Patterns of Emotion Driven by Affect State and Environment

An architecture for a visualized, independent, autonomous, learning agent

Paul G. Joseph and Haim Levkowitz
Computer Science Department
University of Massachusetts
Lowell, MA., USA
pjoseph@gmail.com, haim@cs.uml.edu

***Abstract*—The neuroscientist Jaak Panksepp posits the existence of seven physical systems in the mammalian brain, that when simulated in the laboratory, result in emotions identical to known emotions. These seven systems are SEEKING, RAGE, PANIC, FEAR, LUST, CARE, and PLAY. From the perspective of mathematics, these emotions form the dimensions of a phase space in which every emotion is located and modeled as a dynamical system. Inputs from outside the system interact with internal values and inform the dynamical system, resulting in a simulated affect that in turn drives observed patterns of emotion. This paper discusses a framework built to explore this premise. As a first cut it uses a linear seven-variable dynamical system to drive a search heuristic for a utility-based reflex agent. The results indicate that even this simple linear model shows the basic patterns of emotion observed in mammals. These patterns of emotion facilitate discovery and decision. It provides an independent "always on, automatic" discovery and decision system capable of adaptive goal setting, which works without the need for significant additional cognitive analysis. It offers the basis of a framework that is extendable to include additional sensory information, learning, memory/persistence, and an independent cognitive system.**

Keywords - patterns; affect; emotions; Panksepp; autonomous agent; mammalian emotion.

I. INTRODUCTION

Rosalind Picard [30] used the terms affective computing systems and emotion-oriented computing to refer to systems that consider emotions. Emotion handling is fundamentally important, and a significant amount of research in this area has been carried out in connection with real-life non-simulated pervasive computer systems [4][5][12][23][27][28][33][38]. To date, in the field of software in general, the bulk of the work in this area has been to examine patterns of emotion (facial expressions, bodily movements, language) to infer internal emotions.

In this paper, however, we seek to lay out the basis for a complementary approach—one that models neural systems posited to exist in mammalian brains in order to generate and manage affect. We use dynamical system theory to model these systems and to generate internal affect states that together resulting in an overall "emotion" that in turn drives behavior.

The premise of this paper is that "meaning" substantially involves an emotional evaluation of an object based on our interaction with the object. We use affect values generated by a dynamical system as the basis of a heuristic used by a utility agent to navigate its environment. The premise is that with this approach, it may be possible to attach affect to an object and, thereby, give a computer-based agent a sense of meaning of the objects it is working with.

This, in turn, allows the agent the ability to perform adaptive goal setting. Depending on circumstances and its current needs, the agent will demonstrate appropriate patterns of emotion and act based on these patterns, e.g., if FEAR is high, it will run away from the object causing the fear. If, however, SEEK predominates, then it will seek and if hungry it will specifically seek for nourishment. In short, goals are adaptive and based on the current emotional state.

Though this research explores a model of the mammalian affective system, its goal is not to explore emotions per se. In addition, though it uses the framework of a game to explore the model, the purpose of this work is neither to develop better gaming theory nor to develop a more robust algorithm for the game. Rather, the purpose of this research is to see if, using the model, a computer-based agent can develop affect with respect to other entities it encounters based on interactions with them, and based on this affect, generate appropriate emotion patterns that produce adaptive goal setting, discovery, and decision.

II. BASIS OF APPROACH

The basis of the approach is to model findings from neuroscience using dynamical systems and to use the resulting model as the search heuristic for a utility agent that allows the agent to develop affect for objects in its environment. Subsequently, this affect generates emotional patterns that result in adaptive goal setting and behavior.

 ISBN: 978-1-61208-158-8

A. Neuroscience: Empirically Based Fundamental Emotions

Dr. Jaak Panksepp, a neuroscientist from Washington State University proposed the existence of seven core emotional systems, because they "generate well-organized emotion sequences that can be evoked by localized electrical stimulation of the brain" [29]. We provide a very brief sketch of the Panksepp model (which draws directly from Panksepp, 1998):

SEEKING system: a network that promtes survival activities, making creatures intensely interested in exploring their world, and to get excited when they get what they desire. Neuroanatomically they correspond to the major self-simulation system that courses from the midbrain up to the cortex.

