

Automatically Composing Representation Transformations as a Means for Generalization

Michael B. Chang¹
mbchang@berkeley.edu

Abhishek Gupta¹
abhigupta@berkeley.edu

Sergey Levine¹
svlevine@eecs.berkeley.edu


Tom Griffiths²
tomg@princeton.edu



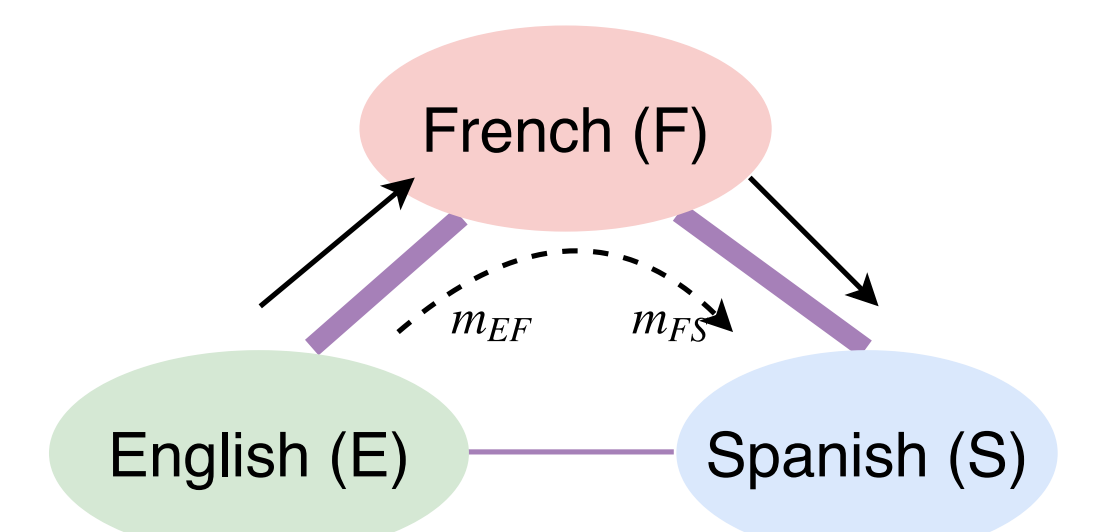
Compositional Generalization Problem

How can we build learners that automatically leverage prior experience to solve more complex problems than they have previously encountered?


Standard supervised learning assumes train and test data are drawn i.i.d. from the same distribution



Compositional Generalization Problem: We want the learner to translate English-Spanish even though it has never seen these translation pairs during training

