

Learning Models of Macrobehavior in Complex Adaptive Systems

Andrew Fast

Department of Computer Science
University of Massachusetts
Amherst, MA 01003
afast@cs.umass.edu

Macrobehaviors in complex adaptive systems are high-level behaviors that arise from a combination of low-level interactions and characteristics of individuals. Over time, individuals participating in these complex systems make choices that influence the behavior of other individuals. These types of interactions are present in both human and computational systems. For example, complex interactive behaviors are demonstrated as audience members enter a talk (Schelling 1978) and software agents negotiate over shared resources (Klein, Metzler, & Bar-Yam 2005). Macrobehaviors arise when groups or aggregations of individuals develop group-level structure through interactions of individual characteristics and actions.

One complex system that is familiar to academics is the pattern of citations in a given field or subfield of science (e.g., (McGovern *et al.* 2003)). Citation networks exhibit many interesting behaviors: researchers rise and fall in prominence, papers get cited at different rates at different times, research topics become popular then fade out over time, etc. One example of a macrobehavior is the rise to prominence and ongoing success of a research group. A typical university research group contains a few faculty (usually one or two), graduates students studying with those faculty, and perhaps additional research or technical staff. Presumably, these researchers have joined together because of shared research interests and the opportunity to build ideas collectively. Citation frequency of both the individual researchers and the whole research group depend on the richness of the interactions within the group. Intuitively, we recognize that groups with the right mix of people and talents will tend to be cited more than the sum of those same individuals working separately.

Macrobehaviors, also called *emergent* behaviors, can lead to mutually enhanced behavior (as in the case of research groups), mutually detrimental behavior, or simply neutral behavior that is otherwise remarkable to system observers. Mutually detrimental behaviors, often called *social pathologies* or *social dilemmas*, are of particular interest to researchers. Jensen and Lesser define a social pathology as a “system behavior in which two or more agents interact such that improvements in local performance do not improve system performance” (Jensen & Lesser 2002). Examples

of social pathologies in computational systems include the “tragedy of the commons” (Turner 1993) and resource oscillation (Klein, Metzler, & Bar-Yam 2005).

Due to the size and complexity of real systems, it is often difficult to identify and predict the occurrence of macrobehaviors. These behaviors arise from transient combinations of local characteristics, the structure of the network guiding the interaction, and adaptations of surrounding entities. Models that are able to represent these behavioral factors would provide a significant improvement for the analysis of complex systems. In this thesis, I propose to extend current techniques for learning in *relational data* to develop (1) models of group structure and attributes and (2) flexible models of temporal dependencies. These two contributions will provide a foundation for models of macrobehaviors that are understandable and easily learned from data.

Recent work in complex network analysis and modeling (e.g., (Watts & Strogatz 1998; Kleinberg 2001)) has led to the adoption of a graph with attributes as the primary choice for representing complex systems. A graph effectively captures the local interactions and can be abstracted to better understand macrobehaviors in sections of the network. Much of the work done recently on learning models of *relational data* can be applied to graphs with attributes. Relational data represent interrelated entities such as people, places, things and events; their defining characteristics (often called attributes); and the structure capturing the relationships between these entities. Examples of relational data include citation data; the network of Hollywood movies, actors, and studios; and hyper-linked data found on the Internet. Much of the work with relational data has concentrated on learning models about single entities either individually (e.g., (Neville *et al.* 2003)) or collectively (e.g., (Taskar, Abbeel, & Koller 2002; Jensen, Neville, & Gallagher 2004; Neville & Jensen 2004)). Recent work has shown that hypothesizing the existence of groups in data can lead to increased performance inferring individual attributes (Neville & Jensen 2005). However, this work is not able to model the characteristics and structure of the groups themselves.

In order to represent macrobehaviors effectively, models must be able to consider the attributes and structure of an aggregation of interrelated entities as a whole rather than a collection of individuals. This is a bottom up approach, inferring the attributes of groups from the structure of the in-

dividual interactions. From a perspective of machine learning, we are interested in variables of the groups (i.e., an aggregation of entities) themselves rather than variables of individual objects. Consider the research group example: It may be possible to identify research groups solely on the network structure among co-authors (i.e., individuals tend to coauthor within the groups more than without); however, the best model of a research group would also consider the roles of individuals (student, faculty, etc.) and constraints on the composition of the groups such as a limit on the ratio of faculty to students in the group. In order to better model groups, I will develop a new class of compositional models for relational data by extending the feature space considered by the *relational probability tree* (Neville *et al.* 2003). In addition to searching over the space of individual attributes and aggregations of individuals, these compositional models will consider features that capture the composition of the network and attribute relationships among individuals.

