Transfer as a Benchmark for Multi-Representational Architectures

Matthew Klenk

Navy Center for Applied Research in Artificial Intelligence
Naval Research Laboratory (Code 5515)
Washington, DC 20375, United States
matthew.klenk.ctr@nrl.navy.mil

Introduction

We argue that transfer of spatial and conceptual knowledge between tasks and domains is an essential benchmark for multi-representational architectures aimed at human-level intelligence. The underlying hypothesis is that spatial relationships provide a natural level of abstraction, highlighting the similarities and differences between situations and domains. Therefore, not only will spatial representations improve domain reasoning and learning, they will also facilitate the transfer of knowledge across domains.

The simulated environments of real-time strategy (RTS) games provide an excellent test-bed for exploring this hypothesis for two reasons: many different RTS domains have been constructed and RTS requires a wide range of reasoning tasks.

We begin by discussing why transfer is an important benchmark, followed by a discussion of RTS games with a couple of illustrative examples.

Why Transfer?

Transfer learning research is motivated by the observation that people improve in their ability to learn new domains based on their experiences in related domains. Faced with a new domain (the *target*), the agent must identify a known related domain (the *source*), determine what similarities exist between the source and target, and transfer knowledge from the source to improve its performance in the target domain.[1] Transfer is especially important for multi-representational architectures focused on human-level performance because it is a hallmark of human reasoning.

Previous research on combining spatial and conceptual reasoning has focused on three primary operations: proposition extraction, reasoning with extracted propositions, and proposition projection (Chandrasekaran 1997). We claim that each of these operations play an important role in transfer. Proposition extraction is useful for identifying similarities between situations and domains. Reasoning over these similarities enables the transfer of knowledge, spatial and conceptual, from the source to the target. One result of this transfer is the projection of spatial entities and relationships into the target.

Few spatial reasoning projects have addressed issues of transfer. A couple of exceptions are Bi-Soar, which employs a diagrammatic reasoning system to perform way finding and memory recall tasks (Kurup 2007), and Companions, which uses analogical model formulation to transfer the spatial and conceptual knowledge necessary to solve mechanical comprehension problems from examples (Klenk *et al.* 2005). While these systems integrated spatial and symbolic representations to perform a variety of tasks, they both fall far short the benchmarks presented below.

Transfer in Real-Time Strategy Games

In recent years, video games have begun to receive more attention from AI researchers (e.g., Laird & van Lent 2000). Real-time strategy simulations provide an excellent testbed for investigating transfer of integrated spatial reasoning for two reasons: 1) Playing RTS games successfully requires a variety of tasks, many of which depend heavily on spatial reasoning, and 2) there are a wealth of RTS games available, facilitating transfer experiments.

Spatial Reasoning in Real-Time Strategy Games

To win RTS games, players must gather resources, research technologies, construct armies, and win military engagements. These games pose many AI challenges including: adversarial planning, continuous planning, temporal reasoning, and acting under uncertainty in dynamic environments. To be successful, representations must include spatial information regarding the environment and relationships between the units. Consider the following examples:

Resource exploitation is an essential component of a successful RTS player. To achieve this, the player must create defensive positions allowing units to safely gather resources. Effective defenses require not only the construction of appropriate units, but also their spatial configuration. An important tactic is to identify a narrow pathway through which enemy forces desire to the travel, called a *chokepoint*. By constructing sufficient defenses at

[1] While most transfer learning research has focused solely on the last two steps, a complete process model of transfer learning includes the identification of a useful source domain.

the chokepoint, the player creates a safe area to carry out actions (e.g., gather resources, repair units).

Another RTS game task is continuous planning, i.e., pursuing and monitoring multiple goals over an extended period of time with unexpected events. For example, while pursuing a resource gathering goal, the player becomes aware of an enemy unit. Now, the player needs to make a decision to either gather forces to engage the enemy or to ignore it. The spatial relationships between the enemy and the player's units are critical in deciding how to respond. If there are adequate defenses between the enemy and the player's resource gathers, then the player can continue pursuing the goal of resource gathering. Otherwise, the player should muster forces and engage the enemy.