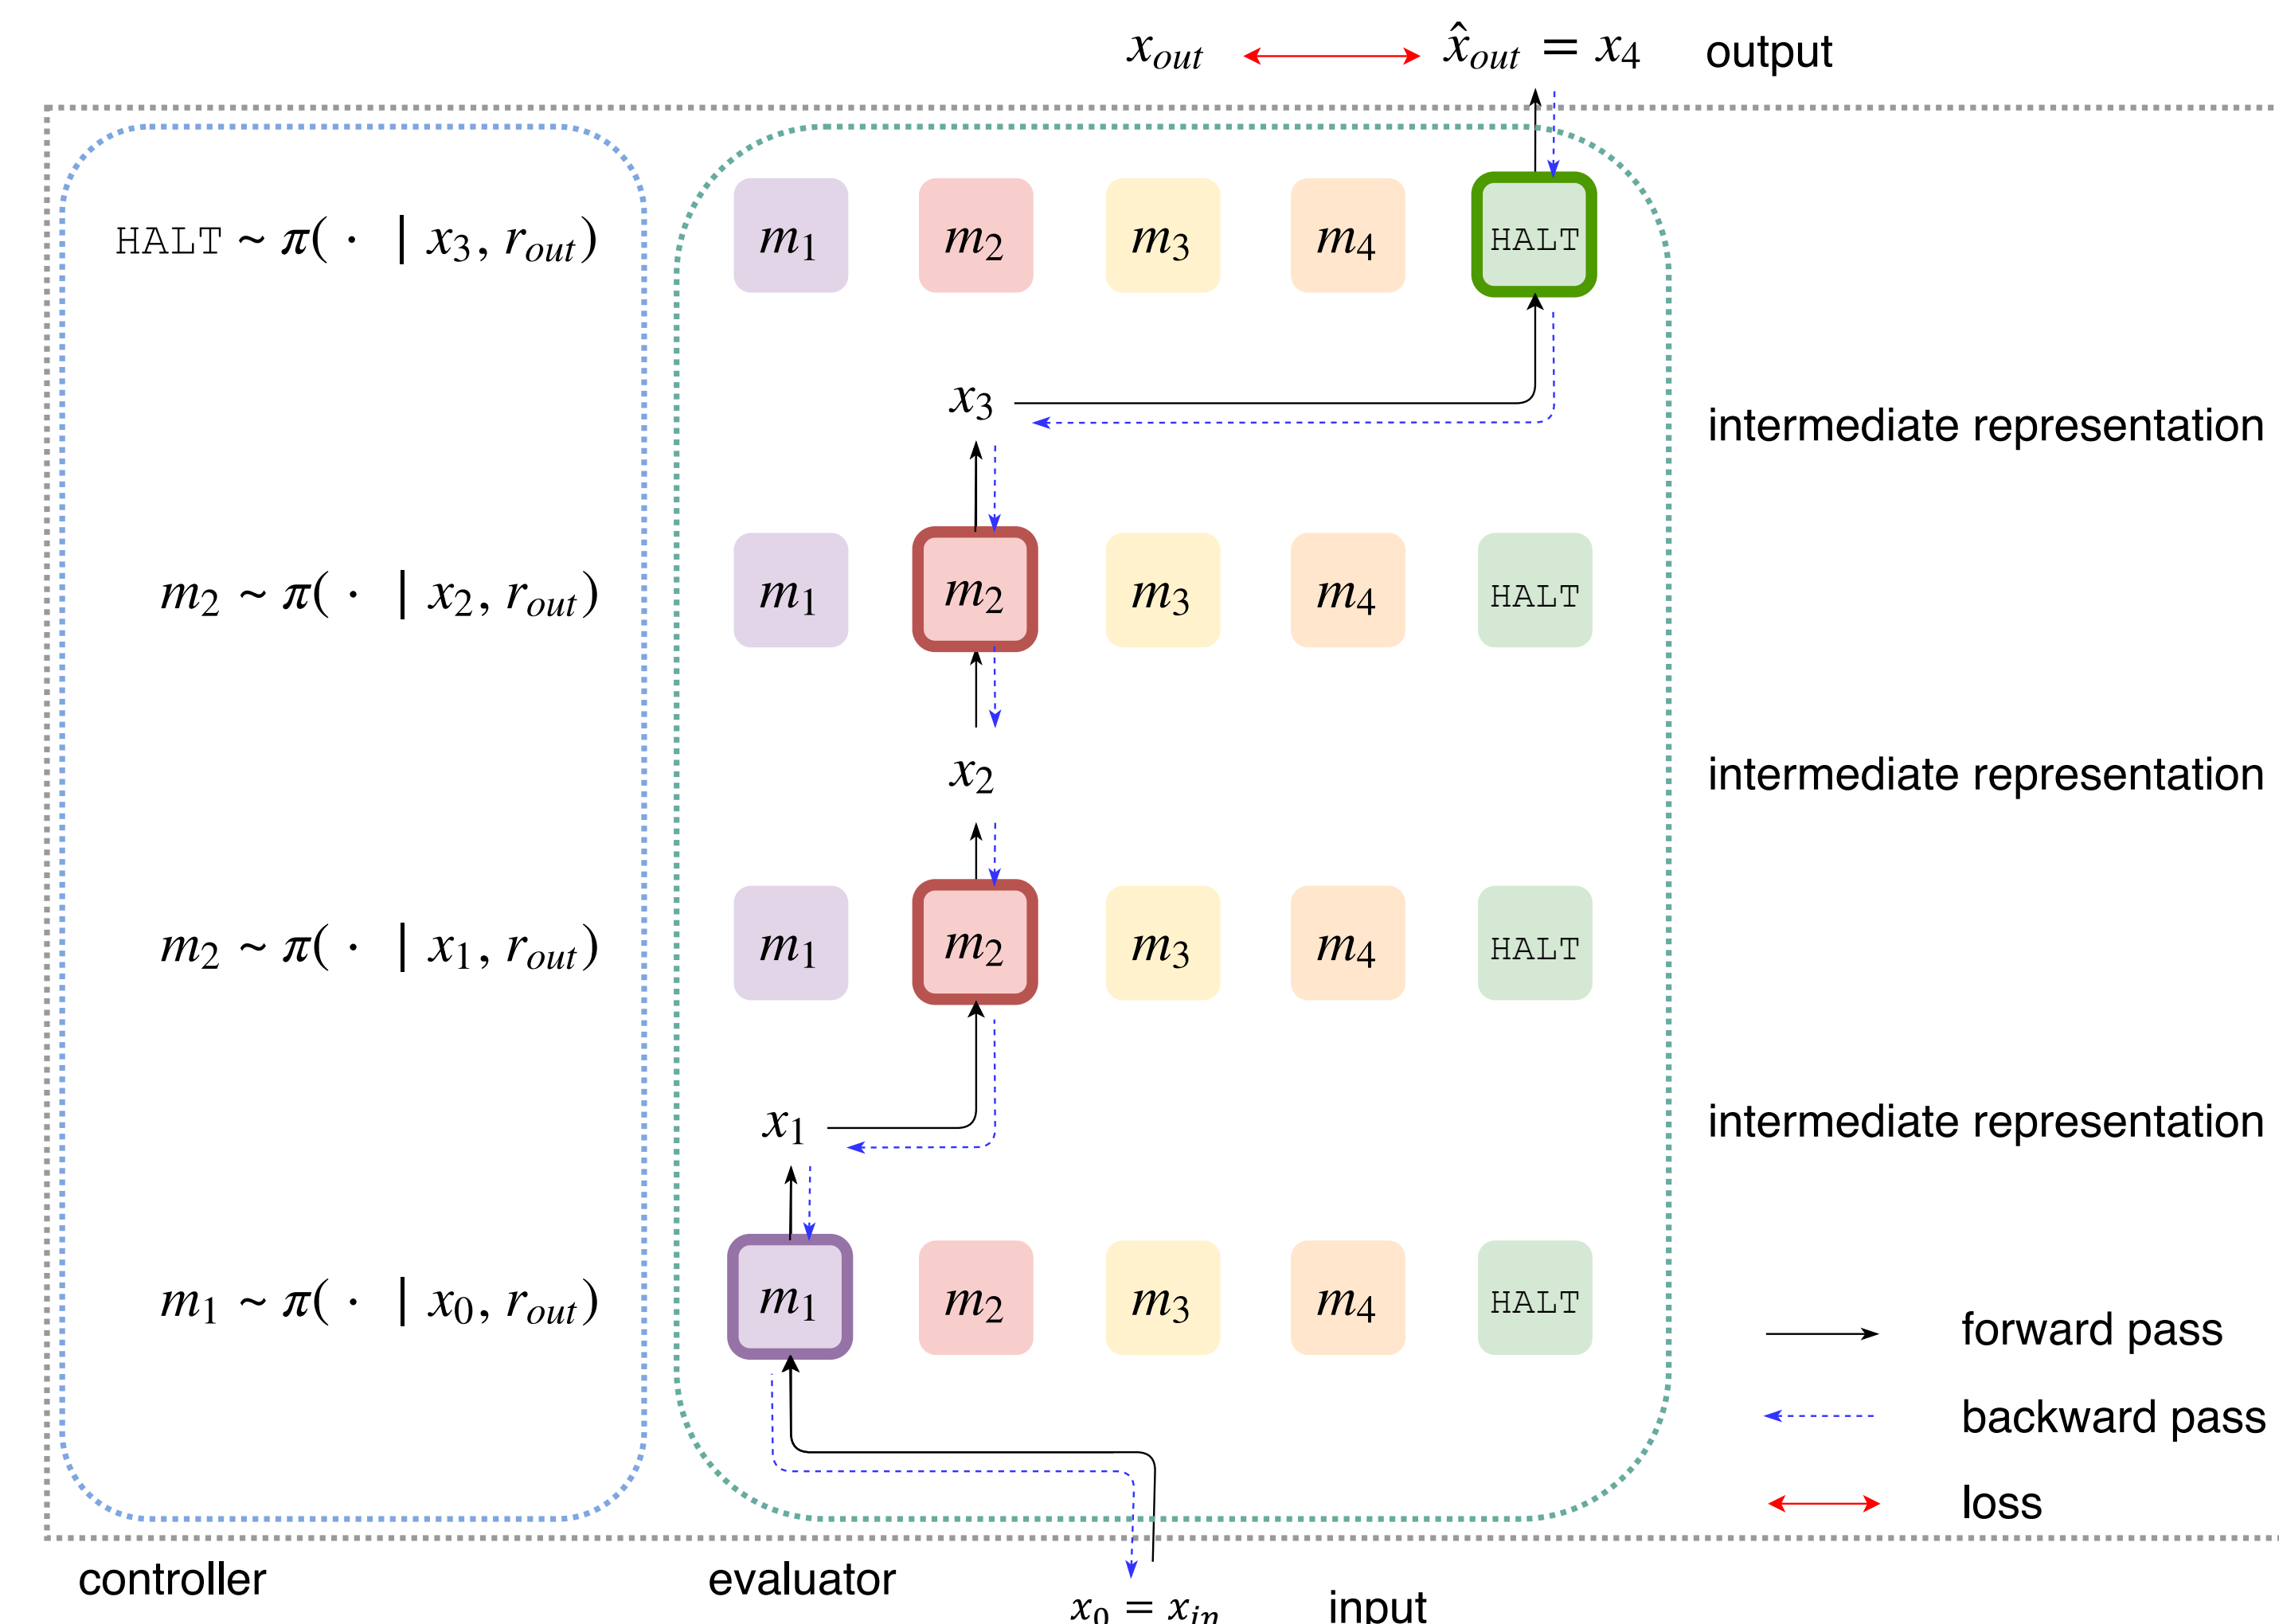


Consider a family of problems whose subproblems are shared within and across problems:



Compositional Recursive Learner