RAGE system: works in opposition to SEEKING, and mediates anger. It is aroused by frustration and attempts to curtail freedom of action. It parallels the trajectory of the FEAR system.

FEAR system: probably designed during evolution to help reduce pain and possibility of destruction. When stimulated intensely, it leads creatures to run away; when stimulated weakly, it causes the creature to freeze.

PANIC system: governs social attachment emotion—specifically related to absence of maternal care when the creature is a baby.

Additionally, there are three "special-purpose" emotions. Panksepp calls these emotions "more sophisticated, special purpose, socioemotional systems that are engaged at appropriate times in the life of all mammals." Currently far less understood than the other four, these are:

LUST system: involves sex and sexual desire, evolved for species propagation

CARE system: maternal love and caring

PLAY system: produces the "roughhousing play of all young animals and humans" at some stage in their development to facilitate learning.

There is by no means universal agreement or acceptance of Panksepp's model. Rebuttals and alternate formulations [1][2][6][7][8] have been proposed. Chiefly these differ in the number and kind of basic emotions. This paper does not seek to prove the correctness of Panksepp's particular formulation. Rather it leverages the fact that one can implement any such model compactly and usefully as a dynamical system.

A key feature of Panksepp's model is its empirical basis for "fundamental" emotions, which distinguishes it from numerous other systems that also claim to have identified so-called "fundamental" emotions. If the physical circuitry for any given emotion does not exist in the brain, then this emotion is a combination of those emotions that do have physical representation.

Thus, for example, "romantic love" would be viewed as some combination of (perhaps) CARE, PLAY, SEEK, and LUST; resentment or indignation as combinations of RAGE, CARE, SEEK. The SEEK response, associated with a large release of dopamine and a sense of well-being would cause emotions like "joy", "satisfaction", "contentment", "enjoyment", etc., variations of the theme of gratification of the SEEKING system and the attainment and material benefit of the goal being sought.

Greene [18][19] describes his theory that the human brain has two modes of responding to situations. One, a set of efficient automatic responses driven by emotions, the other a manual mode used in response to non-standard situations that involves significant cognitive evaluation. LeDeux [25] posits that the former typically operates in timeframes of about 10 ms, while the latter in timeframes closer to 500ms. The non-cognitive bias in the Panksepp based dynamical system search heuristic proposed above allows for an "automatic, always on and very quick" response corresponding to that described by both Greene and LeDoux.

B. Mathematics: Dynamical Systems

At the heart of the framework is the dynamical system model. See [37], for example, for additional details on dynamical systems. Dynamical systems have modeled a wide range of systems including marriage [17] and emotional development [3]. In particular, Lewis [26] used it to bridge neurobiology and emotion theory while Scherer, [35] modeled emotions as emergent processes. However, neither used these approaches for software agents, nor used a model such as the Panksepp model, whereas here, the dynamical system explicitly models in a software agent, the Panksepp systems namely: SEEKING, RAGE, FEAR, PANIC, LUST, CARE, and PLAY.

As per Panksepp's description of these basic emotions, agents when hungry will SEEK food; depending on age, they may PLAY or if they are parents, may CARE or LUST if they are young adults, and do all this using SEEK. Eventually LUST could result in progeny. If a hostile agent is in the vicinity, then the agent experiences FEAR. If trapped, it experiences RAGE. If a child is at a distance from its parent or its parent is not available then it experiences PANIC.

An agent's current emotional state is its existing state modified by incoming changes in affect caused by the meaning of the current objects in its vicinity (defined by how they affect the agent's current emotional state).

