Compositional Constraint Satisfaction Control

Submission ID 3000157

Submission Type Poster

Topic Artificial Intelligence

Status Submitted

Submitter Thomas Ringstrom

Affiliation University of Minnesota

SUBMISSION DETAILS

Presentation Type Either Poster or Oral Presentation

Presentation Abstract Summary Animals need to solve a multitude of different problems to protect themselves from environmental threats and intrinsic physiological instability. The traditional modeling approach to understanding animal behavior in Reinforcement Learning is to encode the rewards sought by the animal as a vector with cardinal values corresponding to desirable or threatening states. We call into question the feasibility of this strategy due to the inflexible coupling between the reward vectors and non-compositional policies. The problem with current RL approaches is that the reward function needs to be engineered for specific tasks, and the policy is tied to a Q-function that is dependent on the reward function. Without a flexible representation of reward, we cannot design artificial agents that have the capacity to respond to new environmental constraints and opportunities; without flexible compositional policies, we cannot design agents that can do this efficiently. We define a reward function that has the semantics of probabilistic constraint satisfaction and apply it to the domain of compositional control algorithms. This common currency is a flexible representation that allows us to map mixtures of constraints onto mixtures of policies, allowing the agent to reuse policies in new situations, and pool preferences for options across logical possibilities.

Co-author Information

* Presenting Author

First Name	Last Name	Affiliation	E-mail
Thomas *	Ringstrom *	University of Minnesota	rings034@gmail.com
Paul	Schrater	University of Minnesota	schrater@umn.edu

Keywords

Keywords
Control Theory

Constraint Satisfaction
Intrinsic Reward Functions
Compositionality
Linearly Solvable MDPs