Wisely Non-Rational

Zippora Arzi–Gonczarowski
Typographics Ltd, Jerusalem 96222, Israel

1998 draft

Abstract

We show how emotions can be naturally and usefully integrated into artificial cognitive perceptions. The embedding is performed into a formal, mathematical, design that has been used before to formalize other aspects of artificial intelligent perceptually based processes. The mathematical infrastructure consists of a category of 'artificial perceptions'. Each 'perception' consists of a set of 'world elements', a set of 'connotations' that stand for embodied sensations, and a three valued predicative connection between the two sets. 'Perception morphisms' describe paths between perceptions. This categorical architecture calls for an increment in the form of a setting for emotive reactions to sensations whenever consulting with higher-level reasoning processes would be impractical. The setting for emotive reactions is conveniently provided by the 'connotations' that stand for primitive impressions of environmental elements. These sense connotations provide a natural grounding link between the sensitive and the sensible aspects of intelligence. Once emotive reactions are entered into the framework, affection pervades and increments the formal model at many levels and becomes an inseparable component of all related cognitive processes, either internal to a single perceiving artifact or inter-perceptual in a society of such artifacts.

1 Introduction

Intelligent artificial agents are typically situated in an external environment, where their intelligence is manifested in the way that they perceive that environment and interact with it. Cognitive processes take as input sensations that are provided by the lower-level sensory-motor-neural apparatus. (Connectionism has been suggested as a candidate for the mechanism that underlies this lower level apparatus [11].) Higher-level reasoning processes perform abstract logical manipulations on primitives that are the labels of the embodied sensations. It follows that perceptually cognitive processes are inherently a blend of the sensitive and the sensible.

The setting of artifacts in an external environment to interact with it exposes them to hazards. Material existences in a real physical environment as well as virtual entities in a 'cyber' environment are in jeopardy. They can be injured and incapacitated. It follows that a prudent theory of perceptive artifacts should include components to safeguard their function and existence. In a dynamic, changing, environment some of the protective measures should be typically reactive: An intelligent, perceptive artifact should be designed to sense danger and react in an appropriate manner to defend itself. A thoughtful reaction that is based on meticulous high-level reasoning that considers and carefully weighs all available aspects of the situation is, of course, desirable. However, urgency often debars a leisurely consultation with higher-level reasoning processes. In some cases there might be no urgency, yet higher-level reasoning processes could perhaps be unable to arbitrate a reaction. Hence an instinctual reaction is required. This is how emotive reactions might enter a theory of artificial perceptions as an essential component: in some situations it is wise not to be rational. Analogous reactive mechanisms are also at the evolutionary roots of human and animal emotions.

Once emotive reactions are entered into the framework, affect pervades almost every aspect of artificial intelligence. Conflicts between simultaneous emotive reactions, also conflicts with long-term goals, need to be resolved. Emotive reactions of the perceiving artifact as well as those of other emotive artifacts in its environment need to be accounted for in higher-level planning and decision making. An intelligent artifact should analyze its emotive reactions in retrospect, learn the lesson of successful as well as unsuccessful reactions, then tune up and update its emotive procedures if necessary. It follows that after a while of experience there is emotional development, so that emotive reactions become hybrids of the sensitive and the sensible just like other perceptual cognitive processes. The tangled knot between the rational aspect of intelligence and the affective aspect of intelligence is in a unity of function that seems to serve the very survival of artifacts that are situated in an environment.

We propose a mathematical categorical architecture of artificial perceptions as a framework that can readily model artificial emotions and their function. These categorical premises have already been shown to provide infrastructure for quite a few high-level cognitive processes [4, 3, 2, 1]. The mathematical primitives are generalized reactions. Sensations enable unmediated embodied interaction with the environment, as well as mental perception, reasoning, and rational judgement. Affect is often linked to perception: ($\imath$) The agent's own emotive reactions are often conjured by perception of something in its environment, as in being startled by an approaching vehicle. ($\imath\imath$) The agent often uses its perception to recognize emotions of other agents. It will be argued that even in cases where there is no connection between an emotion (or a thought) and the currently perceived environment, there is typically a link that can be traced to perceptual processes that are at the roots of things.

