Defining Interactions for Co-Creative Task Transfer

Tesca Fitzgerald[1]**, Ashok Goel**[1]**, Andrea Thomaz**[2]
School of Interactive Computing, Georgia Institute of Technology[1]
Department of Electrical and Computer Engineering, University of Texas at Austin[2]
{tesca.fitzgerald, goel}@cc.gatech.edu[1], athomaz@ece.utexas.edu[2]

Abstract

Embodied creativity introduces several challenges for computational creativity research. Particularly, a creative robot must be able to produce, express, and evaluate creative solutions through action and perception in the physical world. We discuss these challenges by examining the domain of task transfer for a robot which learns from interacting with a human teacher; in order for a robot to reuse task knowledge in a way that is novel to itself and the human teacher, it must utilize a creative process. We discuss a co-creative approach to task transfer: by continuing to interact with the human teacher in order to transfer task knowledge, the robot can address problems which require a creative solution. We introduce two modes of interaction for creative task transfer, and discuss how the robot's interactions with the teacher to request assistance will affect its success in task transfer.

Introduction

Creativity in robotics is often discussed in the context of a robot performing behaviors that typically require human creativity. Gemeinboeck & Saunders (2013) suggested that the embodiment of a robot lends it to be interpreted in the context of human behaviors. The robot's enactment in human environments creates meaning to the observer.

A robot can learn to reproduce actions that typically reflect human creativity; in robots that learn from interaction, a human may teach a robot to repeat a task by providing it with a demonstration (e.g. physically guiding the robot's hand to complete a task) (Argall et al. 2009; Chernova and Thomaz 2014; Akgun et al. 2012). However, a robot that learns to reproduce a demonstrated, creative task is not necessarily creative itself. Bird & Stokes (2006) propose two requirements of a creative robot: *autonomy* and *self-novelty*. According to these requirements, a creative robot's solutions are novel to itself, regardless of their novelty to a human observer. This represents an instance of "psychological creativity": the generation of ideas which are novel to the individual who produced them (Boden 1996).

This distinction from other problems of computational creativity is also evident in a robot that needs to transfer tasks learned in a familiar domain to novel domains. For example, if objects in the new domain (referred to as the *target* domain) are configured similarly to those in the original domain (the *source* domain), the robot may be able to repeat the learned task model in the target domain without producing novel actions. However, if new constraints are present in the target domain, the robot may need to produce behavior which was not taught by the teacher in order to reproduce the task; such processes may require creativity.

We propose the use of *human-robot co-creativity* to address difficult task transfer problems that require the robot to perform a novel behavior. Just as creativity is evident in collaboration between humans (e.g. collaborating to assemble a structure out of blocks), human-robot co-creativity involves the coordination of novel, physical actions to achieve a shared goal. In (Fitzgerald, Goel, and Thomaz 2017a), we have argued that a robot exhibits creativity by (i) reasoning over past task knowledge, and (ii) producing a new sequence of actions that is different from the taught behaviors. We have also argued that for sufficiently difficult task transfer problems (in which the robot must produce an action that is different than that originally taught), creativity is necessary for the robot to perform task transfer successfully. Finally, co-creativity occurs when the robot collaborates with the human teacher to perform task transfer, and is necessary in order to maintain autonomy while addressing a variety of transfer problems. In the rest of this paper, we discuss the function of human-robot co-creativity in task transfer, and propose two modes of interaction for co-creativity.

Grounding in embodied creativity

Creative transfer tasks

In (Fitzgerald, Goel, and Thomaz 2017a) we have argued that a robot which suitably addresses the problem of creative transfer meets three criteria:

- *Autonomy*: Rather than rely on receiving a new demonstration of the entire task, an autonomously creative robot must reason about the task using the representation it has previously learned, while also minimizing its reliance on the human teacher. This criteria does not preclude the robot from deriving new information from human interaction, provided that (i) the robot does not require a full re-demonstration of the task, and (ii) the robot reasons over what information is needed from the teacher and how to request that information. We refer to a robot that

