

# Connecting Cognitive and Physical Worlds with Dynamic Cost Function

## Definition

**Jamie Lennon**

University of Maryland  
Department of Aerospace Engineering  
College Park, MD 20742  
lennon@aic.nrl.navy.mil

### Abstract

Our goal is to mesh the symbolic reasoning capabilities of a cognitive model with the constrained optimization possibilities inherent in optimal controls. We plan to develop and test such a system for several different dynamical models in environments of differing certainty and differing efficiency requirements.

### Problem Description

We desire an intelligent robot that can reason, plan ahead, and make decisions based on its goals, its environment, and the desires of any teammates it might have. A cognitive model capable of symbolic inference can, with a large enough rules set and world information, do this for us. It can make judgment calls not only about what actions it should undertake, but *how* it should perform them - quickly, efficiently, cautiously, aggressively. Furthermore, given feedback, it could fine-tune these behaviors. They exist on a spectrum, after all. "Less aggressively" is not necessarily the same as "cautiously." A "less aggressive" behavior can still be far more reckless than the "least cautious" behavior.

These behaviors, illustrated with possibly emotionally loaded words, are simply stereotypes of classes of trajectories. An "aggressive driver" is one who drives quickly, leaves little clearance between his vehicle and others, and engages in rapid accelerations between and within lanes. We interpret these dynamic characteristics of the vehicle's motion as aggression, regardless of the actual emotional state of the driver. (This is different from the research presented in (Velasquez 1999), where simulated emotions actually trigger reactive behaviors). A "cautious driver" would have different characteristics. He might still drive very quickly on a highway, but would leave larger clearances between his car and others, and would maintain a more constant speed with fewer and more gradual accelerations. How can we elicit these characteristics in position, velocity, and acceleration in a variety of domains?

---

In an aerospace context, for example, we may be dealing with a military air vehicle on a mission through hostile territory. A reconnaissance mission would be flown in a very different manner than a bombing run.

Most existing path planners (Latombe 1991) would need to augment their states to handle more than just path (position only), and be able to compute full trajectories that are optimal with respect to our desired characteristics. The states are likely to quickly become unmanageable, and the resolution on the solution may be poor if discretization is decreased as a result of state explosion. Happily, the field of optimal controls (Kirk 1970) uses the calculus of variations to compute continuous, full-state trajectories. Additionally, the technique allows constraints to be placed on the trajectories. These can include, for example, system dynamics (including real-world concerns such as motor saturation), resulting in a path that the agent is guaranteed to be able to follow. They can also include a multi-term cost function that will allow us to elicit our various behaviors. This is an offline planning technique that is both mathematically rigorous and provably optimal, at the expense of computational complexity. It is therefore quite different from the work done in (Santamaria and Ram 1997), which is for fast, reactive behaviors that do not have a global perspective.

Our optimal control cost function contains three terms: fuel use, clearance from obstacles, and time. Prior to this work, relative cost function weights were set in an ad hoc fashion, often determined experimentally. While we will develop our basic behavior-eliciting weights experimentally, we are working toward rules to dictate weight adjustment "within" as well as "between" behaviors. Perhaps the most analogous work is the hybrid dynamical systems approach taken by (Aaron et al 2002). In this work on low-level navigation, a dynamic "comfort level" is used to adjust the weighting parameters of repulsor fields surrounding environmental obstacles. This is another reactive method, and it does not attempt to calculate or minimize any costs.

Selection of the weights returns us to the cognitive model. Systems like Soar (Newell 1990), EPIC (Kieras and Meyer 1997) and ACT-R (Anderson and Lebiere 1998) all try to model not only the end result of human cognition,

but also the process by which those results are reached. Not only can high-level planning be accomplished, but in theory the implementation should be very "human-like," making user interface more intuitive. This is very desirable if we have a human user critiquing trajectories returned by the system, who may give feedback such as, "No, be more careful." We chose to use ACT-R, although the other architectures could be adopted as well. In ACT-R, procedural rules fire in the presence of certain "chunks" of symbolic declarative memory. This is a serial system, using the bottleneck as a point of coordination among different cognitive modules. This makes it a very attractive option for implementing on a hardware system. Finally, ACT-R also has an extension, ACT-R/S, which supports spatial awareness, which we may leverage to extend the work (see Proposed Plan, below).

