

# Research and Teaching Statement

Theodore P. Pavlic

## Research Interests

I am particularly interested in emergent patterns complex multi/agent 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 transdisciplinary work connects engineering [2, 6, 11, 12, 16], computer science [13–15], and behavioral ecology [8, 10]. Recent projects have included decentralized algorithms for intelligent lighting in the built environment, formal safety verification for autonomous urban vehicles, ecological rationality in solitary foragers, and resource allocation in natural and artificial swarms. Presently, I work in a social/insect laboratory where I do both empirical and theoretical studies to infer how group/level rationality emerges from the aggregated actions of individuals. This work is equally focused on bioinspired design as it is on fundamental biology. In general, I use theories from optimization, parallel computing, networks, dynamical systems, stochasticity, and formal logic to explain or design emergent system/level phenomena. Here, I describe some of my prior research successes and pending work.

### Control of Groups: Cooperation, Optimal Resource Allocation, Safety, and Swarms

Whereas my past work focused on the interactions of individuals with complex environments, my more recent work considers systems of many interacting individuals. In the engineering design cases, I am interested in decentralized implementations of collective group behavior that rely on little communication or coordination between individuals. In the behavioral analysis cases, I am interested in the individual proximate mechanisms that, when combined with the environment, lead to synergistic group behaviors.

**Cooperative task/processing networks:** When a number of agents with similar capabilities are connected on a network, each agent may share its own task/processing resources with its neighbors. Rather than depending on a centralized computational strategy to determine the best cooperation policy, I implemented a fictitious trading economy that induces individual agents to dynamically adjust cooperation frequency based entirely on locally received incentives [4, 12]. Rather than reducing communication ability of each agent and studying the “price of anarchy,” I use a novel individual/level utility function that couples the actions of neighboring individuals so that good performance is possible without explicit communication. As shown by formal theoretical analysis, certain motifs present in the network help to guarantee the stability of the decentralized algorithm to solutions that are qualitatively similar to those of a centralized load balancer. As this system exhibits non/trivial coupling between individual utility and network structure, I have presented these results in complex adaptive systems venues [7]. Moreover, although this approach is focused on flexible manufacturing systems and autonomous/air/vehicle surveillance systems, it is inspired by animal examples of cooperative breeding among unrelated individuals, and I have presented it at sociobiomimicry venues as well [9]. Several extensions of this work are possible.

**Distributed resource allocation under non/separable constraints:** Depending on the optimization criterion, it is often possible for a team of distributed agents to find optimal solutions in a decentralized fashion so long as solution-/space constraints are separable. However, realistic constraints are non/separable and induce tradeoffs between agents; the actions of one agent have to be mirrored in another in order to maintain the constraint. In these latter cases, distributed solution methods are effectively centralized due to necessary communication and coordination. With this problem in mind, I study how to replace explicit communication between agents with anonymous modifications of the shared environment. I have presented results [4, 5] that show how a decentralized bank of lights in a building can achieve multiple illumination constraints using minimum power without any explicit communication between the lights. A similar method could be used for other systems that share access to a physical resource, like distributed power generation systems. I am currently doing empirical work with *Temnothorax rugatulus* ants to determine the decentralized rules that individuals use to simultaneously regulate the intake of several colony/level macronutrients to set points. These rules can be used for bio/inspired design of non/separable distributed resource allocation algorithms. Additionally, this work may provide insight into behavioral factors of colony collapse disorder in honeybees.

**Safety verification in mixed/traffic urban environments:** In the far future, autonomous vehicles may chauffeur human occupants within a highway system with a complicated communication infrastructure that guarantees coordinated motion between the vehicles. During the transition to that future, vehicles on highways will have varying levels

of autonomy, and there may be little/to/no direct communication between them. Whereas my other work focuses on collective optimization with little coordination, here I study collective safety constraints in cases of little coordination. For this task, I have combined physical specifications of the environment with software verification methods that are normally accustomed to guaranteeing safety constraints in discrete/time and discrete/space systems [13, 14]. Prior work that uses similar methods focuses on vehicles that can be specified with homogeneous parameter sets. Instead, I allow for heterogeneity, and the resulting software solutions are less conservative and thus provide more comfortable performance while still ensuring safety.