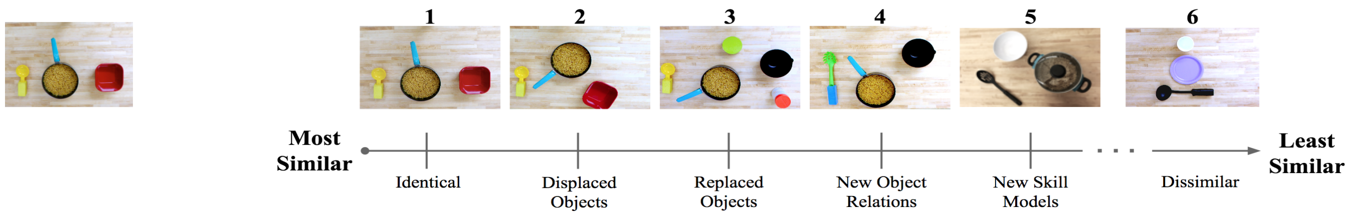

(a) Source

(b) Potential Target Environments

Figure 1: Spectrum of Similarity Between Source and Target Environments

meets these two criteria while collaborating with a human teacher as exhibiting *partial-autonomy*.

- *Novel output*: The robot learns to complete a task with respect to the locations of relevant objects (e.g. pouring is an action which is completed with respect to the location of a bowl and a scoop). By parameterizing the skill models (learned from the demonstration) based on object locations, simple adjustments can be made to objects' locations without altering the skill model itself. However, once a transfer problem requires significant changes to the skill model (either in constraints of the model, or a replacement of the model entirely), it no longer produces the same action. The revised model is reflective of a behavior that is both novel to the human teacher (since it is different than what was originally taught), and novel to the robot (since it is distinct from the output of other skill models the robot may have recorded).
- *Creative reasoning*: A robot may need to derive additional information about the task in the target environment. By interacting with a human teacher to request additional task information, the robot would leverage *co-creativity* in which the robot and human teacher collaborate to produce a novel result. As an alternate approach, a robot can address a target environment by combining aspects of its previous experiences. For example, a robot may know how to pour a mug, and separately, how to pick up a bowl. Knowledge of these two tasks may be combined in order to address a new problem, such as the robot needing to pour a bowl. By performing *conceptual blending* in this way, the robot would leverage a creative reasoning process.

Abstracting and grounding task representations

Not only do robots provide an interesting application for creativity, but they also introduce challenges that are specific to embodied creativity. A robot must represent task information in enough detail to capture the learned motion (represented using an action model) and how it relates to perceptual data (e.g. features of objects used to complete the task). In previous work (Fitzgerald, Goel, and Thomaz 2015; 2017b), we have defined the Tiered Task Abstraction (TTA) representation for tasks learned from demonstrations. The TTA representation contains the following elements:

- Skill Models: The task demonstration is segmented into *task steps*, each of which is represented by a separate *skill model*. These models are parameterized in terms of a start and end location, while maintaining the trajectory "shape" of the demonstrated action.
- Parameterization Functions: These reflect constraints which guide the start and end position of each task step as an offset from an object location. For example, scooping ends with the robot's end-effector 5 cm above the pasta bowl, before continuing with the next task step. The corresponding parameterization function is: $<o_x, o_y, o_z + 5>$, where o is a reference to the relevant object (in this case, the location of the pasta bowl).
- Object Labels: These are the labels which are uniquely associated with each object instance identified in the environment. Each labeled object represents a single object which is consistent over a range of feature values.
- Object Features: These are the feature values associated with each object label. While the label represents a static object, the specific feature values may differ depending on the environment, e.g. object locations, color (based on lighting conditions), spatial configurations, and properties.

Note 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, a task can be represented at a level of abstraction which is common to both the source and target environments. However, once a representation is abstracted, it must be *grounded* in the target environment in order to produce an output which is executable by the robot. A representation is *grounded* in a target environment when each of the TTA elements are present and defined based on information derived in the target environment (either by perception or interaction in the target environment). This challenge of abstraction and grounding is at the core of embodied creativity.

The source and target environments may differ according to the feature values of objects (e.g. dimensions, color, shape), object properties (e.g. a cup may be filled or empty), object locations, and/or spatial relations (e.g. the cup is to the left of the bowl). We have found that transfer problems can be analyzed based on the similarity between the source and target environments, and that this similarity can be expressed on a spectrum (Fig. 1), ranging from identical environments to highly dissimilar environments. As the source and target environments become more dissimilar, the task must be represented at increasing levels of abstraction for transfer to be successful (Fitzgerald, Goel, and Thomaz 2015; 2017b).