A multi-representational architecture should be capable of reasoning about all of the tasks in the domain.

Spatial Transfer in Real-Time Strategy

Many different games have been implemented on top of open-source game engines, e.g., Stratagus (Posen *et al.* 2005). Furthermore, each game is highly tailorable, providing an excellent platform to explore transfer. Experience and knowledge regarding a particular RTS game should improve the performance in other games. This section describes three important ways in which spatial knowledge can participate in transfer.

First, spatial information can be used to help identify similarities between the source and target domains. Proteus (Davies *et al.* 2008) is an extreme example in which spatial information alone drives the mapping process. For example, after studying an example of a chokepoint, one could identify other chokepoints based on similar configurations of obstacles.

Second, spatial information can be included in the transferred knowledge. The NuSketch BattleSpace system (Forbus *et al.* 2002) provides an example of this in hypothesizing enemy location and intent. For example, the appearance of an enemy unit could result in a transfer conjecture of the locations of additional enemy forces.

Third, the transfer could include how to compute spatial relationships. Aha *et al.* (2009) transferred SVMs which quickly recognize opponent's football plays (Aha *et al.* 2009). In RTS games, this could involve identifying a pincer attack based on the separation of the opposing force into two equal sized forces.

A multi-representational intelligent agent should exploit spatial information throughout the transfer process.

Conclusions

This paper presents transfer learning in real time strategy games as an important benchmark for multi-representational intelligent agents. Much previous research on spatial reasoning has focused on primarily spatial tasks, such as wayfinding (Kurup 2007), or on modeling aspects of human perception (Tamborello & Byrne *in press*). RTS games present a range of tasks, which require spatial and conceptual reasoning, making them an excellent platform for exploring multi-representational intelligent architectures. While spatial representations are essential for performing these tasks effectively, we believe that they serve an additional role. Because spatial relationships are shared across domains, they invite comparison, highlighting similarities and differences. Therefore, spatial representations not only improve domain reasoning and learning, but also facilitate the transfer of knowledge across situations and domains.

References

Aha, D.W., Molineaux, M., & Sukthankar, G. 2009. Case-based reasoning for transfer learning. *Proceedings of ICCBR-09*. Seattle, WA: Springer.

Chandrasekaran, B., 1997. Diagrammatic representation and reasoning: some distinctions. In *AAAI Fall Symposium on Diagrammatic Reasoning*. Boston, MA.

Davies, J., Goel, A. & Yaner, P. 2008. Proteus: A Theory of Visual Analogies in Problem Solving. *Knowledge-Based Systems*. 21(7).

Forbus, K., Usher, J. & Chapman, V. 2003. Qualitative spatial reasoning about sketch maps. *Proceedings of IAAI-03*, Acapulco, Mexico.

Klenk, K., Forbus, K., Tomai, E., Kim, H., & Kyckelhahn, B. 2005. Solving everyday physical reasoning problems by analogy using sketches. *Proceedings of AAAI-05*. Pittsburgh, PA.

Kurup, U. 2007. *Design and Use of a Bimodal Cognitive Architecture for Diagrammatic Reasoning and Cognitive Modeling*. Ph.D. Thesis. Computer Science and Engineering, The Ohio State University

Laird, J. & van Lent, M. 2000. Human-level AI's Killer Application: Interactive Computer Games. Invited talk at AAAI-00.

Ponsen, M., Lee-Urban, S., Muñoz-Avila, H., Aha, & D., Molineaux, M. 2005. Stratagus: An Open-Source Game Engine for Research in Real-Time Strategy Games. In D.W. Aha, H. Muñoz-Avila, & M. van Lent (Eds.) *Reasoning Representation, and Learning in Computer Games: Papers from the IJCAI Workshop*. Washington, DC

Tamborello, F. P., II, & Byrne, M. D. (in press). Adaptive but Non- Optimal Visual Search Behavior in Highlighted Displays. *Cognitive Systems Research*.