Because many complex systems exhibit dynamic behaviors, a necessary component of this work will be to introduce flexible, easy-to-learn models of temporal phenomena. Current methods for relational data are limited *a priori* to a fixed set of temporal relations (e.g., a single past time step) (Sanghai, Domingos, & Weld 2003; 2005). Emergent behaviors often arise from the interactions of individual behaviors at various time scales and rates. This necessitates flexible models that are able to identify which time scales and temporal intervals are most relevant for detecting and predicting macrobehaviors in the system being studied. To address this, I propose a probability-tree algorithm based on the relational probability tree (Neville *et al.* 2003) that is able to search over aggregations of objects at multiple scales and intervals to identify the most predictive features.

Both compositional and temporal features search over aggregations at different levels of granularity in the data. This feature-space expansion incurs a burden in both variance and efficiency of the models and may introduce feature selection biases. For a model to be successful, any burdens must be offset by the improved probability estimation of the new models and any biases of the model must be addressed by corrected feature selection techniques.

A primary goal of this work is to improve knowledge discovery and data mining tools by developing compact, understandable models of groups that are applicable to a wide variety of applications. In prior work, I developed an algorithm that was able to capitalize on group structure in a P2P file-sharing network by clustering music files that tend to occur together in users' file-sharing libraries (Fast, Jensen, & Levine 2005). By utilizing this structure, we were able to improve query performance of a standard P2P algorithm. In addition to P2P networks, this work will be applicable to citation networks, multi-agent systems such as unmanned aerial vehicles (UAVs) and other complex systems.

With advances in many fields of distributed computation including networking and multi-agent systems, there has been a proliferation of data capturing the behavior of computational complex systems. These datasets are so large and inherently complex that they are impossible to examine unaided. Without better methods for analyzing and interpreting

the behavior of such systems, we risk stagnated development of new systems. By developing automated tools for the analysis of macrobehaviors in complex systems, we will be able to foster the understanding needed to generate advances for the next generation of computational systems.

References

- Fast, A.; Jensen, D.; and Levine, B. 2005. Creating social networks to improve peer-to-peer networks. In *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Jensen, D., and Lesser, V. 2002. Social pathologies of adaptive agents. In *2002 AAAI Spring Symposium on Safe Learning Agents*.
- Jensen, D.; Neville, J.; and Gallagher, B. 2004. Why collective inference improves relational classification. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Klein, M.; Metzler, R.; and Bar-Yam, Y. 2005. Handling emergent resource use oscillations. *IEEE Transactions on Systems Man and Cybernetics (TSMC)*.
- Kleinberg, J. 2001. Small world phenomena and the dynamics of information. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14.
- McGovern, A.; Friedland, L.; Hay, M.; Gallagher, B.; Fast, A.; Neville, J.; and Jensen, D. 2003. Exploiting relational structure to understand publication patterns in high-energy physics. *SIGKDD Explorations* 5(2):165–172.
- Neville, J., and Jensen, D. 2004. Dependency networks for relational data. In *Proceedings of The 4th IEEE International Conference on Data Mining*.
- Neville, J., and Jensen, D. 2005. Leveraging relational autocorrelation with latent group models. In *Proceedings of the 5th IEEE International Conference on Data Mining*.
- Neville, J.; Jensen, D.; Friedland, L.; and Hay, M. 2003. Learning relational probability trees. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Sanghai, S.; Domingos, P.; and Weld, D. 2003. Dynamic probabilistic relational models. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Sanghai, S.; Domingos, P.; and Weld, D. 2005. Learning models of relational stochastic processes. In *Proceedings of the Sixteenth European Conference on Machine Learning*.
- Schelling, T. C. 1978. *Micromotives and Macrobehavior*. W. W. Norton & Company.
- Taskar, B.; Abbeel, P.; and Koller, D. 2002. Discriminative probabilistic models for relational data. In *Proceedings of Uncertainty in Artificial Intelligence*.
- Turner, R. M. 1993. The tragedy of the commons and distributed AI systems. In *Twelfth International Workshop on Distributed Artificial Intelligence*.
- Watts, D. J., and Strogatz, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature* 393:440–442.