

Special Issue on Structured Knowledge Transfer

Daniel G. Shapiro, Hector Munoz-Avila, and David Stracuzzi

■ *This issue summarizes the state of the art in structured knowledge transfer, which is an emerging approach to the problem of knowledge acquisition and reuse. Its goal is to capture, in a general form, the internal structure of the objects, relations, strategies, and processes used to solve tasks drawn from a source domain, and exploit that knowledge to improve performance in a target domain.*

A Note from the AI Magazine Editor in Chief:

Part Two of the Structured Knowledge Transfer special issue will be published in the summer 2011 issue (volume 32 number 2) of *AI Magazine*. Articles in this issue will include:

"Knowledge Transfer between Automated Planners," by Susana Fernández, Ricardo Aler, and Daniel Borrajo

"Transfer Learning by Reusing Structured Knowledge," by Qiang Yang, Vincent W. Zheng, Bin Li, and Hankui Zhuo

"An Application of Transfer to American Football: From Observation of Raw Video to Control in a Simulated Environment," by David J. Stracuzzi, Alan Fern, Kamal Ali, Robin Hess, Jervis Pinto, Nan Li, Tolga Könik, and Dan Shapiro

"Toward a Computational Model of Transfer," by Daniel Oblinger

Transfer learning is the problem of acquiring knowledge in one domain and using it to improve performance in another. While the field of psychology has studied transfer learning in people for many years, AI has only recently taken up the challenge. The topic received initial attention with work on inductive transfer in the 1990s, while the number of workshops and conferences has noticeably increased in the last five years. This special issue represents the state of the art in the subarea of transfer learning that focuses on the acquisition and reuse of *structured* knowledge. Its goal is to capture, in a general form, knowledge about the objects, relations, strategies, and processes used to solve tasks drawn from a source domain and exploit it in tasks taken from a target domain.

This special issue is motivated largely, but not entirely, by the completion of the DARPA Transfer Learning program, which funded the lion's share of work on structured knowledge transfer and has helped popularize the broader subject of transfer learning as a whole. The issue contains five articles sponsored by the program, plus two outside it, along with three opinion pieces authored by prominent figures in the field. Taken as a whole, the articles paint a picture of structured knowledge transfer as an emerging discipline that is just arriving at an understanding of the problems, methods, and techniques involved.

The articles in this special issue address transfer learning within a three-stage framework: (1) knowledge acquisition in a source task, (2) knowledge mapping between the source and target tasks, and (3) exploitation of the transferred knowledge in the target setting. We include work on near transfer, where the source and target tasks are quite similar (for example, drawn from the same domain), plus instances of far transfer where the tasks come from dissimilar domains that do not even share a common representational vocabulary. The papers describe

transfer within, and between, analysis and synthesis tasks, and they do so by communicating multiple types of knowledge. The objects of transfer include declarative structure (for example, model parameters and relational predicates), procedural knowledge (for example, plans), and more abstract problem-solving strategies (for example, search control heuristics). One of our goals for this special issue was to display this breadth of metaphor.

“An Introduction to Intertask Transfer for Reinforcement Learning,” by Matthew E. Taylor and Peter Stone illustrates transfer in reinforcement learning (RL) settings. It presents multiple methods of mapping actions, states, preferences, control knowledge, and specific value functions from source to target tasks. The work includes examples of both near and far transfer and contributes a broad categorization of transfer tasks. It also provides an insightful analysis of future research paths.

“Automatic Discovery and Transfer of Task Hierarchies in Reinforcement Learning,” by Neville Mehta, Soumya Ray, Prasad Tadepalli, and Thomas Dietterich, also explores transfer in an RL setting but focuses on a deep causal analysis of successful strategies. Here, the object of transfer is a hierarchical solution procedure that generalizes source behavior for use in the target task. Importantly, this work isolates the benefit of structure versus value transfer through lesion experiments.

“The Case for Case-Based Transfer Learning,” by Matthew Klenk, David W. Aha, and Matt Moliniaux illustrates the application of case-based reasoning to transfer tasks. It examines work by multiple authors to show how case similarity and case reuse can perform transfer as a whole or address the component problems of transfer (source knowledge acquisition, knowledge mapping, and knowledge exploitation with learning in the target task). This work advances the characterization of transfer problems and documents the clear but underappreciated relation between case-based reasoning and transfer tasks.

Thomas Hinrichs and Kenneth Forbus focus on far transfer in “Transfer Learning through Analogy in Games.” This work acquires solution strategies for problems set in the General Game Playing framework, and employs abstract action models as the object of transfer. It utilizes powerful analogical reasoning techniques to address the knowledge mapping task and exploits this knowledge in a theory refinement setting. Their system is novel in that it incrementally acquires action models through transfer while operating in the target context, where it explicitly tests, and accepts or rejects those models.

“Knowledge Transfer between Automated Planners” by Susana Fernández, Ricardo Aler, and Daniel Borrajo employs transfer to enhance solu-

tion methods rather than communicate solution elements. It shows how to acquire domain-specific search heuristics from one planning engine and transfer them into another where they serve to guide problem search. The authors demonstrate this technique using near transfer tasks drawn from multiple domains obtained from the International Planning Competition. This work has the flavor of a mixture method; it is interesting because it employs transfer to communicate bias among planning engines such that the net solution shares the best features both engines produce.

