Proceedings of the Twenty-Fifth Innovative Applications of Artificial Intelligence Conference

Scalable Models for Patterns of Life

J.T. Folsom-Kovarik[a], Sae Schatz[b], Randolph M. Jones[a], Kathleen Bartlett[b], Robert E. Wray[a]
[a]Soar Technology, Inc.; [b]MESH Solutions, LLC, a DSCI Company
{jeremiah | rjones | wray}@soartech.com; {sschatz | kbartlett}@mesh.dsci.com

Patterns of Life

Patterns of life (POL) are emergent properties of complex social systems such as neighborhoods or even cities (Schatz et al. 2012). In such systems, observable regularities manifest from the interactions of individuals' behaviors and social norms. Simultaneously, these patterns of life impose structure on individual decisions. For example, a pattern of rush hour traffic arises from drivers' decisions to commute at a certain time. Knowledge of rush hour influences individuals' departure times.

Accurately modeling POL is not only an academic pursuit; complex training and analysis efforts rely upon these models. As an example, social and cultural experts analyzed patterns of life in order to locate Osama bin Laden. These analysts, without any visual confirmation, cited the ways his household deviated from the "normal" pattern of life in the neighborhood, behaving much more secretively than his suburban neighbors (Preston 2012).

Building on such social and cultural expertise, *computational* modeling of POL offers significant potential to continue extending POL's practical application and theoretical study. In industry, POL models might support decisions about marketing, logistics, network security, or building design. Military and law-enforcement models could support observational and small-unit training, intrusion detection, emergency planning, or community relations. Computational POL models will benefit academic researchers by giving them new tools to generate and validate theories about complex sociocultural patterns.

Finally, POL presents an interesting computational modeling challenge for research in many AI fields because of its complex, multi-level interaction. POL modeling differs from standard agent-based modeling in that not only individual behaviors, but also their collective emergent patterns, are targets of study. At the same time, the behavior of any single individual is potentially important.

As a specific example, the Virtual Observation Platform (VOP) is an immersive simulation designed to train military personnel to interpret indicators before a critical incident, such as a terrorist attack, occurs (Schatz et al. 2012). Pre-event indicators are typically ambiguous variations in a typical pattern: a car parked in an odd spot, a sparse crowd in a normally busy and crowded public square. Trainees will practice identifying and reporting regularities of individual behavior, small group interactions, and population gestalt as they watch a simulation play out in real time. They should see subtle variations in interactions between simulated characters and interpret the underlying meaning. To support these needs, the system requires a POL model.

Challenge: Scalable Models of Patterns of Life

Creating computational models of POL, however, is a significant scientific and technical challenge. To meet it, researchers must achieve simultaneous scalability along three key dimensions: population size, intelligence, and automatic behavior specification. Below, we outline the POL scalability challenges and argue that AI offers methods and tools that have been used successfully in solving problems with features similar to POL modeling.

Population size. Generating the behavior of many individuals in real time (or faster) remains a technical hurdle for POL. Even small villages contain hundreds of residents; cities have millions. In the VOP example, trainees might need to see hundreds of simulated characters populating a public marketplace in order to understand the normal range of behaviors and interactions in such an environment. At the same time, many more characters need to populate every surrounding business and city street that the trainees can observe. The large population is vital to providing the required fidelity in background activities and creating an appropriate context for identifying anomalies.

In the current state of the art, researchers contemplate scaling to millions of entities via parallelism (Aaby et al. 2010). However, narrow individual intelligences are required to model large populations today. For example, fluid dynamics efficiently models entity movement in large crowds (Hughes, 2002) but is limited to modeling only physical movements and not the meetings, arguments, and interchanges that take place between friends and strangers.

Individual Fidelity. Scalable POL requires greater breadth in intelligence of individuals, including richness of inputs, behavior choices, and decision complexity in order to produce realistic behavior for individual entities. For example, the displayed behaviors that underlie patterns may need to be highly robust so that they can be carried out despite the interference of human users in the simulation. In cases where humans do not directly influence a simulation, intelligence is still required in order to respond to the many other agents that populate the environment.

Unfortunately, efficient approaches like fluid dynamics models do not model individual intelligence, and therefore do not provide the fine-grained interactions necessary for many POL applications. Abstract models that completely subsume individual decisions in order to display group behaviors lack believability under fine observation of individuals, and if no goal or reason drives an agent, it is impossible to discern a correct meaning of its behaviors. Conversely, executing a fine-grained model of every individual, at all times, is likely to create unacceptable authoring and run-time computation requirements.

