Hierarchical Representations of Behavior for Efficient Creative Search

Christopher Vigorito and Andrew Barto
Department of Computer Science, University of Massachusetts Amherst
{vigorito, barto}@cs.umass.edu

Creative Search

- We adopt a Darwinian view of creativity, in which creative products are assembled through a sequential decision process in a trial-and-error fashion.
- This search process consists of trajectories through the space of creative products and selective retention of highly-valued results (and the means for constructing them).
- An explict population is not maintained, however. Rather, an implict "population" of potential creative products exists at a given time, any member of which can be realized via some behavioral operator.
- One "member" of each "population" is thus selected at each step given the current product (context), and a new "population" results at the next step.

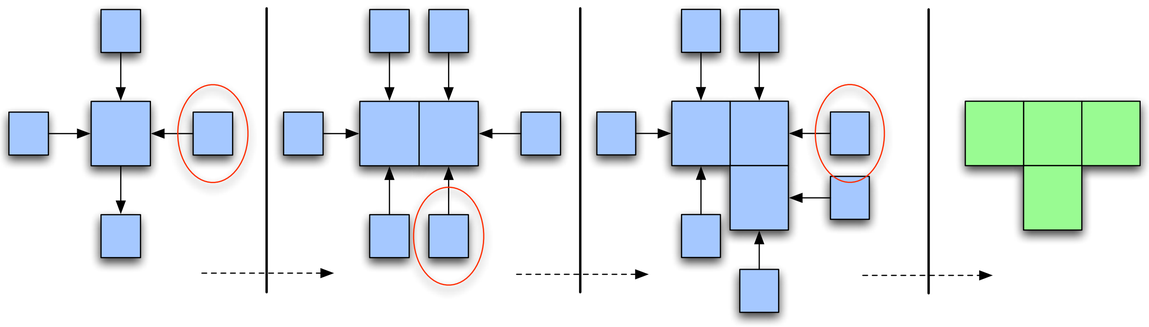

- Although this space can be very large, abstract search operators can effectively reduce its size by affording large (purposeful) jumps through it.
- A hierarchy of these operators allows for navigation through the space at varying levels of abstraction, allowing for facilitated reachability of any creative product.
- Blind variation and selection is only efficient in such large spaces if these macro-operators can be discovered, retained, and employed properly.
- Constructing models of the effects of these operators on the search space allows for efficient planning, facilitating sophisticated variation.

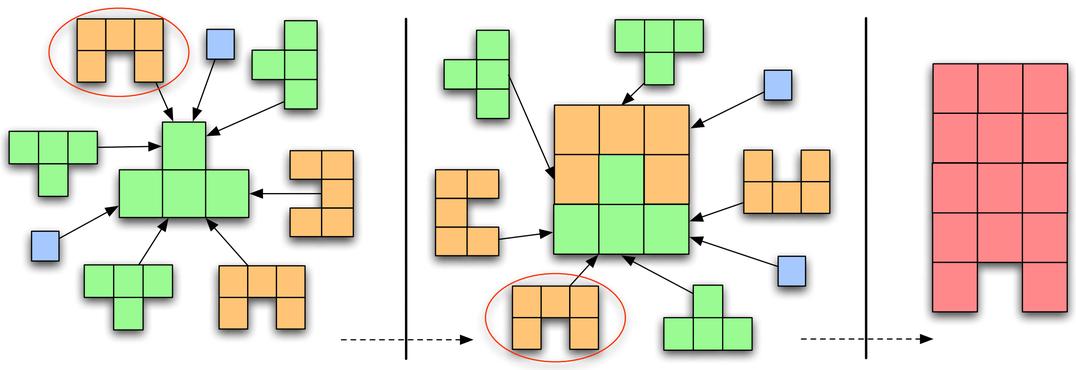

Formalism

- We test our hypotheses in a reinforcement learning framework.
- Macro-operators in this framework correspond to options, which are closed loop behavior policies.
- We assume the environment can be modeled as a Factored Markov Decision Process (FMDP), which allows for exploitation of environmental structure.
- Options (their policies and models) are computed from structural information in the FMDP and used as a behavioral basis for sophisticated variation.
- Given a selection metric, new options can be formed as new structure is discovered and desirable products are generated.
- Option models allow for an agent to treat complex behaviors as primitive actions when predicting action effects, facilitating planning.

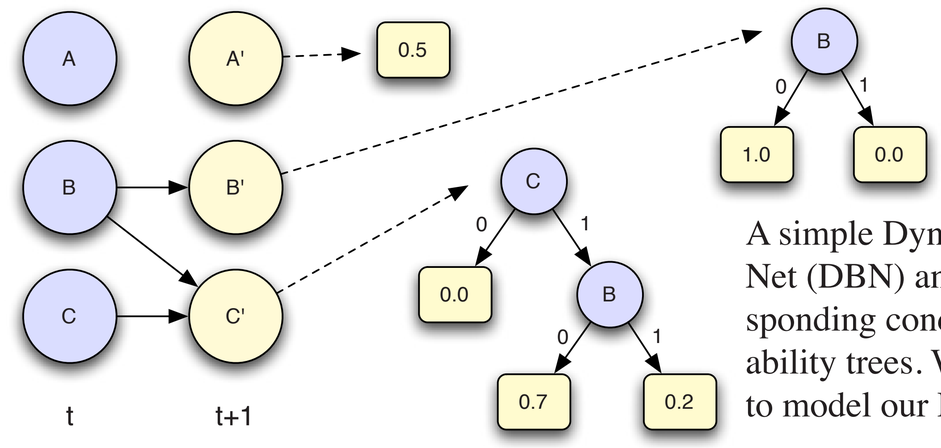

A simple Dynamic Bayes Net (DBN) and its corresponding conditional probability trees. We use DBNs to model our FMDPs.