“Transfer Learning by Reusing Structured Knowledge” by Qiang Yang, Vincent W. Zheng, Bin Li, and Hankz Hankui Zhuo explains how optimization methods and techniques inspired by the concerns of data reuse (principally classification and model fitting) can be applied to extract and transfer underlying knowledge. The work addresses several near transfer tasks, namely localization, recommender systems, and action model discovery. This article is notable because it bridges data and knowledge reuse; it employs techniques normally associated with statistical transfer to communicate more general model structure.

Finally, David J. Stracuzzi, Alan Fern, Kamal Ali, Robin Hess, Jervis Pinto, Nan Li, Tolga Konik, and Dan Shapiro employ transfer to bridge analysis and synthesis tasks. “An Application of Transfer to American Football: From Observation of Raw Video to Control in a Simulated Environment” is a systems piece that chronicles the transformation of a single play from raw video of actual games through player tracking, play recognition, knowledge mapping, and play refinement in simulation. It incorporates sophisticated video processing and demonstrates the use of a cognitive architecture to encapsulate representation, recognition, and execution tasks. The effort is significant because it illustrates transfer across task types given real-world data.

The opinion pieces in the special issue provide unique perspectives on transfer. Jesse Davis and Pedro Domingos examine the prospects for deep transfer using Markov logic networks. They suggest that tools such as second-order logic, network analysis, and predicate invention can facilitate the discovery of abstract relational patterns and causal regularities that link seemingly unrelated domains, like physics and finance. Dan Oblinger provides a revealing analysis of the transfer learning program from a DARPA perspective. He gives a candid assessment of progress relative to program aspirations, including the communities’ ability to frame transfer learning issues and the status / prospects of current solution techniques. Ted Senator challenges us to broaden our focus by conducting research that begins with the discovery of a transfer opportunity and arrives at the identification of

AAAI Executive Council Elections

Please watch your mail for your AAAI Ballot, which will be mailed to all regular AAAI members in April. The membership will vote for four new councilors, who will each serve three-year terms. Ballots will be due back at the AAAI office no later than June 10. The Annual Business meeting for AAAI will be held this summer during AAAI-11 in San Francisco. The exact day and time will be announced on the ballot, in the summer issue of the magazine, and in the conference program.

relevant knowledge and methods only after significant processing.

Viewing the special issue as a whole suggests several lessons. The first is that structured knowledge transfer is effective, as it can profoundly improve performance by shaping the search space for subsequent learning. Mehta, Ray, Tadepalli, and Dietterich illustrate this effect by explicitly measuring the impact of transferred structure over value. Stracuzzi, Fern, Ali, Hess, Pinto, Li, Könik, and Shapiro provide a more qualitative demonstration by employing transfer to seed the play improvement task with an expert football play. This supplies a strategy and a parameter space that significantly expands the horizon for target learning.

A second lesson is to capture and transfer solution structure in general form. This is not at all surprising from a machine-learning perspective, but the issue illustrates many pathways to that goal. Yang, Zheng, Li, and Zhuo capture essential structure in parameters of mathematical models that apply in source and target contexts. Klenk, Aha, and Molineaux represent solution knowledge in cases, selected through similarity metrics, and exploited through case reuse. Hinrichs and Forbus transfer theories connecting objects and action that can be tested in the target context. Fernández, Aler, and Borrajo transfer search bias. Mehta et al. and Stracuzzi et al. transfer hierarchical skills that generalize beyond the source context, obtained either through causal analysis or a form of explanatory reasoning. Each of these methods gains lever-

age over a range of target problems by transferring solution structure expressed in a generalized form.

The third lesson is that work on structured knowledge transfer is in an exploratory stage. We lack a fundamental understanding of the type and difficulty of transfer tasks (as Oblinger notes), and while we agree on some evaluation metrics, we have not settled on evaluation methodologies. For example, Taylor and Stone point out that very few researchers include effort devoted to source knowledge acquisition when assessing transfer effectiveness, while Stracuzzi et al. note that it is even difficult to identify a basis for comparison when transfer enables new behavior. Our metrics are also geared to assess quantitative versus qualitative improvement on target tasks. More broadly, Senator argues that current transfer techniques embed very strong expectations about the task and relevant knowledge. In effect, we are collectively solving textbook exercises in place of the transfer problems that will occur in practice when the setting is far less controlled.

That said, this special issue demonstrates that we have collectively developed an enticing array of transfer technologies that capture and communicate key knowledge structures in general form, and that we can employ those techniques across a spectrum of application tasks. The next generation of transfer systems should exploit deeper shared structure found in more natural application settings, advancing transfer learning research toward general use. It will be a pleasure to see that happen.

Daniel Shapiro is the executive director of the Institute for the Study of Learning and Expertise (ISLE), where he conducts research on cognitive architectures, machine learning, and value driven cognitive systems. Shapiro is the chair for the Innovative Applications of Artificial Intelligence conference in 2011 and a member of the of the board of directors of the Value Driven Design Institute (VDDI). He received his PhD in management science and engineering from Stanford University in 2001.

Hector Munoz-Avila is an associate professor at the Department of Computer Science and Engineering at Lehigh University. Munoz-Avila has done extensive research on case-based reasoning, planning, and machine learning. He is also interested in advancing game AI with AI techniques. Munoz-Avila is recipient of a National Science Foundation (NSF) Career award and two papers awards. He currently holds a Lehigh Class of 1961 Professorship. He has been chair for various international scientific meetings including the Sixth International Conference on Case-Based Reasoning (ICCB-05).

David Stracuzzi is a senior member of the technical staff at Sandia National Laboratories. He received his doctorate from the University of Massachusetts at Amherst in 2006, followed by a postdoc at Stanford and three years as research faculty at Arizona State. His research interests include machine learning, cognitive architectures, and cognitive modeling.