As a consequence, a heterogeneous environment of intelligent entities is necessary in POL systems. This introduces additional run-time challenges: e.g., *blending* relatively few highly intelligent agents with less intelligent entities so that distinctions between them are not obvious to an observer and dynamically *switching* individuals between simple and robust controllers, letting the entire population present large-scale patterns and giving select entities more intelligence only when needed to ensure vital interactions appear. We are taking this approach in investigating efficient POL representations that support planning and reacting for a small city of entities (Jones et al. in preparation).

Behavior specification. POL model synthesis includes authoring, generalization to diverse settings, and run-time adaptation. The complex interaction of components in generating a pattern of life creates tension between human specification (which is costly in time and labor), automated model creation from real-world data (which may be too tied to specific populations and biased by availability of limited data sources), and synthetic models (which may produce unrealistic POL). Researchers must determine which aspects can be automatically generated and how to ease human control over the many details that cannot be automated.

Authoring POL content can draw from real-world data such as video sources, authoritative input from residents or anthropologists, or cultural artifacts like media. Current research in these fields will support improvements to POL, but important issues of understanding such inputs for the purposes of computational modeling remain. Content generalization raises questions of identifying cultural universals, defining cultural differences, and creating relevant model overlays that efficiently capture the differences. Run-time adaptation of the patterns being displayed requires model representations that can influence many individual behaviors with fewer, more abstract directives. Without this abstraction capability, the patterns produced will be unacceptably dependent on individual decisions, making it difficult to identify core causal patterns.

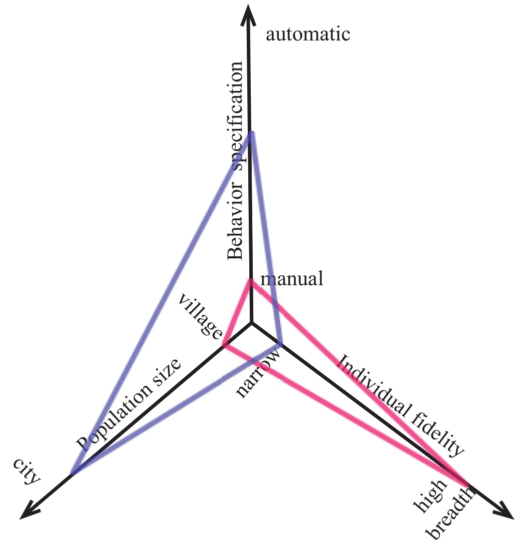

Figure 1: Current models of patterns of life do not scale in all three model dimensions.

Artificial Intelligence Solutions

Current population simulations often have limited application or research impact due to a lack of scalability along at least one of the scalability dimensions. A city-sized traffic model for the VOP could be synthesized automatically from observed traffic patterns but would disallow a spontaneous street celebration due to the limited knowledge of individual entities. A high-fidelity model of the interpersonal interactions that occur with a few village elders during a military patrol in the VOP would not generalize to a town-sized population due to knowledge engineering costs. Modeling patterns of life that scale along all three dimensions requires new approaches and algorithms, built on artificial intelligence foundations. We outline three AI research areas on which future solutions could build.

Pattern recognition. Instantiating an accurate, specific POL model depends on close observation of actual patterns of life. Encoding these observations via knowledge engineering methods is not practical for scalable POL. Automated methods are needed to identify and extract human behavior patterns in order to make POL specification scalable, especially in the size dimension. AI offers many examples of transforming captured human behavior into useful patterns. Data mining in conjunction with social network analysis can reveal criminal networks from text documents (Al-Zaidy et al. 2012). In computer vision, researchers can classify social interactions or crowd movements (Mehran, Ohyama, and Shah 2009).

New challenges in the fields of data mining and pattern recognition for POL will be to extract the more subtle variations of everyday life, to recognize gestalt patterns in addition to individual events, and to incorporate common sense and narrative understanding into diverse settings.

In order for these approaches to identify and then simulate POL regularities and anomalies, researchers need to identify which features to consider despite the breadth and subtleties of human interaction. Identifying these features and data sources to support them is a nontrivial task. For example, in the VOP domain, observing human interactions in a village can be contrasted with the superficially similar problem of protecting a ship in harbor from a small craft attack (DoD 2001). In a port setting there are a large number of interacting individuals (ships). However, in contrast to a POL problem, the movements of ships are more limited in number of possible behaviors and goals, requiring bright lines for force protection rules rather than complex sensemaking. In POL, new metrics are required for recognizing good inputs, outputs, and problems for study.