As mentioned earlier, Panksepp's model is qualified, and is currently the subject of considerable debate in the affective neuroscience community. However, this paper takes the pragmatic approach that it does not really matter if eventually it is found that there are four mechanisms or ten mechanisms—the important finding is that there is a relatively small number of them and that they have both a physical basis and a readily understood meaning.

We distinguish this approach from other AI approaches, such as neural nets where the meaning of the nodes is not well defined or known. This clear definition of the variables and their meaning in turn allows for an intuitive understanding of system emotion. In addition, we distinguish the model from "Braitenberg architectures" [24] in that here, the basis is an actual, specific, neuroscience model—that of the mammalian brain.

 ISBN: 978-1-61208-158-8

III. IMPLEMENTATION

The Pacman game [9] offers a convenient framework for a preliminary exploration of the approach suggested in this paper. Pacman used for its search heuristic, the systems identified by Panksepp, modeled as a dynamical system. The heuristic is coded to make Pacman a utility reflex agent [34] aiming to maximize its overall well-being with actions based on the current location of the ghosts and food. We initialized Pacman's dynamical system with random values between a minimum of 0 and a maximum of 100. Additionally, Pacman had other properties, such as gender, health, age, and strength fixed at "birth," i.e., initialization.

As a first cut, certain relationships are hard-coded into the model. This hard coding is analogous to behavior Pacman was "born" with, i.e., a result of evolution. For example, the model hard coded Pacman's SEEKING emotion as inversely proportional to distance to food, and FEAR to be inversely proportional to the distance to the predator (ghost). In this first cut, we arbitrarily chose the proportionality values to be 1/1000.

Later models will explore if Pacman can learn these distance relationships instead of simply being "born" with them., the learning being assisted by a combination of persistence of emotions with respect to encountered objects and information shared globally across all instances of Pacman.

As a first cut, Pacman used the following simple linear, additive, dynamical system model, with each emotion being described as some linear combination of existing emotions:

pacman.SEEKING += 1000/foodDistance + 1000/pacman.AGE + pacman.HUNGER + pacman.RAGE + pacman.PANIC + pacman.FEAR + pacman.PLAY + pacman.LUST

pacman.FEAR += 1000/ghostDistance + health + 1000/strength + pacman.RAGE + pacman.PANIC

pacman.RAGE += 1000/ghostDistance + 1000/health + 1000/strength + pacman.FEAR + pacman.RAGE + pacman.PANIC

pacman.PANIC += 1/ghostDistance + 1000/health + 1000/strength + pacman.FEAR + pacman.RAGE

pacman.CARE += 1000/foodDistance + 1000/ pacman.RAGE + 1000/pacman.FEAR + pacman.PLAY + pacman.AGE + health + gender

pacman.PLAY += 1000/foodDistance + 1000/pacman.RAGE + 1000/pacman.FEAR + 1000/pacman.AGE + health

pacman.LUST += 1000/foodDistance + 1000/pacman.RAGE + 1000/pacman.FEAR + pacman.PLAY + 1000/pacman.AGE + health

The emotional state of Pacman was then calculated as the simple summation of the "positive" dimensions SEEKING, CARE, PLAY, and LUST , less the "negative" dimensions of FEAR, PANIC, and RAGE. i.e..:

pacman.EMOTION = (pacman.SEEKING + pacman.CARE + pacman.PLAY + pacman.LUST) - (pacman.FEAR + pacman.PANIC + pacman.RAGE)

A simple loop recalculated the model's equations continuously. Figure 1 shows a typical screenshot from the running game. Figure 2 shows values for each of the seven emotions from a sample run of Pacman—the X-axis is the iteration number, the Y-axis the affect value.

The ghost in the game that hunts Pacman used a random algorithm to determine the direction of its next step. Because of this, while for any given iteration number, the values of the dynamical system across hundreds of runs are different, the patterns of emotion driven by the affect states (shown for one run in Figure 2) is the same for all runs.