While some categories of task transfer (represented by discrete similarity levels indicated on this spectrum) do require a co-creative approach, task transfer does not inherently necessitate creativity. To address problems in which objects are displaced in the target environment (image 2 in Fig. 1b), the *object features* element must be grounded in the target environment, while other elements of the original representation can be retained. This grounding occurs by observing the new object locations in the target (Pastor et al. 2009; Fitzgerald, Goel, and Thomaz 2015; 2017b). In transferring a task to a target environment which requires an object mapping (image 3 in Fig. 1b), the robot must first obtain a *mapping* between objects in the source and target environments. With this mapping, the skill model can be re-parameterized according to the correct objects. Thus, the learned skill models are again reused, and so the resulting action is not novel to the robot or human teacher.

In contrast to these examples, consider target environments 4 and 5 in Figure 1b. Target 4 differs from the source in Figure 1a in that objects are: (i) displaced, (ii) replaced, and now (iii) *constrained* because of the new scoop size. The robot's actions must now be constrained such that its end-effector remains higher above the table in order to complete the task successfully. Accordingly, new *parameterization functions* must be identified in the target environment, applying constraints to the learned skill models that are distinct from those of the original demonstration. If a robot can identify the new parameterization functions with some degree of autonomy (e.g. does not simply receive a new demonstration of the task in the target environment), this category of transfer problems meets the criteria for creative transfer: partial-autonomy and novel output.

Target 5 in Figure 1b contains one additional difference: an extra step is needed in order to lift the lid off the pasta pot prior to scooping the pasta. As a result, the original skill models learned in the source cannot be directly transferred. In addition to deriving new parameterization functions in the target environment, this problem also requires that the robot derive or learn a new skill model to account for the missing step. We later discuss potential methods for deriving this information via further interaction with the human teacher. Regardless of what method is used, the robot would (i) autonomously transfer the task representation (since it does not rely on receiving a full re-demonstration of the task), (ii) produce action that is novel to both the robot and the human teacher, and (iii) utilize a creative reasoning method (by blending previously and newly learned skill models). Therefore, a robot that successfully completes transfer problems of this kind meets the criteria for creativity.

Figure 1 illustrates that without addressing problems of creative transfer, task transfer methods are limited to addressing a narrower set of transfer problems: those which do not require novel behavior or reasoning to address (targets 1-3 in Fig. 1b). By proposing human-robot co-creativity as a framework to address problems of creative transfer, we broaden the range of problems that a robot can address from transferring a single task demonstration.

Co-creative Task Grounding

To address problems in which objects are replaced in the target environment (e.g. target 3 in Fig. 1b), both the object features and object labels must be grounded in the target environment. We have demonstrated a method for grounding this information by inferring an *object mapping* from guided interaction with the human teacher (Fitzgerald et al. 2016). An object mapping indicates which objects in the source environment correspond to each object in the target environment, and is used to ground object labels in the target environment. By asking the teacher to assist in the object mapping by indicating the first object the robot should use in the target environment, the robot can attempt to infer the remainder of the object mapping. We have found that a robot can accurately infer an object mapping within the first 1-3 mapping-assistive interactions, and could then repeat the rest of the task autonomously (Fitzgerald et al. 2016).

To similarly abstract and ground the task representation in order to address problems of creative transfer, two elements of the TTA representation must be grounded in the target environment: the *parameterization functions* (for both categories of creative transfer problems) and *skill models* (for creative transfer problems involving new skill models). These two elements also contain the most high-level information about the task: the constraints between the robot's hand and objects in the environment, and the skill model which preserves the trajectory shape of the demonstrated action, respectively. Because these represent high-level information and are informed by the goal of the task, they cannot be grounded by the robot with complete autonomy. Presuming that the human teacher is aware of the goal of the task, and how that goal should be met in the target environment, we propose that the teacher is available to assist the robot in reaching that goal. The aim of this *co-creative* approach is enable the robot to (i) exhibit partial-autonomy, (ii) collaborate with the human teacher to infer information about the task in the target environment, (iii) produce parameterization functions and/or skill models that can ground an abstracted task representation, and (iv) produce and execute a trajectory in the target environment.

Grounding Parameterization Functions

In order to address problems in the New Object Relations category, it will be necessary to ground parameterization functions. The robot should interact with the teacher so that it infers the necessary information to ground missing elements of the task representation, without requiring too much information and time from the human teacher (so as to maximize the robot's autonomy).

We propose a method for grounding parameterization functions in a manner similar to object mapping. Rather than evaluate only the object mapping confidence at each step of the task, the robot should also verify its confidence in using the next step's parameterization function. One method of measuring confidence may be to compare the objects used in the next step to those which the robot would have used in the source environment. Assuming that similarly-shaped objects can be manipulated in similar ways, dissimilar objects may need to be manipulated differently (despite serving the same purpose). If the robot is not confident in this similarity (meaning its confidence value is below some threshold β), it can request the human teacher to align its end-effector in preparation to complete the next step of the task.

