Research Statement

Theodore P. Pavlic

Research Interests

I am particularly interested in decision-making in single- and multi-agent complex systems. I study both how to design effective autonomous decision-making systems as well as the rationale behind decisions of living agents in natural systems. Consequently, my work has been published in peer-reviewed journals in fields as disparate as engineering [e.g., 3, 9] and behavioral ecology [e.g., 5, 8]. Presently, I study software verification for hybrid and stochastic cyber-physical systems that emerge in urban environments consisting of both autonomous and human-driven vehicles; this application area requires certifiably safe technology that is both robust and complementary to human decision making. In general, I use theories from optimization, parallel computing, networks, dynamical systems, and formal logic to explain, design, or verify emergent system-level phenomena. Here, I describe some prior research successes and pending work of mine.

Biomimicry: Solitary Foraging Theory for Engineering (Complete)

There is a convenient homomorphism between autonomous task-processing agents (e.g., military patrol vehicles) and solitary foraging animals. In both cases, the agent (forager) faces an incoming queue of tasks (prey) according to a merged Poisson process with a specified rate. Processing tasks (feeding) and searching are both costly activities, but the agent must process tasks in order to accumulate value (calories) necessary for proliferation by the designer (natural selection). A good design will choose which encountered tasks to process and how long to process each task until returning to search so as to maximize the value the agent will accumulate over its runtime (lifetime).

Modular objectives: Stochastic dynamic programming (SDP) is the natural solution to the solitary agent problem. However, realistic static optimization objectives can be used to increase analytical tractability or reduce real-time computational complexity. Models from behavioral ecology make the simplifying assumption of long lifetimes to justify the use of asymptotic average rate of gain as a static performance metric. However, not only do engineering parameters come from coarser sets than the analogous ecological parameters, but engineering applications may have qualitatively different objectives (e.g., maximize efficiency instead of accumulated value) than those considered in ecology. Moreover, even when accumulated value is being maximized, static optimization objectives that do not require the assumption of long runtimes are desirable. So I generalized the intuitive analysis methods used in behavioral ecology to apply to a broad class of objective functions [1]. A compact version of these generalizations recently appeared in a robotics journal [9], and that work shows how an existing foraging-inspired decision-making strategy can easily be modified for finite-lifetime performance that is better than its ostensibly optimal bio-inspired strategy.

Speed and accuracy: Due to the speed-dependent probability of detecting a task, increasing search speed does not necessarily increase the encounter rate with tasks. However, the choice of optimal search speed should not be influenced by tasks that will be ignored due to their low profitability relative to other tasks in the environment. Thus, choosing search speed and which task types to ignore are coupled problems. Ecologists have studied this problem via computer simulation for a small number of task types with detection functions tailored to specific species of foragers. Instead, I formulated prototypical detection functions so that I could analytically solve the coupled speed–type choice problem for an arbitrary number of task types. The resulting algorithm was then validated on a simulated unmanned air vehicle [3].

Ostensibly Irrational Decision-making in Human and Non-Human Animals (Complete)

Impulsiveness: Behavioral ecologists conventionally assume that natural selection favors animal decision-making rules that maximize long-term energetic accumulation rate. However, when animals in a laboratory are given a mutually exclusive binary choice between two food items, they often prefer the one with the lower handling time regardless of the associated energetic reward. Based on these results, some ecologists suggest that the rate-maximization assumption may be a flawed. However, I have found theoretical evidence that the laboratory environment itself can generate suboptimal behaviors [5]. Because simultaneous encounters in a Poisson process occur with zero probability, binary-choice behaviors are outside of the traditional optimal foraging theory formulation. Using a low-complexity algorithmic implementation of foraging-like behavior on a robot as inspiration, I show how a simple decision rule can maximize a forager's long-term energetic rate in natural Poisson-consistent scenarios but can appear to irrationally favor short handling times when facing repeated mutually exclusive binary-choice trials. Moreover, this deleterious effect is predicted to be enhanced by starvation. Because the conventional operant laboratory experiment uses training by starvation followed by binary-choice trials, the laboratory apparatus may be the source of the strange behavior.

Sunk costs: The generalized foraging analysis I developed also helps explain sunk-cost effects observed in nature that are conventionally thought to be irrational [8]. For example, even though economic models predict that residence time should be independent of entry costs, the time swans spend underwater foraging for tubers is positively correlated with the energy expended breaking the surface of the water. However, I show that when a forager moves into an area where there are an abundance of food sources that have a much higher entry cost, it is beneficial for the forager to reduce the frequency of the accumulation of these costs. So additional in-patch foraging displaces time otherwise spent entering costly patches. This result is not obvious to conventional foraging theory because of the overly simplified gain functions used in the models; it is the generalization to arbitrary gain functions that explains the interesting behavior.