IV. RESULTS

The single bands shown for the emotion patterns driven by SEEKING, LUST and PLAY are expected; Pacman is generally SEEKING and has no companions to PLAY with or LUST after! Likewise, the dual banding of emotion patterns driven by FEAR is the result of high FEAR when the ghosts were nearby and a switch to low FEAR when the ghosts were beyond a certain "threshold distance." The pattern of emotion for CARE varied inversely with that of FEAR and so showed a similar dual banding. The triple banding of emotion driven by PANIC and RAGE reflects an average value with fluctuations to either side when the ghosts were closer or further than some threshold value.

The last graph in Figure 2 reflects Pacman's fluctuating emotion, driven by his likewise fluctuating emotional state as it tried to stave off death.

It became clear that given several of the physical factors with which Pacman was initialized (age, fear, health, gender etc.), that even this simple linear model offered nearly infinite possibilities, implying that mathematically, one may not need to use anything more complex than a linear dynamical systems model in order to see rich behavior.

The non-cognitive bias of the model results in a very lightweight decision engine analogous to findings described by LeDoux [25] that emotional responses are at least an order of magnitude quicker than conscious, cognitive evaluations. Additionally, by definition, dynamical system models appear to closely correspond to key concepts from philosophy of mind including: total net affect state (emotion), being some combination of various magnitudes of basic affect states, it is "ineffable," a requirement specified for qualia [10]; the principle of marginal control where a higher level (emotion) is both defined by and controls a lower level (the constitutive emotions) [31]; correspondence between the model's ability to change values while at the same time have a definitive actionable value, and William James's concept of the transitive and substantive [22].

V. DISCUSSION

The thrust of this paper has been to use a model that leverages current findings from neuroscience on how the human brain develops affect with respect to objects or other living entities in its vicinity. The premise is that if this can be replicated in a computer or embodied in a robot, then the computer or robot can develop affect for an object and thereby a sense of meaning of the object. Additionally, expected to hold true is the reverse—previously understood sense of meaning for an object in turn regenerate affect values previously associated with the object based on experience and the relative immediacy of the current interaction.

The drawback in the initial version of the equations used above, a "catch 22" as it were, is the hard coding of the distance variables. This considerably weakens the argument as possibly, without emotions, Pacman could do a better job of avoiding the ghost than with emotions. However, the point of the paper is not whether Pacman will do a better job with a "pure distance" based heuristic, but rather whether Pacman can be made to develop "affect" for the ghost, i.e., whether, at the end, Pacman has obtained a "meaning" of the ghost. Future work will aim at removing the hard coding of these distance variables and to see if through a combination of memory/persistence, using multiple instances of Pacman and the ghost, and observing ghosts consuming other instances of itself, whether Pacman can evolve similar distance relationships.

Another question is whether having values on affect dimensions for different objects in the vicinity suffice to assign true affect to that object, or does affect come into play only when there is a "body" to feel. Additionally, if needed, must a body be made of organic material such as flesh and blood or would a machine with hydraulics, sensors, and other artificial materials suffice.

Numerous researchers have raised questions related to embodiment, specifically, the symbol-grounding problem [20][36], the frame problem, the common-sense problem [21], and the rule-described/expertise problem [13]. It is the subject of much research (see for example [11][14][15][16]).

This paper takes the pragmatic approach that embodiment helps "seed" the system with meaningful attributes to real objects and that subsequently, this embodiment provides "agency" to the system; that mammals are, among other things, "organic implementations of algorithms" with senses acting as analog counters of values monitored to evaluate Panksepp well-being. The reason for the various differences in senses (counters) is merely to distinguish between different input data types. This implies that digital implementations of environment sensors would be analogous to organic implementations of the mammalian senses, or that both philosophically and practically, robots may have sufficient embodiment to support an equivalent of emotions. However, robots presently lack the dexterity needed to interact with an environment in a rich manner and consequently our approach is to simulate embodiment in a virtual environment using visualized agents.