2 Contextualization

Ever since the Garden of Eden the human society has been documenting and analyzing its experience with emotive reactions. It is therefore hard to say anything new about human emotions that has not been said before. AI related research about emotions, including computational models, typically has the goal of understanding, illuminating, imitating, and synthesizing human affect processes. Picard [18] discusses requirements for emotionally intelligent computers, applications of affective computing, moral and social questions raised by the technology, and the design and construction of affective computers. The approach is wide ranging, yet from a human angle. [6, 5, 14, 12, 19, 21], among others, describe architectures and computational models for believable synthetic agents as well as for learning and adaptation. (A review of a few more computational proposed models appears in [17].)

The proposed mathematical categorization of artificial perceptions offers a neutral approach to cognition and affect. Emotional artificial behavior is not necessarily based on the imitation of human emotions. As just argued, emotive reactions are needed, hence natural, in perceptually cognitive artifacts that sense an environment with which they interact. It may well be that the sensation, the drive and the resulting activity are totally foreign to human experience, yet it is still an emotive reaction. It is proposed to start from generalized basic reactive mechanisms that typically vary for particular agents and environments. Building on these basic emotive reactions, categorical constructions provide a broad spectrum of generalized affective processes, with complex emotional behavior that runs parallel to higher level cognitive processes.

This study introduces a novel approach in the introduction of a neutral categorical mathematical architecture, where no such framework already exists. Most aspects of emotions that it captures are already in the literature, typically with respect to human emotions. (As Magnan and Reyes have observed [15], categorical constructs generalize, providing means to avoid over-determinations.) Sloman et al. [20] argue that there are multiple layers of emotional phenomena. Damasio [9] also makes a distinction between primary and secondary emotions. The idea that there are different levels of emotions is also supported by the categorical architecture that is proposed here. (Whether it supports parallel layers to those in the cited literature remains to be looked into.) The idea that there is a close connection between perception and emotions is emphasized by Picard [18].

The advantages of mathematical formalizations, as argued in the introduction to [8], include clarity, precision, versatility, generalizability, testability, allowance to model complex phenomena that are far too complex to be grasped by a verbal description, and allowance to use results of a well developed science. Lawvere states [13] that: 'Even within mathematical experience, only [category] theory has approximated a *particular* model of the general, sufficient as a foundation for a *general* account of all particulars' The Mathematical categorization provides a setting for the approximation of affective, cognitive, behavior by describing particular models, as well as foundations, for an account of such behavior.

3 Basic Mathematical Constructs

Definition 1 *A* Perception *is a three-tuple $\mathcal{P} = \langle \mathcal{E}, \mathcal{I}, \varrho \rangle$ where $\mathcal{E}$ and $\mathcal{I}$ are finite, disjoint sets, and ϱ is a 3-valued predicate $\varrho : \mathcal{E} \times \mathcal{I} \to \{t, f, u\}$.*

The set $\mathcal{E}$ represents the environment which is perceived. Anything which could perhaps be discerned by perception could be a legitimate element of $\mathcal{E}$, and hence a *world element* (*w-element* for short).

The set $\mathcal{I}$ stands for internal sensations that are conjured by world elements. Its elements have a subjective existence dependent on the machine. These are the building blocks of subjectively experienced, referable, sensations. Anything which may be stored and manipulated by a cognitive artifact (words, symbols, icons, etc.) could represent a legitimate element of $\mathcal{I}$, and hence a *connotation*.

The import of this study is based on the coupling of reactions with connotations. Connotational reactions are conceived to be independent of other intellectual and reasoning processes that might be impelled by higher-level reasoning processes following complex logical reasoning. Any reaction could count as such, including, but not restricted to, a simple acknowledgement of the emotive connotation that has been conjured (the human analog is a verbal expression of the emotion). Another typical emotive reaction could involve a shift in the focus of attention. We will show that this augmentation of the categorical premises provides basis for a broad spectrum of emotional activity.