Given the agent's current goals, its environment, its state, and possibly other factors, the cognitive model is well-equipped to decide both the actions it needs to take and how it should perform them. "How" is linked to a weight set, which is then passed to the optimal controls trajectory planner. The trajectory planner then returns a trajectory that displays the dynamic characteristics sought by the cognitive model.

### Proposed Plan of Research

Our first step has been to characterize three "extreme" weight sets, where one of the three weights dominated the cost function. Using a numeric solver (Matlab's `bvp4c`) and a simple dynamic model, we generated sample curves for one and for five obstacles. Certain experimental parameters - for instance, the height and extent of the potential field used to create a penalty for approaching an obstacle - have been determined. Single-obstacle results generalized well to a multi-obstacle field.

Our next steps in the optimal controls side of the research are to extend the problem for free end time, to characterize behaviors resulting from less "extreme" weight sets, and then to group the weights into behaviors, so that they can be adjusted by the cognitive model.

When the optimal control trajectory planner is well-characterized, it will be time to begin programming the cognitive model. This will include extensive rules for recognizing features in the environment and how they impact goal acquisition. We also seek to add a domain-specific understanding of fuel economy: space satellite maneuvers, for example, must be more fuel-thrifty than easily refueled aircraft or ground robots. Further, we want the cognitive model to be able to collaborate with a human user, making adjustments to the trajectory based on the human's inputs.

Finally, we would like to leverage some of the perspective-taking work done at the Naval Research Laboratory (Trafton et al. 2004), enabling the agent to recognize a "hostile" obstacle and then to take cover behind a "good" obstacle.

### Metrics

We will perform a complexity analysis of the algorithms used to compute the optimal trajectories. We expect that worst-case examples will be too costly to run in real-time, but hope that average cases can be run on the order of seconds. Replanning in the field becomes a more tenable position in that case.

To compare our work with others', we will develop at least two dynamic models: a wheeled ground robot and an autonomous spacecraft. We will test the system across environments with varying degrees of uncertainty, as well as cases where fuel conservation is more or less important. We expect to find crossover points where, on one side, a regime of reactive behaviors is more successful, and on the other, the optimal trajectory planner is superior.

### Acknowledgements

would like to thank my advisor, Dr. Ella Atkins, for her continuing guidance and direction. This work was performed at the Naval Research Laboratory under funding from the Office of Naval Research under work order N0001404WX30001.

### References

- Aaron, E., Sun, H., Ivancic, F., and Metaxas, D. 2002. A hybrid dynamical systems approach to intelligent low-level navigation. In *Proceedings of Computer Animation 2002*: 154-163.
- Anderson, J.R. and Lebiere, C. 1998. *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum.
- Kieras, D., and Meyer, D.E. 1997. An overview of the epic architecture for cognition and performance with application to human-computer interaction. *Human Computer Interaction* 12: 391-438.
- Kirk, D.E. 1970. *Optimal Control Theory: An Introduction*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Latombe, J.C. 1991. *Robot Motion Planning*. Norwell, MA: Kluwer Academic publishers.
- Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Ram, A., Arkin, R., Boone, G., and Pearce, M. 1994. Using genetic algorithms to learn reactive control parameters for autonomous robotic navigation. *Adaptive Behavior* 2(3): 277-304.
- Santamaria, J.C. and Ram, A. 1997. Learning of Parameter-Adaptive Reaction Controllers for Robotic Navigation. In *Proceedings of the World Multiconference on Systemics, Cybernetics, and Informatics*.
- Trafton, J.G., Schultz, A.C., Perzanowski, D., Bugajska, M.D., Adams, W., Cassimatis, N.L., and Brock, D.P. 2004. Children and robots learning to play hide and seek. *Cognitive Systems Research*, under review.
- Velasquez, J.D. 1999. An emotion-based approach to robotics. In *Proceedings of the 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems*: 235-240.