Having few, specific and known, labels for the key variables allows for a physical feel for and an intuitive understanding of emotion, and contrasts against approaches such as neural networks. However, the question remains—how do we know it "works?" No clear answer is presently available—only the general and vague criterion that emotion seems "realistic." Popper [32] defined a hypothesis as scientific only if it is falsifiable. At present, we do not meet this criterion, i.e., the subject at hand does not presently qualify as "science."

Another (relatively minor) problem is related to the richness of the environment—it was thought that having only Pacmen, ghosts and food would not be sufficient and that additional objects would be needed in order to elicit rich emotions from the agents. However, compared to the earlier objections, this issue is of relatively secondary importance. By giving ghosts their own heuristic, by giving agents "memory" of objects they have interacted with together with emotional values attributed to these objects and the ability to store/persist them and share this knowledge with others of their "species", by allowing agents to "breed", have varying "strengths" etc., a large amount of richness can be introduced as needed. The hope is that by leveraging these additional (yet unused) options, Pacman will automatically evolve behaviors presently hard coded in the equations.

Lastly, studies by Greene [18][19] suggest a dual process theory in the human brain with one set of brain structures dealing with affect based decisions and another with utilitarian reasoning. In the case of strong competition/conflict between the two, this conflict is resolved in the ventromedial prefrontal cortex, resulting in a unified decision. There is no reason why a robot or agent cannot have a similar mechanism—a parallel and independent, analytical engine for utilitarian reasoning, with the final decisions reached by weighing recommendations based on affect against those based on utility.

VI. Conclusion

A simple, linear, lightweight, "automatic/always on" dynamical systems model, capable of adaptive goal setting, discovery, and decision. and resting on a combination of basics from dynamical systems, computer science, and neuroscience, appears to demonstrate seemingly reasonable patterns of emotion.

Four key concepts underlie the approach. First, to prevent ad hoc decision models, the approach attempts to follow closely theories of the mammalian brain. Second, the model has a strong non-cognitive bias. Third is the use of a dynamical systems model, which by definition, meets key requirements from philosophy of mind. Fourth, is the attempt to assign entities interaction dependant affect, as a way to approach the meaning/qualia problem.

Future work will: Evolve the distance relationships; Refine the model (specifically, further reduce arbitrariness and determine the parameters/coefficents) using additional findings from neuroscience; Add learning using information sharing and persistence; Explore use in a visualized agent as a lightweight decision system able to adaptively set goals and with a parallel cognitive system.

Acknowledgment

The authors thank Professors F. Martin and H. Yanco of the Computer Science Department, Professor W. Kaufmann of the Philosophy Department, and Professor J. Graham-Eagle of the Mathematics Department, all at the University of Massachusetts, Lowell, MA, for support, detailed reviews, and feedback. Thanks also to Mr. S. Cronin and Mr. R. Cole for valuable comments.

References

[1] Barrett, L. 2006. Are emotions natural kinds? Perspectives on Psychological Science 1, 28-58.

[2] Barrett, L., Mesquita, B., Ochsner, K. N., and Gross, J. J. 2007. The Experience of Emotion. Annual Review of Psychology, Volume 58, 373-403.

[3] Camras, L. A. and Witherington, D. C. 2005. Dynamical systems approaches to emotional development. Developmental Review, 25(3-4), 328-350.

[4] Camras, L. A. and Shutter, J. 2010. Emotional Facial Expressions in Infancy. Emotion Review Vol. 2 no. 2, 120-129.

[5] Chen, C. Y., Forlizzi, J., and Jennings, P. (2006) ComSlippter: An Expressive Design to Support Awareness and Availability, Alt.CHI Paper in Extended Abstracts of Computer Human Interaction (ACM CHI 2006).

[6] Damasio, A. 1994. Descartes Error: Emotion, Reason, and the Human Brain, New York, New York: Putnam Publishing.

[7] Damasio, A. 1999, The Feeling of What Happens: Body and Emotion in the Making of Consciousness, New York, New York: Harcourt Brace and Company.