Group Task Processing: Cooperation and Distributed Optimal Resource Allocation

Task-processing networks: I have also studied task-processing patterns in groups of agents [2]. In one example [6], tasks arrive at different agents in a network, and those agents can request processing assistance from their neighbors. Those neighbors must choose how often to answer calls for assistance. Such systems occur with cooperative breeders, flexible manufacturing systems, and human organizations. To minimize communication and coordination constraints, I assume that each agent optimizes a local objective function that encodes the cost and value of processing tasks. I give conditions on the network topology and on the local utility functions that guarantee existence of a unique Nash equilibrium as well as totally asynchronous convergence of distributed numerical algorithms to the equilibrium. Moreover, when certain motifs exist in the network topology, the corresponding Nash equilibrium has reduced cooperation willingness at relatively loaded agents and increased cooperation willingness at neighbors of relatively loaded agents. Hence, the competitive Nash equilibrium leads to a collectively good solution. Specially tailored versions of these results have been presented at complex systems [4] and sociobiomimicry [7] conferences, and the general work has been submitted to an engineering journal [6]. Several extensions of the work are possible.

Distributed resource allocation under constraints: I also study resource allocation models inspired by problems from the emerging field of smart lighting [2]. For example, an array of autonomously controlled lights may be required to meet several given illumination thresholds using minimum power across the array. These inseparable constrained optimization problems can be solved using numerical dual-space methods on a centralized controller. However, I have shown that these problems are generalizations of the classical ideal free distribution (IFD) augmented with nutrient constraints. Moreover, the lighting problem can be solved cheaply in a decentralized fashion by mimicking eusocial insect foraging under colony-level constraints. Other resource allocation problems, like economic dispatch in power systems, also fit within this framework.

Software Verification in Mixed-Traffic Urban Cyber-Physical Systems

My most recent work uses formal methods from computer science and electrical engineering to verify safety properties in cyber-physical systems. I have been focusing on autonomous vehicles in urban environments. One problem of particular interest is the Hoare-like verification of the safety of adaptive cruise control software on highways were there is limited vehicle-to-vehicle communication. Prior work has been limited to model checking, is overly conservative, and assumes vehicles are limited to emergency braking maneuvers. Our work incorporates specifications for both physical models and software code. Moreover, we verify smooth maneuvers that guarantee safety [10, 11]. I have also been collaborating with peers on the real-time synthesis of correct-by-design autonomous controls and the hybrid-state estimation of human driver intent.

References

[1] Theodore P. Pavlic. Optimal foraging theory revisited. Master's thesis, The Ohio State University, Columbus, OH, 2007. URL `http://www.ohiolink.edu/etd/view.cgi?acc_num=osu1181936683`.

[2] Theodore P. Pavlic. *Design and Analysis of Optimal Task-Processing Agents.* PhD thesis, The Ohio State University, Columbus, OH, August 2010. URL `http://rave.ohiolink.edu/etdc/view.cgi?acc_num=osu1281462093`.

[3] Theodore P. Pavlic and Kevin M. Passino. Foraging theory for autonomous vehicle speed choice. *Engineering Applications of Artificial Intelligence*, 22:482–489, 2009. doi:10.1016/j.engappai.2008.10.017.

[4] Theodore P. Pavlic and Kevin M. Passino. Cooperative task processing. In *Proceedings of the ICAM 2009 Symposium: Emergence in Physical, Biological, and Social Systems IV*, Ann Arbor, Michigan, November 13, 2009. Poster abstract.

[5] Theodore P. Pavlic and Kevin M. Passino. When rate maximization is impulsive. *Behavioral Ecology and Sociobiology*, 64(8):1255–1265, August 2010. doi:10.1007/s00265-010-0940-1.

[6] Theodore P. Pavlic and Kevin M. Passino. Cooperative task processing. *IEEE Transactions on Automatic Control*, 2010. Submitted.

[7] Theodore P. Pavlic and Kevin M. Passino. Design and analysis of cooperative task processing agents. In *Proceedings of the Third Annual Frontiers in Life Sciences conference – Social Biomimicry: Insect Societies and Human Design*, Tempe, Arizona, February 18–20, 2010. Poster abstract.

[8] Theodore P. Pavlic and Kevin M. Passino. The sunk-cost effect as an optimal rate-maximizing behavior. *Acta Biotheoretica*, 59(1):53–66, 2011. doi:10.1007/s10441-010-9107-8.

[9] Theodore P. Pavlic and Kevin M. Passino. Generalizing foraging theory for analysis and design. *International Journal of Robotics Research [Special Issue on Stochasticity in Robotics and Bio-Systems Part 1]*, 30(5):505–523, April 2011. doi:10.1177/0278364910396551.

[10] Theodore P. Pavlic, Paolo A. G. Sivilotti, Alan D. Weide, and Bruce W. Weide. Comments on 'Adaptive cruise control: hybrid, distributed, and now formally verified'. Technical Report OSU-CISRC-7/11-TR22, The Ohio State University, 2011. URL `ftp://ftp.cse.ohio-state.edu/pub/tech-report/2011/TR22.pdf`.

[11] Theodore P. Pavlic, Paolo A. G. Sivilotti, Alan D. Weide, and Bruce W. Weide. Verification of smooth and close collision-free cruise control. In *Proceedings of the 2011 Symposium on Control and Modeling: Cyber-Physical Systems*, Urbana, Illinois, October 20–21, 2011. Poster abstract.