The three-valued predicate ϱ is the *Perception Predicate* (*p-predicate* for short) which relates w-elements and connotations. Provided that *we now couple connotations with reactions*, then if $\varrho(w, \alpha) = t$, then w is perceived as having connotation α, and this *impels an associated reaction*, say α_t. Likewise, if $\varrho(w, \alpha) = f$, then w is perceived as lacking connotation α, and this *impels an associated reaction*, say α_f. If $\varrho(w, \alpha) = u$ then it is undefined whether perception of w should impel any of these reactions. Reactions could be empty, and in that case one gets back to the original premises of the cited works. The issue of *why* the p-predicate has any one of the three values at a certain point simply warrants no discussion: sensations are not necessarily explained, not at the (virtual) level of the proposed categorical architecture. Actual $\mathcal{E}$ and $\mathcal{I}$, *with reactions*, and the values of the p-predicate, once given, provide a definition of a particular perception, which is a mathematical object that virtually stands for an embodied affective-cognitive perceptual state. This captures the intuition that perceptions and sensations are innate to agents: their gestalt perceptions, mental imagery, neural-sensory-motor apparatus, function, internal organization *affective state*, etc. Perceptions are high-level in the sense that they reside higher than pixels on the screen or waveforms of sound. They are object centered and therefore happen at and above the level of recognition of cohesive wholes, where meaning and conscious cognizance begin to play a role.

The flexible dynamic flow between different high-level perceptions is formalized by morphisms:

Definition 2 *Let $\mathcal{P}_1 = \langle \mathcal{E}_1, \mathcal{I}_1, \varrho_1 \rangle$ and $\mathcal{P}_2 = \langle \mathcal{E}_1, \mathcal{I}_2, \varrho_2 \rangle$ be two perceptions. $h : \mathcal{P}_1 \to \mathcal{P}_2$ is a* Perception Morphism *(*p-morphism *for short) if h defines the following set mappings with a structure preserving condition: $h : \mathcal{E}_1 \to \mathcal{E}_2$, $h : \mathcal{I}_1 \to \mathcal{I}_2$, and* No–Blur *is the structure preservation condition: For all $\mathrm{w} \in \mathcal{E}_1$, and for all $\alpha \in \mathcal{I}_1$, whenever $\varrho_1(\mathrm{w}, \alpha) \neq u$ then $\varrho_2(h(\mathrm{w}), h(\alpha)) = \varrho_1(\mathrm{w}, \alpha)$.*

Composition of p-morphisms and the identity p-morphism are defined at the level of set mappings, and a theorem follows that perceptions with p-morphisms make a mathematical category, designated $\mathcal{P}rc$. This provides a well developed mathematical infrastructure for a 'theory of artificial perceptions' (in the same manner the category of groups is basis for group theory).

Various types of p-morphisms have been shown to generalize various types of transitions between perceptions (improving, communicative, analogical, and others). In the case of affective perceptions, consider a perception $\mathcal{P} = \langle \mathcal{E}, \mathcal{I}, \varrho \rangle$ with connotations $\mathcal{I} = \{accept, reject\}$, that are coupled with suitable reactions. A simple example p-morphism $h : \mathcal{P} \to \mathcal{P}'$ could be based on the identity mapping of w-elements, and the following mapping of connotations: $h(accept) = reject$, $h(reject) = accept$ with *rigid* structure preservation: $\varrho'(h(\mathrm{w}), h(\alpha)) = \varrho(\mathrm{w}, \alpha)$, namely $\mathcal{P}'$ rejects exactly whatever $\mathcal{P}$ accepts, and vice versa. In this example h is a p-isomorphism. This captures a simplistic transition, and systematizes the intuition that conversion between two extremes, both affective and cognitive, is typically easier than a total redefinition of perception and attitudes. The pre-theoretical intuition is that transitions between perceptions that are based on p-morphisms are, in a sense, 'smooth'. The mathematical properties of the p-morphism (isomorphism, one to one, many to one, onto, impossible etc.) carry meticulous information about how close these perceptions are, and pinpoint differences between them. We later provide further examples of correlations between the relative complexity of mathematical constructs and the intricacy of affective-cognitive processes.