[8] Damasio, A. 2003, Looking for Spinoza: Joy, Sorrow, and the Feeling Brain, Boston, Mass: Houghton Mifflin Harcourt, Inc.

[9] DeNero, J. and Klein, D. 2010. Teaching Introductory Artificial Intelligence with Pacman. Symposium on Educational Advances in Artificial Intelligence.

[10] Dennett, D. 1988. Quinning Qualia. In Consciousness in Modern Science, A. Marcel and E. Bisiach, eds, Oxford: Oxford University Press.

[11] deVega, M., Glenberg, A., and Graesser, A. eds. 2008, Symbols,Embodiment, and Meaning, Oxford: Oxford University Press.

[12] Doulamis, N. 2006. An Adaptable Emotionally Rich Pervasive Computing System, Procedings of the Eusipco Conference.

[13] Dreyfus, H. 1992, What Computers Still Can't Do: A Critique of Artificial Reason, Cambridge, Mass: MIT Press.

[14] Franklin, S. 1997a. Autonomous Agents as Embodied AI Cybernetics and Systems. Special issue on Epistemological Issues in Embodied AI, 28:6, 499-520

[15] Franklin, S. 1997b. Artificial Minds, Cambridge, Mass: MIT Press.

[16] Glenberg, A., Havas, D., Becker, R., and Rinck, M. 2005. Grounding Language in Bodily States: The Case for Emotion. In The Grounding of Cognition: The Role of Perception and Action in Memory, Language and Thinking. Cambridge, UK: Cambridge University Press.

[17] Gottman, J. M., Murray, J. D., Swanson, C. C., Tyson, R., and Swanson, K. R. 2005. The Mathematics of Marriage: Dynamic Nonlinear Models, Cambridge, Mass: MIT Press.

[18] Greene, J.D. 2010, Neuropsychology and Ethics. American Psychological Association Conference, Boston, Mass.

[19] Greene, J. D. 2004. Why are VMPFC patients more utilitarian? A dual-process theory of moral judgment explains. Department of Psychology, Harvard University, Cambridge, Mass.

[20] Harnad, S. 1990. The Symbol Grounding Problem. Physica D, 42, 335-346.

[21] Horgan, T. and Tienson, J. 1989. Representations Without Rules. Philosophical Topics, 17 (Spring), 147-174.

[22] James, W. 1892, Psychology—The Briefer Course, Notre Dame: Indiana: University of Notre Dame Press 1985.

[23] Kalkanis, C. (last accessed June, 19, 2011) Affective Computing—Improving the performance of context-aware applications with human emotions. http://cswww.essex.ac.uk/Research/digital/CK_pres.pdf

[24] Lambrinos, D. and Scheier, Ch.. 1995. Extended Braitenberg Architectures. Technical Report, No. 95.10, AI Lab, Computer Science Dpartment, University of Zurich, Zurich, Switzerland.

[25] LeDoux, J. 2003. Synaptic Self: How Our Brains Become Who We Are, New York, New York: Viking

[26] Lewis, M. D. 2005. Bridging emotion theory and neurobiology through dynamic systems modeling. Emotional and Brain Sciences, 28(2), pp. 169-245.

[27] Lorini, E. 2008. Agents with emotions: a logical perspective, ALP Newsletter, Vol. 12, No. 2-3, August

[28] Neyem, A., Aracena, C., Collazos, C. A., and Alarcon, R. 2007. Designing Emotional Awareness Devices: What one sees is what one feels, Ingeniare, Revista chilena de ingenieria, Vol. 15 No 3, 227-235

[29] Panksepp, J. 1998, Affective Neuroscience: The Foundations of Human and Animal Emotions, New York, New York: Oxford University Press.

[30] Picard, R. 1997. Affective Computing, Cambridge, USA: The MIT Press

[31] Polyani, M. 1966, The Tacit Dimension, Chicago, Ill: University of Chicago Press, 2009.

[32] Popper, K. 1959, The Logic of Scientific Discovery, London, UK: Routledge, 2002.