Representational abstraction. A "missing link" in modeling POL is the lack of general, stable representations of population behavior. However, finding and exploiting computationally precise, useful representations at different levels of abstractions is a hallmark of AI, from the study of the processes of mind and structure of games in AI's infancy to recent advances in market abstractions and agency.

POL researchers have defined representations, such as social networks and crowd models, which reflect AI's representational heritage. These results, while important, are comparable to Simon's (1996) individual watch parts. Progressing directly from these scattered elements to a useful and scalable POL would be extremely difficult.

We recommend a shift to identifying higher-level abstractions (representations and composition rules) that provide the "stable subsystems" to bridge the gap between individual components and the working watch of a scalable POL simulation. Hierarchical representations of crowds, groups, and individuals (e.g., Musse and Thalmann 2001) suggest a path toward this goal, but significant extension will be required to include the many factors that influence patterns of life.

Behavior generation. Computational models of human behavior (e.g., Anderson and Lebiere 1998, Laird 2012) and intelligent agents (e.g., Wooldridge 2000) have also been major thrusts in AI. Generative behavior models are needed in order to reify and explore new findings about POL. Researchers who develop generative POL models will contribute to validation and verification of work in the new field and provide important utility to applied users.

In this area, crowd-integrated cognitive models (Pelechano et al. 2005) and agent-based systems that can plan ahead, reason, and react (Laird 2012) are relevant to generating complex behaviors that can both support and potentially control emergent patterns. Prior work in heterogeneous computational modeling of individuals (Lebiere et al. 2002, Wray et al. 2005) may offer initial points of exploration for heterogeneous behavior generation for POL.

In conclusion, contributions from many fields of AI will be welcomed to meet key challenges and realize new opportunities by creating scalable models for patterns of life.

Acknowledgement

This work is supported in part by Office of Naval Research (ONR) contract N00014-11-C-0193. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ONR or the US Government. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

Aaby, B.G.; Perumalla, K.S.; and Seal, S.K. 2010. Efficient simulation of agent-based models on multi-GPU and multi-core clusters. In proceedings of the 3rd International ICST Conference on Simulation Tools and Techniques.

Al-Zaidy, R.; Fung, B.; Youssef, A.M.; and Fortin, F. 2012. Mining criminal networks from unstructured text documents. *Digital Investigation* 8(3–4): 147–160.

Anderson, J. R. & Lebiere, C. 1998. *The atomic components of thought.* Mahwah, NJ: Erlbaum.

DoD [Department of Defense]. 2001. DoD USS Cole Commission Report. Washington, DC. January 9, 2001.

Hughes, R.L. 2002. A continuum theory for the flow of pedestrians. *Transportation Research Part B* 36(6), 507–535.

Jones, R.M.; Folsom-Kovarik, J.T.; McLaughlin, P.; and Frederiksen, R. (in preparation). High-density pattern-of-life modeling.

Laird, J.E. 2012. *The Soar Cognitive Architecture.* Cambridge, Mass.: MIT Press.

Lebiere, C., Biefeld, E., Archer, R., Archer, S., Allender, L., & Kelley, T. (2002). IMPRINT/ACT-R: Integration of task network modeling architecture with a cognitive architecture. In proceedings of the Advanced Technologies Simulation Conference.

Mehran, R.; Oyama, A.; and Shah, M. 2009. Abnormal crowd behavior detection using social force model. Proc.of IEEE Conference on Computer Vision and Pattern Recognition, 935–942.

Musse, S.R. and Thalmann, D. 2001. Hierarchical model for real time simulation of virtual human crowds. *IEEE Transactions on Visualization and Computer Graphics* 7(2): 152–164.

Pelechano, N.; O'Brien, K.; Silverman, B.; and Badler, N. 2005. Crowd simulation incorporating agent psychological models, roles and communication. In proceedings of the First International Workshop on Crowd Simulation, 21–30.

Preston, S. W. 2012. CIA and the Rule of Law. Harvard Law School, Cambridge, MA. April 10, 2012. Public address.

Schatz, S., Folsom-Kovarik, J.T., Bartlett, K., Wray, R.E., & Solina, D. 2012. Archetypal Patterns of Life for Military Training Simulations. In proceedings of the Interservice/Industry Training, Simulation & Education Conference.

Simon, H.A. 1996. *The sciences of the artificial.* Cambridge, Mass.: MIT Press.

Wooldridge, M. 2000. *Reasoning about Rational Agents*. MIT Press.

Wray, R.E.; van Lent, M.; Beard, J.; and Brobst, P. 2005. The design space of control options for AIs in computer games. In proceedings of the 2005 IJCAI Workshop on Reasoning, Representation, and Learning in Computer Games, 113–118.