Experiments

- We conducted experiments in a simple stochastic domain consisting of 20 binary variables (approximately 1 million states).
- Each variable is also a primitive action.
- Blue lights can be toggled primitively.
- Other lights have dependencies that must be satisfied before they can be toggled.
- Variables influence each other in a hierarchical fashion (blue -> red -> green -> yellow).
- We ran two agents, one with primitives only and one with primitives + options computed from a given structural model.
- We then looked at the frecquency of change of each variable under a uniformly random policy.

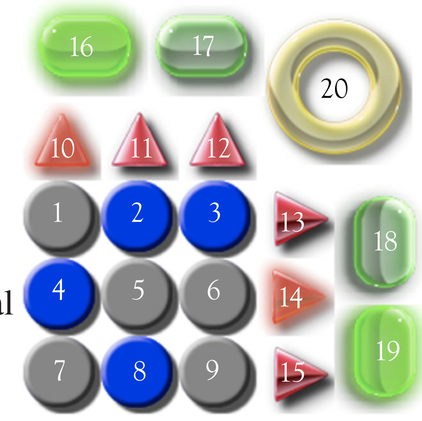

The Light Box Domain

Results

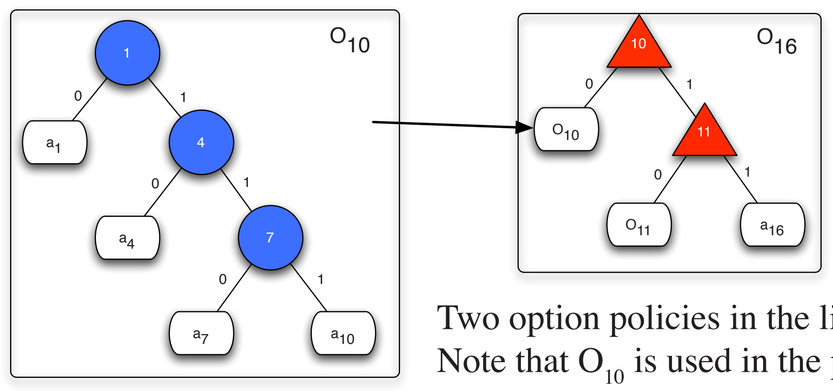

Two option policies in the light box domain. Note that O_{10} is used in the policy of O_{16}.

- Both agents alter the blue lights as frequently as each other.
- The agent with options is able to change every variable equally well.
- The agent with only primitives fails to change higher numbered variables frequently, since often their dependencies are not satisfied.
- Options clearly allow for more sophisticated variation, even when the variation is blind.

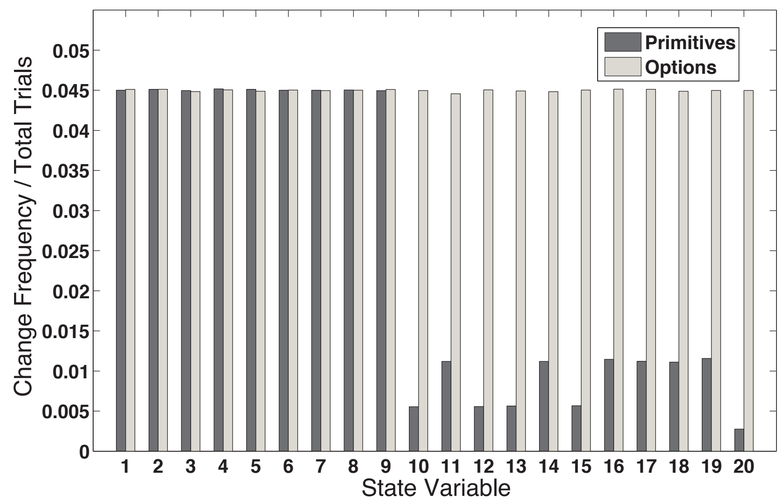

Conclusions and Future Work

- Blind variation with options can seem purposeful because of the large jumps in the creative search space options afford.
- This is a hallmark of evolution, that blind variation and selective retention of highly-valued products result in increasingly complex structures.
- The same increasing complexity in behavior is observed in our approach, without the need to maintain explicit populations of creative products.
- Agents with hierarchies of options are capable of highly sophisticated behavioral variation and are thus more adept at complex creative behavior than agents without options.
- Current work consists of incremental methods for structure discovery and motivational systems for maximizing the rate of this discovery.
- Future directions include incoporation of perceptual abstraction in addition to the behavioral abstraction afforded by options.

The work presented here was supported by the National Science Foundation under Grant No. IIS-0733581. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.