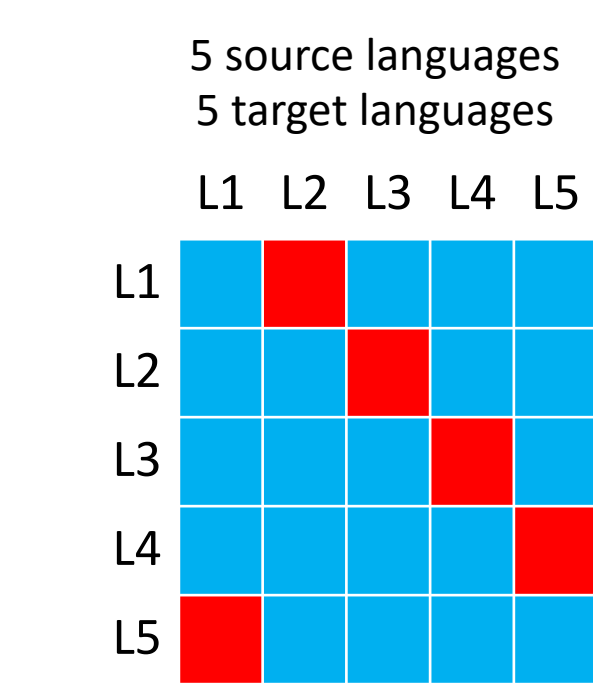
The **Compositional Recursive Learner (CRL)** is a **self-organizing learner** that learns the parameters of primitive transformations as well as their means of composition.




Multilingual Arithmetic Extrapolation

Task: Evaluate an modular arithmetic expression given in the source language and return the value in the target language