**Stochastic strategies for the maintenance of teams within swarms:** My most recent work focuses on swarms of independent agents where each agent must allocate itself to one of a number of tasks. For example, a swarm of robots on Mars may be tasked with collecting rocks of different sizes and carrying them back to a central station for sophisticated analysis. Some rock loads will require more team members to carry than others. So to maximize the chances of finding many interesting specimens, it is desirable to allocate only as many agents to carrying as is necessary. Rather than relying on communication between the agents, each agent may be viewed as an entity of a gas which attaches and detaches from the object according to certain probabilities that are designed *a priori*. Just as chemical reaction networks tend to concentrations that are well/described by kinetic theory, agents will tend to teams with sizes that are determined by “programming” the “reaction” probabilities. By comparing to empirical data, my colleagues and I have shown that such stochastic/hybrid/system models are good fits to the phenomenon of collective food transport observed in ants, which can form appropriately sized teams around food items and carry those items back to their home nest [2]. The allocation of agents around loads is not unlike classical adsorption processes on surfaces in chemistry that have equilibrium surface allocations that depend on the pressure of the gas. For robotic applications, we have introduced enzyme/inspired robot/to/robot interaction rules that ensure equilibrium surface allocations are independent of such environmental parameters. The result is a decentralized stochastic algorithm for formation of teams with stable and predictable sizes that do not vary with total number of robots or geometry of the space [16]. Ongoing work involves the use of such methods for nano/scale medical applications.

### **Solitary Agents and their Environments: Optimization and Irrationality**

I summarize some key results from my past work, which focused less on the complex systems comprised of groups of individuals and more on the systems formed by a single individual and its complex environment. Most of these results are connected to foraging theory that describes how solitary animals make choices in an uncertain environment. 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). Good designs 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 are too specific to be directly applied to engineering problems, and these engineering problems often use different optimization criteria than those used in ecological modeling of natural selection. So I generalized the intuitive analysis methods used in behavioral ecology to apply to a broad class of objective functions more amenable to the engineered ecology of mobile robotics [3, 11]. Moreover, I simplified models of speed/dependent sensor accuracy so that traditional models of speed–accuracy tradeoffs in predation could be used to design optimal speed–task combinations for mobile autonomous agents [6].

**Ostensibly irrational decision making:** Behavioral ecologists and human psychologists have often shown that human and non/human animals make irrational choices in some laboratory environments. Rather than assuming that natural selection has failed to adapt these animals to their environment, I have shown that impulsiveness in the laboratory can be the effect of simple behavioral heuristics that induce optimal behaviors only when simultaneous encounters with mutually exclusive options are rare [8]. So animals may simply be maladapted for the operant laboratory. Similarly, by using more realistic modeling approaches, I have shown that sunk/cost effects (i.e., where past costs of a task lead to increased persistence in continuing that task) can also be optimal in some harsh environments [10].

## 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<sup>1</sup> 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<sup>2</sup> 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 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.

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

<sup>2</sup>Mercurial source code repositories for course material available at <http://hg.tedpavlic.com/>.

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<sup>3</sup>. 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.

---

<sup>3</sup>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](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. URL <http://search.asee.org/search/fetch?url=file%3A%2F%2Flocalhost%2FE%3A%2Fsearch%2Fconference%2F27%2FAC%25202003Paper513.pdf>.
- [2] Ganesh P. Kumar, Aurélie Buffin, Theodore P. Pavlic, Stephen C. Pratt, and Spring M. Berman. A stochastic hybrid system model of collective transport in the desert ant *Aphaenogaster cockerelli*. In *Proceedings of the 16th ACM International Conference on Hybrid Systems: Computation and Control*, pages 119–124, Philadelphia, PA, April 8–11, 2013. doi:10.1145/2461328.2461349.
- [3] Theodore P. Pavlic. Optimal foraging theory revisited. Master’s thesis, The Ohio State University, Columbus, OH, 2007. URL [http://rave.ohiolink.edu/etdc/view?acc\\_num=osu1181936683](http://rave.ohiolink.edu/etdc/view?acc_num=osu1181936683).
- [4] 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?acc\\_num=osu1281462093](http://rave.ohiolink.edu/etdc/view?acc_num=osu1281462093).
- [5] Theodore P. Pavlic. Physical stigmergy for decentralized constrained optimization: an intelligent lighting example. In *Proceedings of the 4th International Conference on Cyber-Physical Systems (ICCPs 2013)*, Philadelphia, PA, April 8–11, 2013. Poster abstract.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] Theodore P. Pavlic and Stephen C. Pratt. Superorganismic behavior via human computation. In Pietro Michelucci, editor, *Handbook of Human Computation*, chapter 74, pages 911–960. Springer, 2013. doi:10.1007/978-1-4614-8806-4\_74.
- [13] 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>.
- [14] 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.
- [15] Theodore P. Pavlic, Sai Prathyusha Peddi, Paolo A. G. Sivilotti, and Bruce W. Weide. Poster abstract: getting out of the way – safety verification without compromise. In *2012 IEEE/ACM Third International Conference on Cyber-Physical Systems (ICCPs 2012)*, Beijing, China, April 17–19, 2012. doi:10.1109/ICCPs.2012.51.

- [16] Theodore P. Pavlic, Sean Wilson, Ganesh P. Kumar, and Spring Berman. An enzyme-inspired approach to stochastic allocation of robotic swarms around boundaries. In *Proceedings of the 16th International Symposium on Robotics Research (ISRR 2013)*, Singapore, December 16–19, 2013.