Research and Teaching Statement

Theodore P. Pavlic

Research Interests

I am particularly interested in distributed decision-making in single- and multi-agent complex systems. Recently, I have shifted my focus to software verification in hybrid and stochastic cyber-physical systems. My goal is to combine work from these two areas for the verification of safety and liveness properties that emerge from distributed decentralized systems. In general, I combine theories from optimization, parallel computing, networks, non-linear dynamical systems, and formal logic to explain, design, or verify emergent system-level phenomena. I prefer problems that are multi-disciplinary or have solutions that can be applied in varied contexts. Consequently, my results have been published in peer-reviewed journals in fields as disparate as engineering [e.g., 4, 10] and behavioral ecology [e.g., 6, 9]. Here, I describe some of my prior research successes and some 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 [2]. A compact version of these generalizations recently appeared in a robotics journal [10], 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 [4].

Ostensibly Irrational Decision-making in 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 [6]. 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 [9]. 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: Cooperative Nash Equilibria and IFD Extensions

Task-processing networks: I have also studied task-processing patterns in groups of agents [3]. In one example [7], 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 [5] and sociobiomimicry [8] conferences, and the general work has been submitted to an engineering journal [7]. 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 [3]. 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 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 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 [11, 12]. 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.

Teaching Interests

Facing an audience composed of mostly millennials, classroom teachers compete for attention with various forms of interactive and highly available content. Staying relevant in this atmosphere becomes more difficult as students see less incentive to physically attend classes that cover content that they can find on-line. Some universities have embraced self-paced e-Learning not only as a way to appease these students but also to free up time for faculty to do research. However, rather than competing directly with purely on-line schools, universities can leverage faculty experience to create unique in-person offerings that are complemented by an engaging on-line presence. Along with the opportunity to do research with faculty, instructors can use inquiry-based teaching methods to augment traditional classroom learning with self-paced opportunities that prepare students to do future research. I focus on creating educational experiences tailored to attracting and retaining this new audience. This approach has involved using open-source course content that evolves with student feedback, using inquiry in small-room laboratories to generate synergistic emergent educational experiences, and generating course experiences from novel open-ended short-term problems that students must solve in either classroom or extracurricular settings.

Open-Source Collaborative Course Content

Rather than using textbooks as the primary reference material in the courses I have taught, I have produced hundreds of pages of shorter documents[1] that are not only tailored for each course but provide multiple ways to engage the student. Each document is adequate for printing, but its electronic version also includes hyperlinks that: assist in navigation throughout the document, link to additional third-party sources of information that I have reviewed on the web, and link to other course documents. These electronic versions are fully searchable and include professional-quality graphics that can be used by students in their formal reports. Because undergraduate students may have little experience writing technical reports, each document is written to serve as an example. Hence, these documents simultaneously guide each student through the classroom experience, provide her help finding additional technical resources, and serve as examples of language and format that she should use in her course submissions. Additionally, because the students know that these are living documents that I have contributed to, they are more likely to contact me about questions they have about the content. Moreover, I can easily make adjustments to the living documents based upon the responses I receive from the students.

It is important to me to introduce students to tools that allow them to be successful in their courses but also reduce barriers to entering into other academic pursuits. Consequently, every course document of mine is produced entirely with freely available LaTeX (including graphics, which are produced natively in LaTeX with PSTricks as opposed to using a GUI tool that may be less available) and distributed freely on-line with its source code[2] under an open-source Creative Commons license. Hence, not only are the documents available for the students to learn from, but other instructors can reuse and modify them (e.g., to adapt them to their courses or to correct any mistakes I might have made). I also provide homework and *curriculum vitae* LaTeX templates to assist students in their other work. These templates have received a great deal of attention on the Internet; for example, my LaTeX homework template is the #1 result from a Google search for "latex homework template," and my other course documents and templates appear as page-1 results for other Google queries like "latex CV template," "transistor basics," and "rotary electrodynamics." I frequently receive positive feedback about these efforts from students, instructors, and engineering professionals.

Using Laboratory Inquiry to Prove Rules by Exception

The scientific method was not explicitly taught to me as young engineering student, and I viewed engineering as separate from science. Later, it became clear to me that the process of generating knowledge about newly created technology is a science and is subject to scientific inquiry. To experiment with inquiry in the classroom, I took advantage of a one-year graduate fellowship developing inquiry-based scientific instruction curriculum for fourth graders and helped implement that curriculum with teachers from a local inner-city elementary school. Since then, my teaching methods have been greatly influenced by the experimental

---

[1] Samples can be found on any of the course web page mirrors archived at http://www.tedpavlic.com/teaching/osu/.

[2] Mercurial source code repositories for course material available at http://hg.tedpavlic.com/.

scientific method. In engineering classes in particular, I work to strengthen students' understanding and appreciation of the abstract mathematical topics by highlighting how experimental deviations from theory follow from violations in model assumptions.

The spontaneous theoretical deviations that occur in the laboratory provide an opportunity to actually validate that theory using scientific inquiry. In the laboratory courses that I have taught, many students either forget or are unaware of the mathematical assumptions taken for granted by the theory being used. When students tell me that observed differences from theory are from "noise" or "parameter variations from the manufacturing process," I work with them to design and execute an experiment to make these blunt statements precise. This process typically leads them to re-examine how few assumptions are actually necessary for common engineering methods and how easy it can be to find which assumptions are violated. Moreover, by testing different hypotheses for why the deviations occur, the students sometimes correct experimental errors and restore the expected results, and they often gain valuable insight into the scientific method itself.

It is not practical to provide a hands-on laboratory experience to every student about every topic. In these cases, I find that simple first-principle-based simulation models can be distributed to students that help them with their own out-of-class inquiry. For example, it is rare that I find younger students who can gain insight from a SPICE circuit schematic. Even if they have SPICE access, there are barriers to opening the schematic and choosing the analyses to run. However, most engineering students have access to tools like MATLAB and are comfortable programming, and so I distribute scripts that implement step-by-step simulations built up from first principles and also demonstrate sample analyses[3]. Students can then examine how simple components introduced in their other classes can be built into complex systems. Furthermore, they can make and test hypotheses about system components.

Learning by Creating

Although guided instruction is an important part of post-secondary education, I believe that engineering students greatly benefit from collaborating with their peers to create new solutions to unsolved problems. As described by Freuler et al. [1], I have spent several years on the instructional staff of a program for first-year engineering students that provides them a cohesive year-long experience that introduces them to basic mechanical, electrical, and computer engineering tools and then requires teams to design, build, and present autonomous robots to complete novel tasks. This kind of tight organization between classroom instructors throughout the academic year is not practical for older students, and so I also served as a team leader in an extracurricular group that challenged these older students with similar but more advanced design and build problems. In both cases, novel design challenges were valuable examples of real engineering and helped students gain exposure to topics that are not easily covered within the classroom.

---

[3] One example is the set of DC–DC boost converter simulation codes and results available at http://www.tedpavlic.com/teaching/osu/ece327/#lab_voltreg. This example was produced at the request of a student.

References

[1] Richard J. Freuler, Michael J. Hoffmann, Theodore P. Pavlic, James M. Beams, Jeffrey P. Radigan, Prabal K. Dutta, John T. Demel, and Erik D. Justen. Experiences with a comprehensive freshman hands-on course – designing, building, and testing small autonomous robots. In *Proceedings of the 2003 American Society for Engineering Education Annual Conference & Exposition*, 2003.

[2] 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`.

[3] 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`.

[4] 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.

[5] 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.

[6] 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.

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

[8] 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.

[9] 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.

[10] 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.

[11] 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`.

[12] 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.