

Neurocognitive Patterns: Using Brain, Behavior, and Context to Infer User Intent

Webb Stacy

Aptima, Inc.,
12 Gill Street, Suite 1400, Woburn, MA 01801
wstacy@aptima.com

Abstract. Neurocognitive Patterns is a system that will offer execution options to users as soon as they form an intention to act. It will accomplish this by combining neural signals, user behavior, and contextual knowledge to determine when a user has a goal, and what that goal is. Because it will leverage the user's neural signals and behavioral history, the options it will provide to the user will be available quickly. Because it will leverage real-time contextual and background knowledge, its estimates concerning the user's goal will be accurate. Our initial target domain is UAV operators, but we expect it will be of use to other military decision-makers in Command and Control settings. We also expect that Neurocognitive Patterns will be a useful tool in Cognitive Neuroscience in general for interpreting neural signals in the presence of salient contextual information.

Keywords: Neural Signals, Behavioral Measures, User Intent, Using Contextual Information.

1 Introduction

Modern military operations are increasingly complex and increasingly electronic, and this often has the effect of increasing operator workload. System developers have incorporated many technological developments, but they have not kept up with our increasing understanding of the cognitive capabilities of the operators themselves. Though there is a recently-renewed, increased interest in establishing a system's usability and utility during development in order to mitigate the increasing technology-workload mismatch, there are more dramatic possibilities for improvement.

One such possibility is to use neural signals from the operator as the ultimate unobtrusive form of input: the operator merely has to think about something to make it so. A solution like this uses neurocognitive signals from operators to determine their intended actions so that computers can automatically develop plans to execute those actions. Unfortunately, traditional brain-computer interface approaches are not suitable. They use stereotypic and non-specific neural signals, do not scale with task complexity, and cannot handle changing task environments and contexts. What is needed is a way to leverage the neural signals of the operators' natural cognitive processes that can adapt to the dynamics of modern complex military environments.

An immediate difficulty is that, with Cognitive Neuroscience's current understanding of the meaning of neural signals from the brain, it is extremely difficult to provide general interpretations of underlying cognitive activities (i.e., of specific "thoughts"). There has been progress in this regard, of course. For example, Luu and his colleagues [1] identified the neural signature of visual intuition from dense-array electroencephalography (dEEG) data, but it is currently not possible to know the content or correctness of that intuition. As another example, Researchers at CMU identified functional Magnetic Resonance Imaging (fMRI) correlates of semantic components of simple nouns such as "apple." [2] Both studies represent impressive advances in Neuroscience, but they are a far cry from identifying user intent in a complex task environment. It follows that a solution that provides useful neurocognitive control of human-machine systems will need to supplement neural data from the operator with additional contextual knowledge and situational information.

2 Towards a Solution

A system currently under development called Neurocognitive Patterns will provide such a solution. It is a system that will combine neural data, user behavior, and contextual knowledge to determine when a user has a goal, and what that goal is. It will plan ways to accomplish the goal, present them to the user, and execute the plan that the user chooses.

As illustrated in Figure 1, Neurocognitive Patterns will collect the neural signals and user behavior logs and will combine these with mission context, situational updates, and background knowledge in long term memory (LTM) to infer the user's goal. Neurocognitive Patterns will then automatically plan one or more ways to accomplish the user's inferred goal, and will present them to the user for approval. Once approved, it will interact with the workstation to execute the plan. Users will wear equipment to collect neural signals such as dEEG but will otherwise engage with their workstations in the normal manner.

Emerging research in Neuroscience has been able to identify neural correlates associated with a user intending to do something, especially in the motor domain, though it has proved difficult to determine exactly what that intention might be. The key idea in Neurocognitive Patterns is to leverage those neural correlates of intent to identify *when* the user intends to take action, and to use high-level contextual knowledge from sources like the mission, operator behavior, and previously-established context to determine *what* the user intends to do.

The identification of an "intent" signature in the operator's neural signals will indicate that the operator would like to accomplish a newly-formed goal. This will prompt Neurocognitive Patterns to identify the nature of that goal.

A naïve approach to identifying user goals will not solve the problem. For example, simple aggregation of user behaviors like keystrokes or mouse clicks is rarely productive as the meaning of those interaction behaviors is obscure without some understanding of context, such as the computer applications that are receiving those keystrokes and mouse clicks, and the overall mission of the user. Clearly, context must be part of the equation.