[33] Russell, J. A., Bachorowski, J. A., and Fernandez-Dols, J. M. 2003. Facial and vocal expressions of emotion. Annual Review of Psychology, 54, 329-349.

[34] Russell, S. and Norvig, P. 2009. Artificial Intelligence, A Modern Approach, New Jersey: Prentice Hall.

[35] Scherer, K. R. 2009. Emotions are emergent processes: They require a dynamic computational architecture. Philosophical Transactions of the Royal Society of Biological Sciences, 364 (1535), 3459-3474.

[36] Searle, J. 1980. Minds, Brains, and Programs. Emotional and Brain Sciences, Vol. 1, 417-424.

[37] Strogatz, S. (1994), Nonlinear dynamics and chaos, New York: Perseus Publishing Company.

[38] Zhou, J., Yu, C., Riekki, J., and Karkkainen, E. 2007. AmE Framework: a Model for Emotion–aware Ambient Intelligence, The Second International Conference on Affective Computing and Intelligent Interaction (ACII2007): Lisbon, Portugal

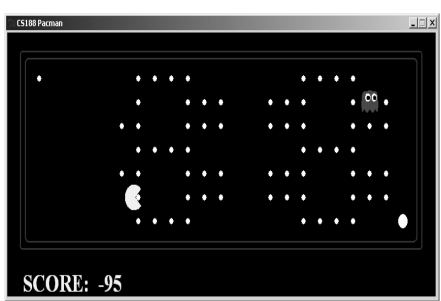

Figure 1. Snapshot from a sample run (Pacman is semi-circular, ghost is square, and food is oval)