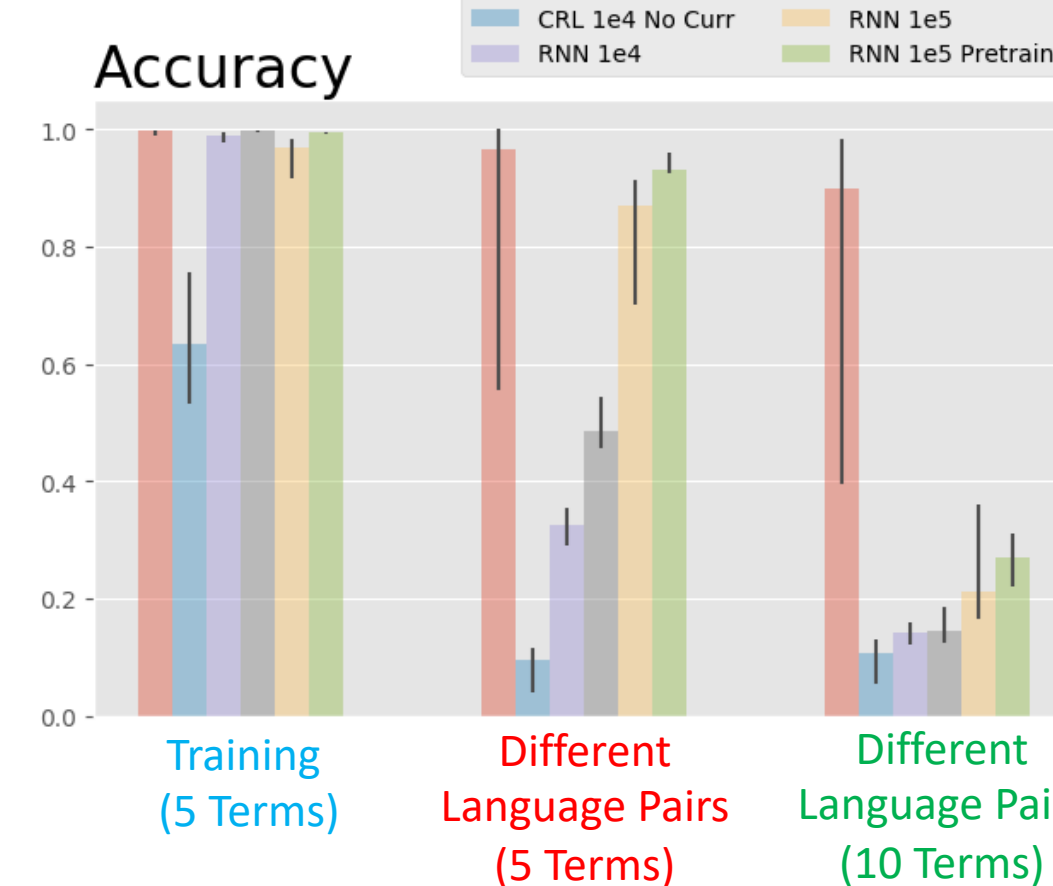
Goal: Extrapolate to **different language pairs** and **longer sequences** of different language pairs compared to training



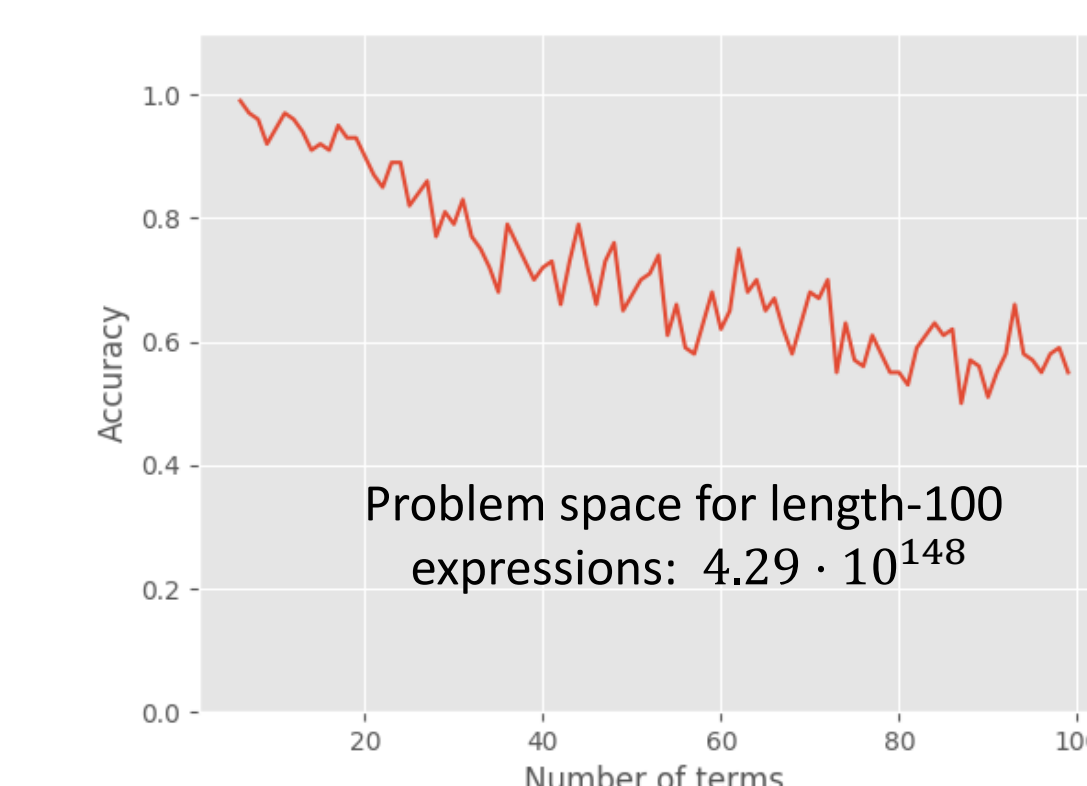
Example Execution Trace




CRL extrapolates better to **new language pairs** and **longer expressions**




How far can we push extrapolation?