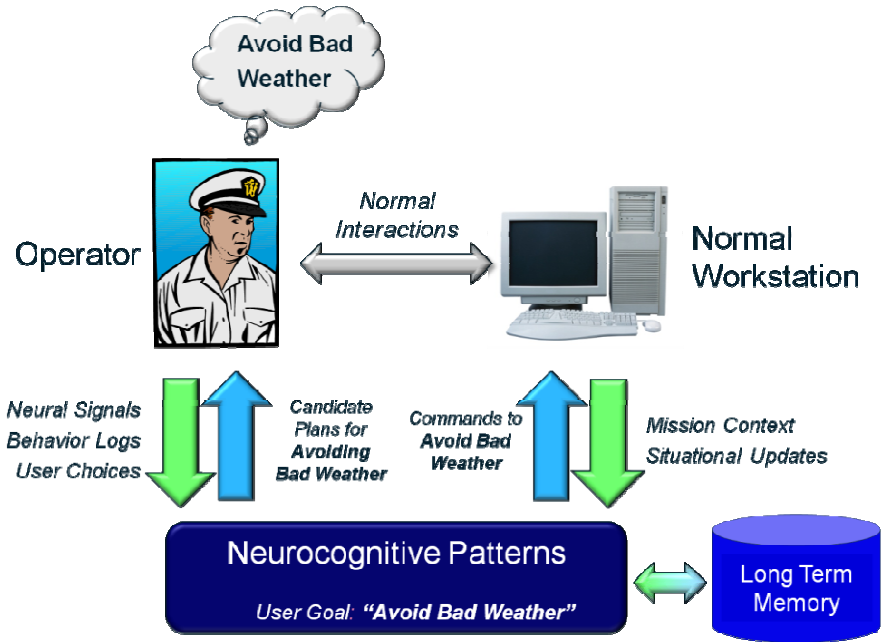


Fig. 1. Neurocognitive Patterns will speed decisions and allow users to make them easily and unobtrusively

Neurocognitive Patterns will incorporate context via an approach called Top-Down/Bottom-Up (TDBU) processing. In TDBU processing, hypotheses (“top down”) about the meaning of sensor signals (“bottom up”)—neural signals, keystrokes, mouse clicks, and possibly other sensors—are generated and then tested against actual sensor data. The hypotheses originate from contextual knowledge and from previously-perceived aspects of the situation.

The TDBU approach is often effective for machine perception tasks. For example, a machine vision task might be to identify visual objects from a camera image. The TDBU approach might use the fact that the setting is indoors in a kitchen, and that a counter, a sink, and a floor have already been identified, to generate and test the hypothesis that a certain set of pixels on the counter represents a microwave oven. In Neurocognitive Patterns, the task is to identify operator intent, and the TDBU approach might use facts like this:

- The mission is for an unmanned aerial vehicle (UAV) operator to identify possible weapons caches without being spotted.
- A group of people has been identified on the ground.
- The operator’s previous behavior indicates that he has been trying to get a closer look by flying at a lower altitude.

to generate the hypothesis that the operator intends to take action to make the vehicle less visible. This hypothesis will then be tested against the operator “sensor signals.” Neurocognitive Patterns will consider the operator’s neural signals and their recent interaction with their computer system.

Neurocognitive patterns will thus treat the interpretation of neural signals as a top-down/bottom-up machine perception problem. The effect will be a system that provides the opportunity for automated planning assistance to operators without the operators having to explicitly indicate their intent. This will provide an unprecedented level of system efficiency and effectiveness. Further, the general approach of treating the interpretation of neural signals as a machine perception problem can find wide applicability in the field of Cognitive Neuroscience in general.

3 Overview of the System

Neurocognitive Patterns will have four main modules, shown in Figure 2, which will enable the system to detect the nature and timing of user goals, to automatically develop plans to accomplish them, to present those plans to the user, and to execute the user's choice. The process will start with the Goal Identification Module, which will use neural signals to detect the timing of when a user intends to do something, will draw on a variety of contextual and background information to determine the set of goals the user is likely to be pursuing, and will examine this set of goals in relation to the user's recent behavior to deduce their current most likely intent. Given those candidate goals, the Goal Planning Module will automatically construct a plan to accomplish them, and will present the plans to the user via the Graphical User Interface (GUI). If there is a plan with which the user concurs, they will indicate it using the GUI, and that plan will be executed by the Goal Execution Module. If there is no plan with which the user agrees, the user will have the option to indicate whether the goal was wrong or whether the plan was not acceptable, and Neurocognitive Patterns will learn from this feedback.