 ISBN: 978-1-61208-158-8

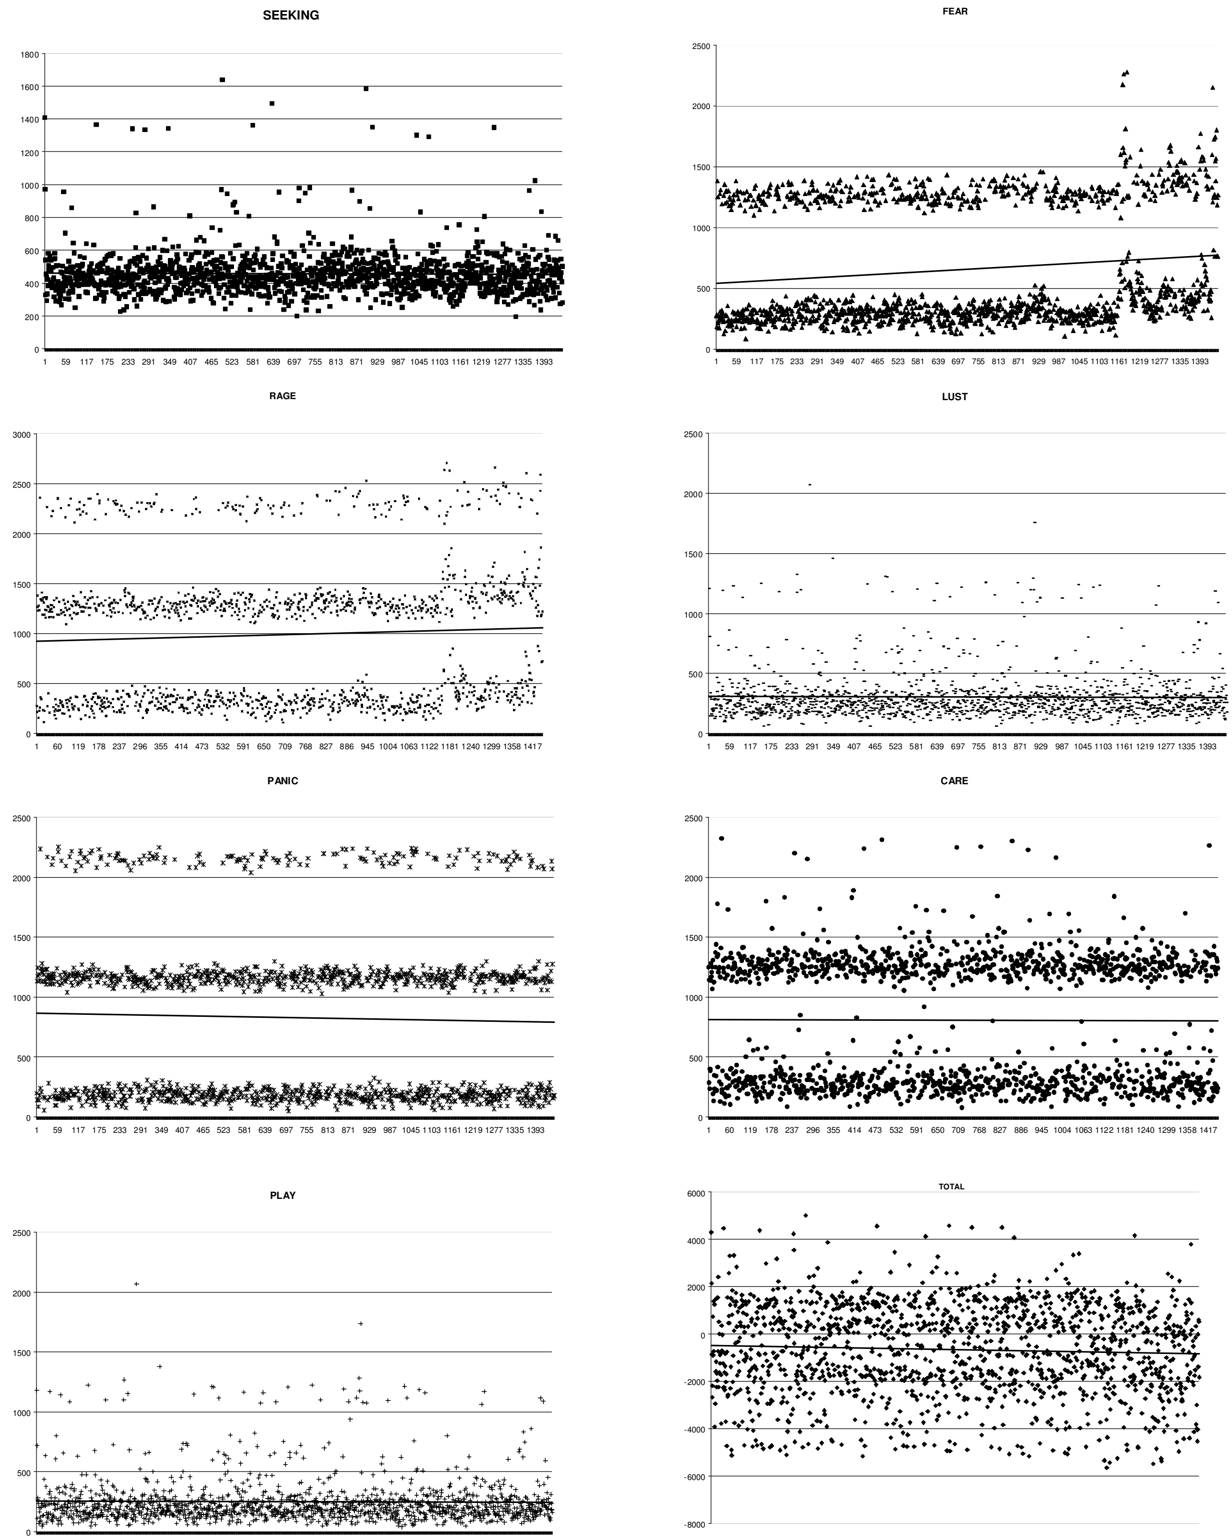

Figure 2. Affect values vs. iteration steps for the Panksepp variables modeled with a simple linear dynamical system

 ISBN: 978-1-61208-158-8