Spontaneous Specialization of Reusable Primitive Transformations





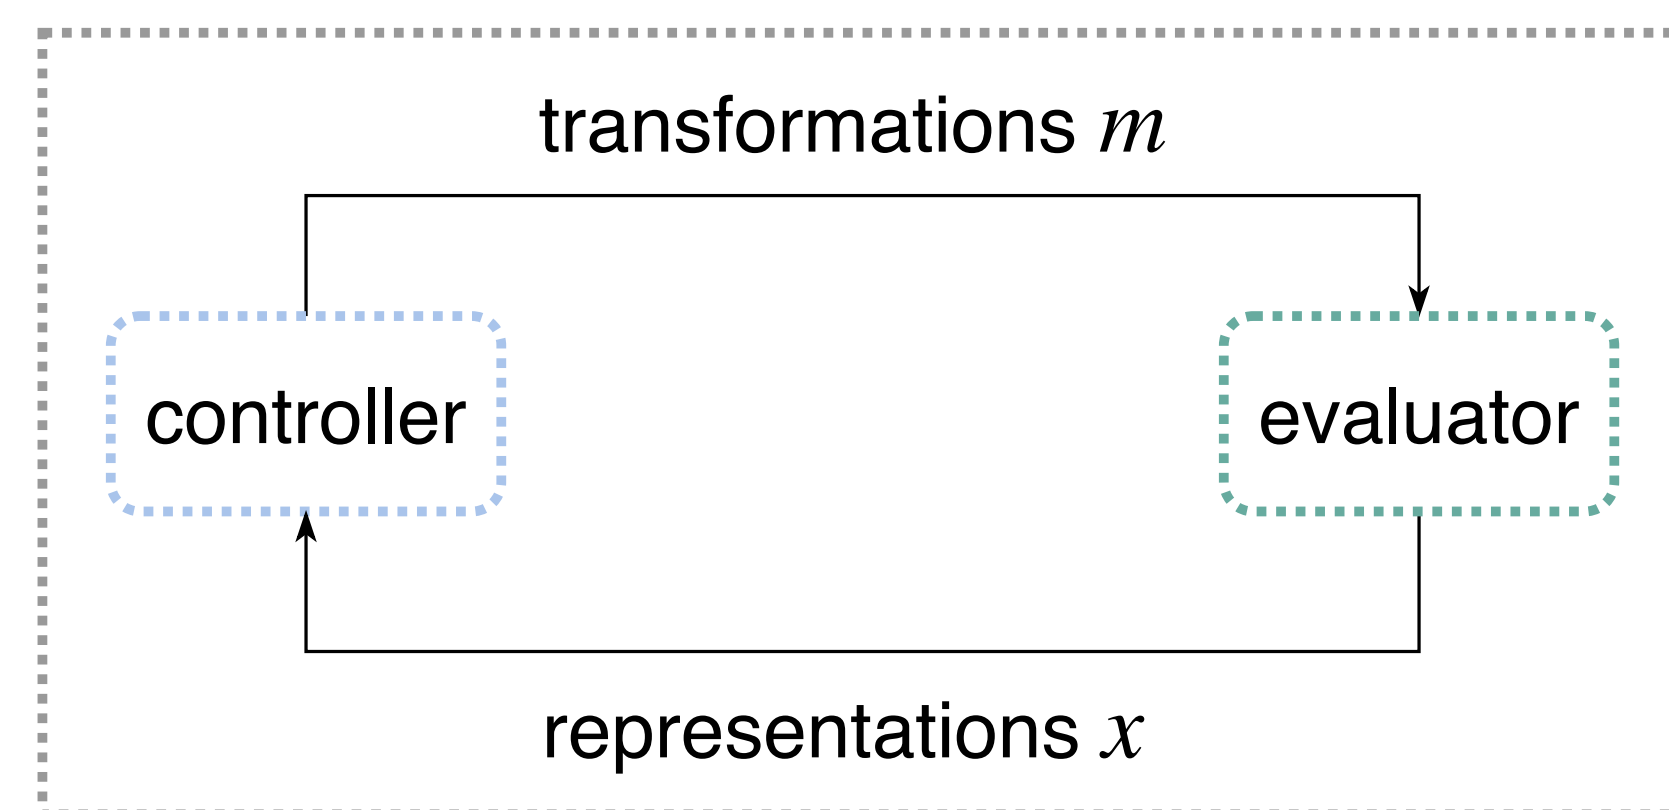
CRL re-represents new problems in terms of old problems, thereby taking more computation steps