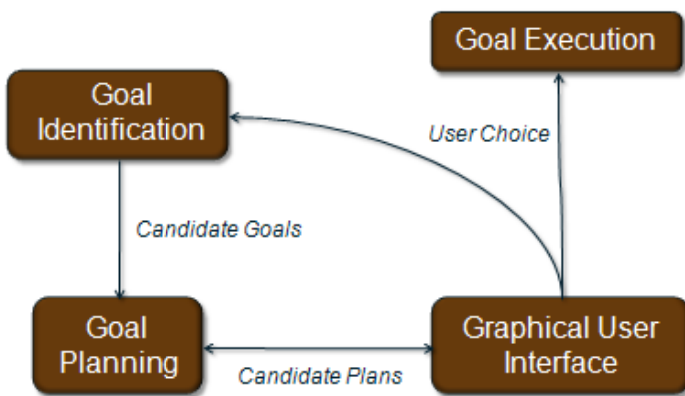


Fig. 2. Top-level view of Neurocognitive Patterns

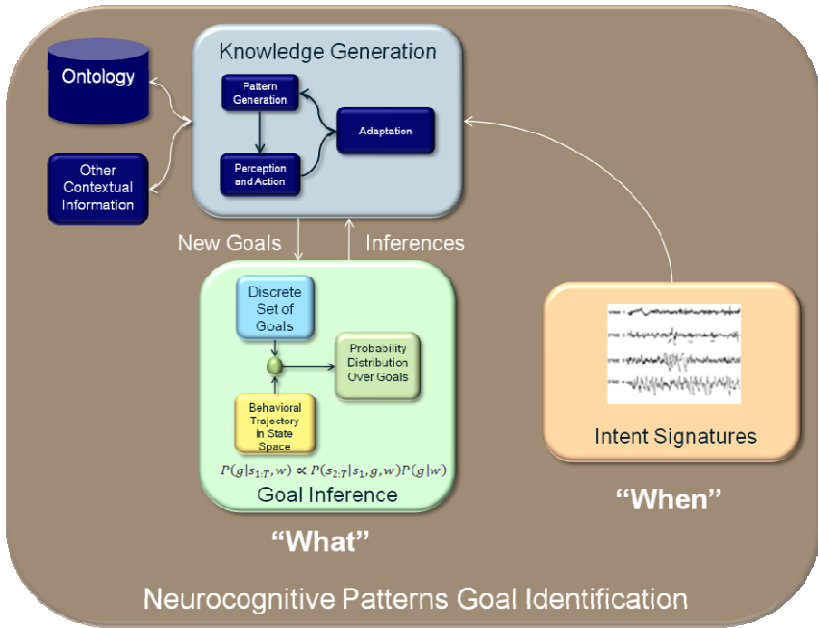


Fig. 3. Components of the Goal Identification Module

As Figure 3 shows, the Goal Identification Module for Neurocognitive Patterns is heart of the system that interprets neural signals. It will accomplish the timely detection of the user’s intent to act by finding signatures of intent (the Intent Signatures Component), and will determine the user’s goals using two components, the Knowledge Generation Component and the Goal Inference Component.

The Knowledge Generation Component, based on a system called Cognitive Patterns [3], will reason over available context and background information to narrow down the set of probable user goals, together with associated user behaviors (the state space) and their probabilities. When confronted with a novel situation, as represented in the contextual and background information, it will have the capability of generating a potential new goal and probabilistic state space, giving it the capability of behaving sensibly in the face of what would otherwise be puzzling and uncertain information.

Cognitive Patterns is an adaptive system based on theoretical concepts from Cognitive Science, such as Conceptual Blending [4], dynamically interpreted perceptual symbol systems [5], and a strong theoretically and empirically justified model of the human eye-brain visual system [6].