4 Higher Level Constructs

Given a perception in an environment, simultaneous excitation of multiple emotive reactions is not unexpected. This could be due to several distinct w-elements, and/or one w-element that conjures multiple connotations. Simultaneous reactions could possibly conflict. Competition for a focus of attention is one possible example. There are higher-level reasoning and decision making methods that could sometimes be applied to resolve such conflicts. It is often the case, however, that higher level processes cannot provide a decisive arbitration. In other cases, as argued before, limited resources debar a leisurely consultation with higher-level reasoning processes. It remains to resort to solutions that are essentially affective because they are not based on obvious rational analysis.

The idiosyncratic way in which multiple emotive reactions are put together could vary. Emotive reactions could cancel one another, or tone one other down or up, or take priority one over the other, or put the entire system in a deadlock... The manner in which compound emotive reactions are handled would capture own inclinations and temperament with respect to a broad spectrum of affections. A few possible examples: Does an ambition to achieve a goal precede caution? Does competition precede a prior promise? Does institutional legality precede social obligation?

In the proposed formal model of artificial perceptions, affective-cognitive control of multiple emotive reactions can be featured in the subcategory of *Boolean perceptions*, where sets of connotations are closed under Boolean operations, namely the $\mathcal{I}$'s are Boolean algebras. [3] provides a deductive apparatus for the computation of the three-valued p-predicate for Boolean combinations of connotations. A Boolean combination of connotations is then itself a connotation, and hence now also coupled with an emotive reaction that can mediate between the multiple emotive reactions. A w-element that happens to be a computer program could be perceived, for instance, as $\varrho(w, \neg \textit{identified} \wedge \textit{it\_knows\_the\_password}) = t$. The complex connotation impels control of the multiple emotive reactions. This should consist of a resolution of the conflict: Should the perceiving artifact perhaps trust the stranger, or shut down immediately, or linger to trace an intruder...

[3] provides a categorical natural transformations that captures methodical classification and logical representation formation: Starting from basic neural-sensory-motor perceptions, organize and shape structured representations of labeled sensations that can be further used for higher-level processes. Basic neural-sensory-motor perceptions are mapped into their Boolean closures. With

added emotive reactions that are now impelled by connotations, the cognitive process is also an affective process. The contents of the complex connotation comes coupled with a compound emotive reaction. The mediated reaction could, for example, be based on: ($\imath$) Predetermined precedence between basic emotive reactions, capturing predetermined 'traits of character'. (Random precedence, for example, would capture an unpredictable character.) ($\imath\imath$) Fast heuristic decision making, capturing a form of 'gut feeling'. (Freeze all reactions to wait and 'think' is, in this context, an example of a reaction.)

Two canonical free constructions from the category of perceptions (or some subcategory) into the Boolean subcategory formalize reasonable ways to go about producing the Boolean closure. One is based on a general free construction. It provides for all possible Boolean combinations of connotations and associated reactions in a general, methodical manner. The second construction is more 'perceptually acute', and it answers to a criterion of mathematical completeness. This gives rise to even more cognitive-affective contents, because this construction provides for internalization of perceived lawlike patterns. For example, if, for all w-elements w, the p-predicate is such that $\varrho(w, \mathit{half\_full}) = \varrho(w, \mathit{half\_empty})$, then this construction merges the connotations *half_full* and *half_empty*, capturing a cognitive internalization of this pattern. The two connotations may be, indeed, associated with different, possibly conflicting, emotive reactions. The process thus gives rise not only to an emerging compound emotive reaction, but also to cognitive awareness of the conflicting emotions.

P-morphisms in the Boolean subcategory are based on Boolean homomorphisms, capturing a structural similarity between the perceptions involved.

5 Perceiving Reactions

In the sections above we focused on emotive reactions of a single perceptive artifact. An affective artifact $\mathcal{Q}$ could constitute a perceptible w-element in the environment of another affective perception $\mathcal{P}$. In that case $\mathcal{P}$ could perhaps be able to perceive emotive reactions of $\mathcal{Q}$, for instance $\varrho(\mathcal{Q}, \mathit{tired}) = f$. Fatigue could stand for lack of energy, a state that impels the agent to reach for a recharging socket – an overt action that can be perceived.