Composing Representation Transformations

Key Idea

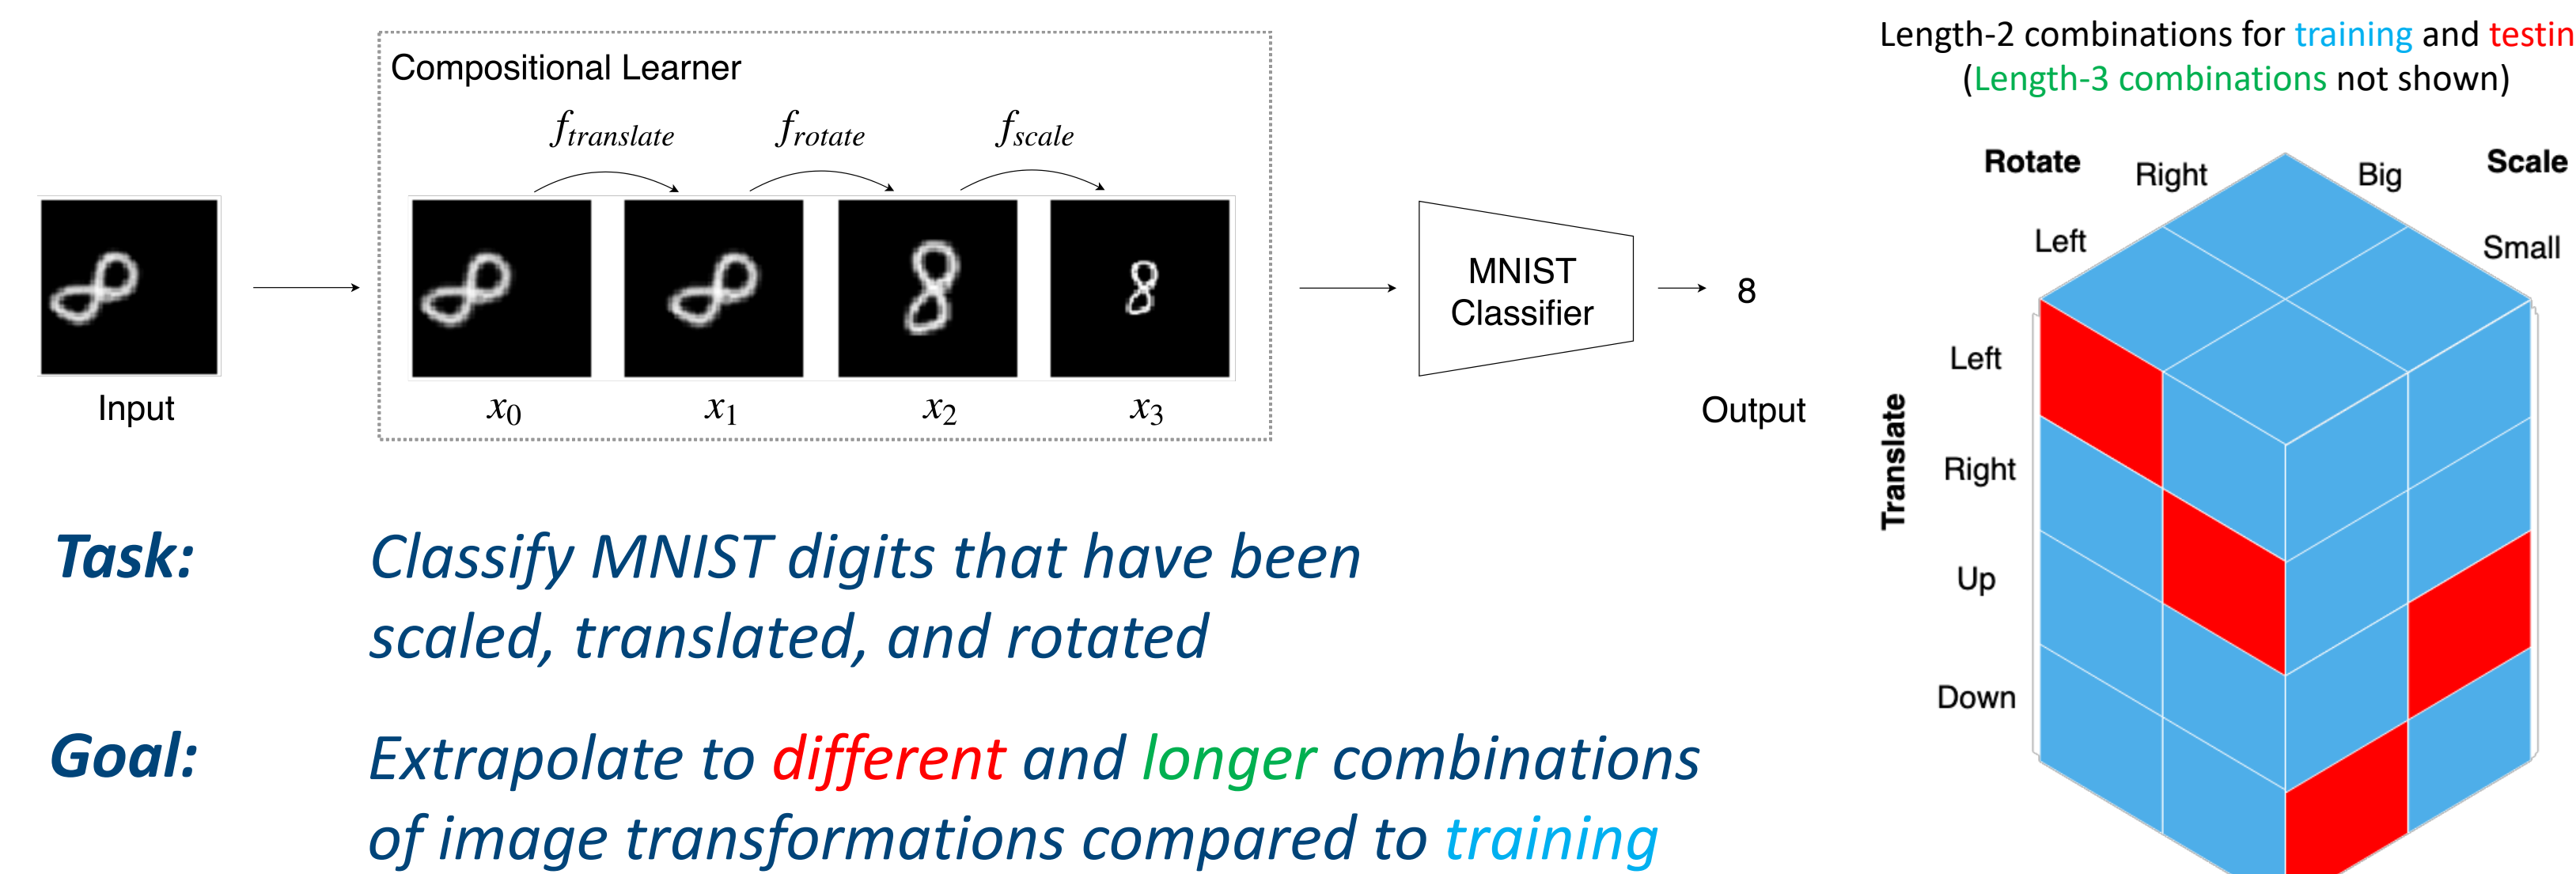
By exploiting the compositional structure of the task distribution, we can learn algorithmic procedures over **representation transformations** as a means for extrapolation in compositional problems



Standard approaches for learning multiple tasks either train a separate learner per task or train a single learner for all tasks. **Both have difficulty generalizing for compositional problems.**

The **compositional recursive learner** explicitly **re-uses previously learned sub-solutions** to solve new problems by **composing computational modules in new ways.**

Classification Extrapolation



Task: Classify MNIST digits that have been scaled, translated, and rotated

Goal: Extrapolate to **different and longer combinations** of image transformations compared to training

