

Metareasoning for Adaptation of Classification Knowledge (Extended Abstract)

Joshua Jones
Georgia Institute of Technology
Technology Square Research Building
85 5th Street NE
Atlanta, GA 30332
jkj@cc.gatech.edu

Ashok Goel
Georgia Institute of Technology
Technology Square Research Building
85 5th Street NE
Atlanta, GA 30332
goel@cc.gatech.edu

Categories and Subject Descriptors

I.2.1 [Artificial Intelligence]: Applications and Expert Systems—*Games*; I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods; I.2.6 [Artificial Intelligence]: Learning

General Terms

Design, Reliability

Keywords

Classification, Hierarchical Classification, Metareasoning

1. INTRODUCTION

AI research on metareasoning for agent self-adaptation has generally focused on modifying the agent's reasoning processes (e.g., [2]). In this paper, we describe our ongoing work on the equally important problem of using metareasoning for modifying the agent's domain knowledge. Since classification is a ubiquitous task in AI, we consider the problem of using meta-knowledge for repairing classification knowledge when the classifier supplies an incorrect class label. More specifically, we consider the subclass of classification problems that can be decomposed into a hierarchical set of smaller classification problems; alternatively, problems in which features describing the world are progressively aggregated and abstracted into higher-level abstractions until a class label is produced at the root node. This subclass of classification problems is recognized as capturing a common pattern of classification [1] [3]. We will call this classification task *compositional classification*, and the hierarchy of abstractions an *abstraction network* (AN).

In particular, we consider the problem of retrospective adaptation of the content of the intermediate abstractions in the abstraction network (and *not* its structure) when the classifier makes an incorrect classification. Note that structural credit assignment is a core problem in making this adaptation: given the error at the root node, the structural credit assignment problem is then to identify the intermediate abstractions in the abstraction network responsible for the error. Here we propose and explore the following hypothesis for using metareasoning for self-adaptation of domain knowledge: if the semantics of domain concepts that form the intermediate abstractions in a classification hierarchy can be grounded in predictions about percepts in the world, then meta-knowledge in the form

Cite as: Metareasoning for Adaptation of Classification Knowledge, (Extended Abstract), Joshua Jones, Ashok Goel, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1145–1146
Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

of verification procedures associated with those domain concepts is useful for addressing the structural credit assignment problem. The verification procedures explicitly encode the grounding of intermediate abstractions in percepts from the environment.

2. ABSTRACTION NETWORKS

To make the problem concrete, we will present an example from the turn-based strategy game called FreeCiv (www.freeciv.org). On each turn in a game of FreeCiv, a player must select a compound action that consists of setting various parameters and moving units such as military units and "worker" units, called settlers, that can improve terrain or build new cities. Building new cities on the game map is a crucial action, as each city produces resources on subsequent turns that can then be used by the player to further advance their civilization. The quantity of resources produced by a city on each turn is based on various factors, including the terrain and special resources surrounding the city's location on the map, and the skill with which the city's operations are managed.

We have written an automated agent that plays FreeCiv. When our agent selects the action for a unit that is to build a city, a crucial decision is whether the location on the game map currently occupied by the unit is suitable for the placement of the new city. We will judge the quality of a potential city location based upon the quantity of resources that we expect a city built in that location to produce over time. This decision is an example of a compositional classification task. Figure 1 illustrates a knowledge hierarchy for this task used by our FreeCiv game-playing agent.

To more formally describe compositional classification, let T be a discrete random variable representing the class label. Let $S = \{s : s \text{ is empirically determinable and } h[T] > h[T|s]\}$, where $h[x]$ denotes the entropy of x . S is a set of discrete random variables that have nonzero mutual information with the class label and are *empirically determinable*, meaning that there is some way to interact with the environment to determine which value has been taken by each member of S , though in general this interaction will not be possible until some time after the classification has been produced. For example, in the classifier depicted in Figure 1, we can predict the population growth of a city as part of the classification of a map location, but we cannot observe the actual population growth that occurs at a city until some time later, after the city has been built. Empirical determinability captures the notion of perceptual grounding of concepts, indicating that each equivalence class represents some verifiable statement about the world. We call the problem of predicting T in such a setting *compositional classification*.

The knowledge structure we use for compositional classification contains a node for each $s \in S \cup T$. These nodes are connected

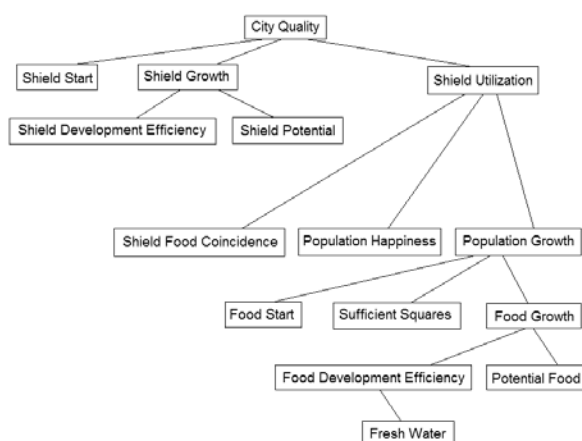


Figure 1: A game playing agent's abstraction network for the classification of city quality in FreeCiv.