Perceiving emotions in other agents provides, among others, the necessary basis for empathy, which is, in turn, a foundation for forms of moral, ethical, and social behavior. Consider a situation where w is a w-element that stimulates an emotive reaction of repulsion in $\mathcal{Q}$ (it can be observed that $\mathcal{Q} = \langle \mathcal{F}, \mathcal{J}, \tau \rangle$ avoids proximity to w, $\tau(w, \mathit{repulsion}) = t$). Consider another emotive perception, $\mathcal{P} = \langle \mathcal{E}, \mathcal{I}, \varrho \rangle$, that perceives $\mathcal{Q}$ and its reaction to w: $\varrho(\mathcal{Q}, \mathit{avoids}\ w) = t$. An empathic reaction of $\mathcal{P}$ would be either to also avoid w, or to behave in a pro-active, pro-social, manner and 'help' $\mathcal{Q}$ by getting w out of its way. Either reaction is not necessarily desirable. $\mathcal{P}$ should distinguish its own emotions from those of others, and preserve its own interests and goals even within a pro-social, harmonious society. Empathic reactions could be regulated within the proposed formal framework of perceptions with compound emotions as explained above: If $\mathcal{P}$ also perceives that $\varrho(\mathcal{Q}, \mathit{It's\_me}) = f$, then this should invoke a competing emotive reaction to subdue the empathic reaction to $\mathcal{P}$'s own interests. The extent of the pro-social (perhaps even altruistic) behavior that should result depends on the idiosyncratic resolution of the emotional conflict, as discussed before.

In the general case $\mathcal{P}$ could perhaps perceive an entire collection of other emotive perceptions such as $\mathcal{Q}$. If its perception is able to distinguish among them and identify each one separately, then it should also be able to feature a different degree of interested empathy with each and every one of them, formalizing a broad spectrum of inter-agent relationships.

Sensible intelligence is typically marked by a discerning perception and understanding of its environment. Another pre-theoretical intuition is that observation and internalization of the peculiarities and the special patterns of a given environment are also a perceptual task. There are lawlike patterns which can be observed: one sensation may suggest another, some sensations come always together while others never do, and there are other possible connections as well. Given a perception $\mathcal{P} = \langle \mathcal{E}, \mathcal{I}, \varrho \rangle$, $\mathcal{E}$ and ϱ introduce a structural element into the internal representation $\mathcal{I}$ of connotations/sensations.

The proposed theory of artificial perceptions provides means for formalizing observation of lawlike patterns between connotations and between Boolean combinations of connotations. The example *half_full* = *half_empty* appeared above. A general lawlike pattern between connotations is defined by [3]:

Definition 3 *Let $\mathcal{P} = \langle \mathcal{E}, \mathcal{I}, \varrho \rangle$ be a perception, and let $\alpha, \beta \in \mathcal{I}$. α* **subsumes** *β, denoted $\alpha \trianglelefteq \beta$, if for all w in $\mathcal{E}$, $\varrho(\mathrm{w}, \alpha) = t \Rightarrow \varrho(\mathrm{w}, \beta) = t$ and $\varrho(\mathrm{w}, \beta) = f \Rightarrow \varrho(\mathrm{w}, \alpha) = f$.*

A perception can observe the patterns of its own connotations, enhancing a certain sense of 'self awareness'. In yet other cases, if α and β in the definitions describe perceived reactions of other artifacts, then their patterns of reaction can be observed as well. Categorically, there is no difference between the following two observations: ($\imath$) Let $\mathcal{E}$ be an environment with containers of gaz. The pattern *compressed*$\trianglelefteq$*heated* is an expected observation. ($\imath\imath$) Let $\mathcal{E}$ be an environment of infants. The pattern *hurt*$\trianglelefteq$*tears* is an expected observation. We humans distinguish the first observation as 'cognitive' and the second as 'affective' because we react with empathy only to the crying infant. An observing artificial perception would make such a distinction only if it were equipped with distinct reactions to heated gaz and to crying infants.