The Knowledge Generation (KG) module adds a layer of flexible, temporarily-generated knowledge to the reasoning system. Because it can generate knowledge easily, it can accommodate anomalies and novel events when necessary. The KG component comprises three major internal subcomponents: Pattern Generation, Perception and Action, and Adaptation, which together enable it to create situationally-relevant abstract patterns, to match sensory input to a suitable abstract pattern in a multilayered TDBU fashion (similar to the mechanisms used for visual perception in the brain), and generate new abstract patterns corresponding to “intuitive leaps.”

Armed with this set of goals and associated information, the Goal Inference (GI) Component, based on a system called MIMIC [7], will use a mathematical model of user intent to determine the probability distribution over that set of goals, given the user's recent (bottom-up) behavior, thus completing the TDBU strategy.

The GI Component will take as input a model of priorities and goals which will be provided by the KG component will provide. It will attempt to infer (1) the *goals* that are most likely to have generated the human's actions and (2) the most likely *priorities* for the person, of which each goal is a concrete realization. It will do this by combining the model with the user's behavioral history using Bayesian inference to invert this causal relation to infer the most probable goals and priorities that gave rise to the action sequence [8].

To perform this inference, we compare all possible goal hypotheses against each other in terms of how well they explain observed action sequences. Specifically, for each goal hypothesis, it will compute the posterior probability, given observed Actions and the Environment:

$$\begin{aligned} &P(\text{Goal} \mid \text{Actions, Environment}) \\ &\propto P(\text{Actions} \mid \text{Goal, Environment})P(\text{Goal} \mid \text{Environment}) \end{aligned} \quad (1)$$

Goals with a sufficiently high probability will be considered to have been identified, and these will be passed to the Planning Module to determine the options to be presented to the user.

4 Conclusion

The field of Cognitive Neuroscience is making progress at an unprecedented rate. This new knowledge will lead to dramatic improvements both in the kinds of automated support we can provide our warfighters and in the technology for interpreting neural signals in meaningful and useful ways. We hope that this effort represents a contribution in both directions.

Acknowledgements. I would like to acknowledge the support of the Office of the Secretary of Defense and the Office of Naval Research, especially LCDR Joseph Cohn, Ph.D., for strong support of this effort. I would also like to acknowledge strong collegial encouragement of Alexandra Geyer, Ph.D., and Brian Riordan, Ph.D., both of Aptima. The opinions expressed here are those of the author and do not necessarily reflect the official position of the sponsors or the Department of Defense.

References

- [1] Luu, P., Geyer, A., Fidopiastis, C., Campbell, G., Wheeler, T., et al.: Reentrant Processing in Intuitive Perception. PLoS ONE 5(3), e9523 (2010), doi:10.1371/journal.pone.0009523
- [2] Just, M., Cherkassky, V., Aryal, S., Mitchell, T.: A Neurosemantic Theory of Concrete Noun Representation Based on the Underlying Brain Codes. PLoS ONE 5(1) (2010), DOI: 10.1371/journal.pone.0008622

- [3] Stacy, W., Cohn, J.V., Geyer, A., Wheeler, T.: A Cognition-Based Control System for Autonomous Robots. In: Proceedings of the Human Factors and Ergonomics Society 54th Annual Meeting, vol. 54, pp. 1473–1477 (2010)
- [4] Fauconnier, G., Turner, M.: The Way We Think: Conceptual Blending and the Mind's Hidden Complexities. Basic Books, New York (2002)
- [5] Barsalou, L.W.: Abstraction as dynamic interpretation in perceptual symbol systems. In: Gershkoff-Stowe, L., Rakison, D. (eds.) Building object categories. Carnegie Symposium Series, pp. 389–431. Erlbaum, Mahwah (2005)
- [6] Kveraga, K., Ghuman, A.S., Bar, M.: Top-down predictions in the human brain. *Brain and Cognition* 65, 145–168 (2007)
- [7] Riordan, B., Bruni, S., Schurr, N., Freeman, J., Ganberg, G., Cooke, N.J., Rima, N.: Inferring user intent with Bayesian inverse planning: Making sense of multi-UAS mission management. In: Proceedings of the 20th Behavior Representation in Modeling and Simulation Conference (BRIMS), Sundance, Utah (2011)
- [8] Baker, C.L., Saxe, R., Tenenbaum, J.B.: Action understanding as inverse planning. *Cognition* 113(3), 29–349 (2009)