [Cui *et al.*, 2021]. As a result, the relationship between (i) the level of abstraction used to represent the task and (ii) the information needed to ground the abstracted representation thus **provides guidance for selecting the interaction modality that is best suited for a particular transfer problem.**

6 Conclusion

This article analyzes task transfer in the data-sparse context of interactive robot learning, and results in four key insights:

1. The more dissimilar the source and target environments are, the more that the source task representation must be abstracted in order to be successfully transferred.
2. There is a correlation between (i) the degree to which the task representation is abstracted and (ii) the amount of data that is needed to ground the abstracted representation in the target environment.
3. As a result of #1 and #2, there is a tradeoff between the generality of a task representation (e.g. the range of transfer problems that it can successfully address) and the data requirements that must be met to ground the abstracted task representation in the target environment.
4. The modality of the interaction (e.g., gestures, corrections, or dialogue) between the robot and human teacher has a direct impact on the information the robot can derive from the teacher's feedback. As a result, the grounding requirements of the abstracted task representation must be taken into consideration when selecting the interaction modality used to obtain that grounded data.

Our work is the first to both apply an abstraction-based approach to the problem of robotic task transfer and ground it in action, perception, and interaction. For more details, see our full paper [Fitzgerald *et al.*, 2021a].

Acknowledgments

This material is based on work supported by the NSF Graduate Research Fellowship under Grant No. DGE-1148903.

References

[Argall *et al.*, 2009] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.

[Chernova and Thomaz, 2014] Sonia Chernova and Andrea L Thomaz. Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3):1–121, 2014.

[Cui and Niekum, 2018] Yuchen Cui and Scott Niekum. Active reward learning from critiques. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6907–6914. IEEE, 2018.

[Cui *et al.*, 2021] Yuchen Cui, Pallavi Koppol, Henny Admoni, Scott Niekum, Reid Simmons, Aaron Steinfeld, and Tesca Fitzgerald. Understanding the relationship between interactions and outcomes in human-in-the-loop machine learning. In *Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.

[Fitzgerald *et al.*, 2017] Tesca Fitzgerald, Ashok Goel, and Andrea Thomaz. Human-robot co-creativity: Task transfer on a spectrum of similarity. In *International Conference on Computational Creativity (ICCC)*, 2017.

[Fitzgerald *et al.*, 2018] Tesca Fitzgerald, Ashok Goel, and Andrea Thomaz. Human-guided object mapping for task transfer. *ACM Trans. Hum.-Robot Interact.*, 7(2):17:1–17:24, October 2018.

[Fitzgerald *et al.*, 2019] Tesca Fitzgerald, Elaine Short, Ashok Goel, and Andrea Thomaz. Human-guided trajectory adaptation for tool transfer. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1350–1358. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

[Fitzgerald *et al.*, 2021a] Tesca Fitzgerald, Ashok Goel, and Andrea Thomaz. Abstraction in data-sparse task transfer. *Artificial Intelligence*, 300:103551, 2021.

[Fitzgerald *et al.*, 2021b] Tesca Fitzgerald, Ashok Goel, and Andrea Thomaz. Modeling and learning constraints for creative tool use. *Frontiers in Robotics and AI*, 8, 2021.

[Pastor *et al.*, 2009] Peter Pastor, Heiko Hoffmann, Tamim Asfour, and Stefan Schaal. Learning and generalization of motor skills by learning from demonstration. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 763–768. IEEE, 2009.

[Sadigh *et al.*, 2017] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. In *Robotics: Science and Systems*, 2017.

[Schaal, 2006] Stefan Schaal. Dynamic movement primitives-a framework for motor control in humans and humanoid robotics. In *Adaptive Motion of Animals and Machines*, pages 261–280. Springer, 2006.

[Scheutz *et al.*, 2017] Matthias Scheutz, Evan Krause, Brad Oosterveld, Tyler Frasca, and Robert Platt. Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017.

[Taylor and Stone, 2009] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *The Journal of Machine Learning Research*, 10:1633–1685, 2009.