in a hierarchy reflecting direct dependence relationships organized according to background knowledge. Each node will handle the subproblem of predicting the value of the variable with which it is associated given the values of its children. For this purpose, each node has a supervised classification learner associated with it. We have experimented with Artificial Neural Networks, k-Nearest Neighbor learners and rote learners within AN nodes. The experimental results described here use rote learners within nodes. Each node also has an *Empirical Verification Procedure* (EVP) associated with it. An EVP is an arbitrary, possibly branching sequence of actions in the environment and observations from the environment concluding with the identification of the correct output class that should have been returned by the relevant node during the last classification attempt. There is a strong notion of the predictivity of classification in AN learning; we expect that both the top-level class and intermediate classes translate into some expectations of the agent about the world. EVPs encode the means by which the agent can, at some future time after classification, reflect upon its knowledge given what has been observed in the world and alter knowledge that led to violated expectations.

In order to produce a classification, the values of the AN leaf nodes are first fixed by observation. Each node with fixed inputs then produces its prediction. This is repeated until the value of the class label is predicted by the root of the hierarchy.

At some time after classification, the true value of the class label is obtained by the monitoring process. If the value produced by object-level reasoning was correct, the agent's expectations are met and no further action is taken. If the value is found to be incorrect, a self-diagnosis and repair procedure is followed. The specifics of this procedure are dependent upon the characteristics of the learner types that are used within nodes and the classification problem setting. A very simple diagnostic procedure that always works, though it may not be maximally cost-effective in all settings, is to simply invoke all EVPs and potentially perform learning at all AN nodes during each adaptation.

Our work on use of metareasoning for structural credit assignment in compositional classification is related to past work on tree-structured bias (TSB) [3][4]. In TSB, a concept hierarchy like those represented by ANs is used to limit the hypothesis space that must be searched by a learner. However, TSB has dealt only with binary classifications at all nodes in the hierarchy, while ANs can deal with multivalued classifications. More importantly, TSB research does

not have the concept of EVPs, which encode the meta-knowledge used in our self-diagnostic procedure, instead relying on carefully constructed queries to the environment to learn the functions at internal nodes. Thus, rather than using explicitly represented meta-knowledge to perform self-diagnosis, TSB has a fixed training procedure that implicitly relies upon a given type of query.

3. FREECIV EXPERIMENT

In order to verify that the diagnostic technique described above allows for correction of faulty knowledge in an AN, we have experimented in the FreeCiv domain. We used the AN depicted in Figure 1, with table-based rote learners within each node. This AN was used to produce outputs from a set containing three values, corresponding to predictions of poor, moderate and good resource production for a city built on a considered map location. Specifically, the values correspond to an expected degree and direction of deviation from a logarithmic baseline resource production function that was manually tuned to reflect roughly average city resource production. The empirical verification procedures simply discretize observed game features. The content of each rote learner was initialized to zeros, which was known to be incorrect in some cases for each of the learners. In each trial, a sequence of games is run, and learning and evaluation occurs on-line. The learners are trained on sequences of 49 games. We segment these sequences of games into multi-game blocks for the purpose of evaluation; the block sized used is 7 games. Each game played used a separate randomly generated map, with no opponents. The agent always builds a city on the first occupied square, after making an estimate of the square's quality. We observed a 52% decrease in the error rate of the learner, averaged over 60 independent trial sequences, when comparing the first block of examples to the 7th block. This result is positive in that it demonstrates that the self-diagnosis and repair procedure is successful in correcting faulty classification knowledge.

4. CONCLUSIONS

In this paper, we considered retrospective adaptation of the content of intermediate abstractions in an abstraction network used for compositional classification. Retrospective adaptation is triggered when the classifier makes an incorrect classification. We showed that if the intermediate abstractions in the abstraction network are organized such that each abstraction corresponds to a prediction about a percept in the world, then meta-knowledge in the form of verification procedures associated with the abstractions can be used by introspective metareasoning to perform structural credit assignment and then adapt the abstractions.

Acknowledgements

This research is supported by an NSF (SoD) Grant (#0613744) on Teleological Reasoning in Adaptive Software Design.

5. REFERENCES

- [1] T. Bylander, T. Johnson, and A. Goel. Structured matching: a task-specific technique for making decisions. *Knowledge Acquisition*, 3:1–20, 1991.
- [2] J. W. Murdock and A. K. Goel. Meta-case-based reasoning: self-improvement through self-understanding. *J. Exp. Theor. Artif. Intell.*, 20(1):1–36, 2008.
- [3] S. J. Russell. Tree-structured bias. In *AAAI*, pages 641–645, 1988.
- [4] P. Tadepalli and S. J. Russell. Learning from examples and membership queries with structured determinations. In *Machine Learning*, volume 32, pages 245–295, 1998.