To do this, we propose that the robot asks the teacher "Please move my hand to the next object," or "Can you show me where to move my gripper?" The teacher would then provide assistance by moving the robot's gripper to the correct location, which the robot would record as the new offset between its gripper and the utilized object.

Grounding Action Models

To address tasks requiring new skill models (such as image 5 in Figure 1b), the robot will need to ground the same elements as before (object features, object labels, and parameterization functions) in addition to the new skill models. To do this, we hypothesize that the robot can again evaluate its confidence for completing each step of the task. We introduce an additional threshold to this evaluation process: if object similarity is below a second threshold $\alpha < \beta$, then the robot searches for other previously-learned task demonstrations which contain the unfamiliar object. Additionally, an unfamiliar object which does not share the same affordances as its counterpart in the original demonstration will need to be manipulated differently in order to achieve the same task goal. If there exists another demonstration of any task which also uses the new object, the robot should then evaluate the similarity between (i) the task step involving the object in the original source environment and (ii) the task step in the newly-retrieved demonstration that involves the new object. If the two task steps are similar, then the newly-retrieved step may be applicable toward the task in the target environment.

If they are not similar, the robot may ask the teacher to re-demonstrate the task step, e.g. "Can you show me the next step of the task?" The teacher would respond by moving the robot's hand to complete the task step, during which the robot would record the trajectory of its gripper. This form of assistance would require the most involvement from the teacher, but also provide the robot with the most detailed information. A challenge of this approach will be clearly indicating that the teacher should demonstrate a full task step, rather than simply reposition the robot's hand.

Conclusions and Future Work

In (Fitzgerald, Goel, and Thomaz 2015; 2017b), we identified a relation between (1) source and target similarity and (2) the level of abstraction at which the task should be transferred. A result of this relation is that more dissimilar transfer tasks require more information about the target environment in order to ground the task abstraction in that environment. We now view this problem through a lens of co-creation; how should the robot interact with a human teacher in order to derive the information necessary to ground its representation? Furthermore, how should it request this assistance from the human teacher?

In accordance to our findings in (Fitzgerald, Goel, and Thomaz 2015; 2017b), we hypothesize that transfer problems can be addressed once the appropriate amount of assistance is provided, based on the level of abstraction at which the task is represented and transferred. We also hypothesize that the extent of a teacher's assistance varies according to (1) how the robot requests assistance and (2) the type of interaction enabled. Additionally, we expect that there is a correlation between (1) the amount of assistance provided, and (2) the effort by the teacher in order to provide that assistance. As such, we expect that assistance requests should be selected based on the particular transfer problem. Our next steps will be to evaluate these interaction methods through a user study, analyzing the information gained via each assistance method and the types of transfer problems which they enable. By evaluating its own confidence in its representation of the transfer task, the robot may be able to select the assistance type which provides the necessary information for the task, while also maximizing its own autonomy.

Acknowledgments

This work was supported by the NSF Graduate Research Fellowship under Grant No. DGE-1148903.

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.

Bird, J., and Stokes, D. 2006. Evolving minimally creative robots. In *Proceedings of the Third Joint Workshop on Computational Creativity*, 1–5. IOS Press, Amsterdam.

Boden, M. A. 1996. *Dimensions of creativity*. MIT Press.

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

Fitzgerald, T.; Bullard, K.; Thomaz, A.; and Goel, A. 2016. Situated mapping for transfer learning. In *Fourth Annual Conference on Advances in Cognitive Systems*.

Fitzgerald, T.; Goel, A.; and Thomaz, A. 2015. A similarity-based approach to skill transfer. In *Workshop on Women in Robotics at Robotics: Science and Systems*.

Fitzgerald, T.; Goel, A.; and Thomaz, A. 2017a. Human-robot co-creativity: Task transfer on a spectrum of similarity. In *Eighth International Conference on Computational Creativity*.

Fitzgerald, T.; Goel, A.; and Thomaz, A. 2017b. Similarity and abstraction in task transfer. In *submission*.

Gemeinboeck, P., and Saunders, R. 2013. Creative machine performance: Computational creativity and robotic art. In *Proceedings of the 4th International Conference on Computational Creativity*, 215–219.

Pastor, P.; Hoffmann, H.; Asfour, T.; and Schaal, S. 2009. Learning and generalization of motor skills by learning from demonstration. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, 763–768. IEEE.