Observed lawlike patterns can be entered into the internal perceptual cognitive image. This is formalized by the canonical free construction (mentioned above) that incorporates observed patterns into the Boolean structure. A cognitive perception that is observing and perceptually acute can thus be 'emotionally intelligent' [10]. It can be both affectively and cognitively aware of its emotive reactions and the emotive reactions of others, and eventually use these observations in higher-level reasoning. Planning or decision making that are based on an internal cognitive structure into which the observed emotions are internalized would naturally incorporate this knowledge into the reasoning processes. Affective-cognitive perceptions would, for example, anticipate a reaction of tears whenever a child is hurt, and take precautions before, say, administration of an injection. Nurses who do this routinely become less sensitive, and do not react with such painful visceral feeling to the tears, yet there remains the cognitive apprehension of the pattern with the resulting considerate precautions, and nobody would have said that these nurses are not affective.

Perhaps the counterpart to empathic perceptions would be observing perceptions in a competitive society, that perceive 'weaknesses' of emotive artifacts and harness the observations to their own benefit. An observant perception could, for instance, record the pattern $\mathit{surprise} \trianglelefteq \mathit{distraction\_of\_focus}$, then conceive and devise an environmental cause of surprise that should, in turn, bring about the anticipated distraction of focus: make the adversary blink first. Affective perception is not necessarily 'nice'.

6 Emotional Development

Emotive connotations do not need a symbol or a collocation to impel action [7], however a symbol and a collocation provide means of reference to the sensation, the drive, and the activity. This enables later retrospective analysis and update of the conjured emotive reaction. The proposed theory of artificial perceptions suggests that perceptions should generate internal representations that are schemed and organized, individually for every perception and environment, as Boolean lattices generated over the basic sensory-motor-neural apparatus. Emotive reactions are hence labeled and collocated on the Boolean lattice. This provides a natural setting for 'emotional development'.

Emotive reactions can be revised by external catalysts and/or learning from experience. Just as one expects that an intelligent artifact should benefit from its experience when it forms an internal cognitive image of its environment [16], it is also reasonable to expect that an intelligent artifact should benefit from its experience when it adapts and models its emotional behavior. After having endured an 'emergency', or maybe even parallel to the stimulation of an emotive reaction, it should be good practice to also summon higher level reasoning processes that consider and carefully weigh all available aspects of the situation. These could be applied to evaluate the performance of the emotive reaction. This provides an opportunity to sort out the emotions involved, draw retrospective conclusions, then either reinforce a successful reaction or update and modify a less successful one. As a result the emotive reactions are improved by learning from experience and by reasoning processes, and the updates are collocated in their natural 'slot'.

Memory capabilities of artificial perceptions can benefit emotional development by keeping track, for every emotive reaction, of its history of stimulations and resulting updates. (Of course, memory tasks are typically performed more rigorously by artificial perceptions then by humans.) Retrospective evaluation processes should be able to access the emotional history through the embodied 'slots' to reason about the causes of certain reactions.

The internal lattice graph, offering sloping lines with eventual shortcuts like 'ladders and ropes', can serve as basis for mental processes. By sliding along the gratings of the Boolean trellis, internal 'thought' processes can activate the emotive reactions. Activation of emotive reactions by thought processes is, indeed, a familiar human experience. It can be supported by the proposed categorical architecture.

7 Emotions and Conceived Perceptions

Creative design processes are examples of internal 'mental' processes that could conjure emotive reactions. The proposed theory of artificial perceptions provides categorical tools for perceptual transitions from perceptions of actual environments to an internal conception of environments that are based on categorical natural transformations of actual perceptions [1]. An 'internal' perception of a conceived environment may conjure emotive reactions as well. In that case there seems to be no connection between an emotion (or a thought) and the currently perceived environment, however there is a link that can be traced to perceptually based processes that are at the roots of things. Imaginative design processes are typically formalized by conceived perceptions, that can conjure, for example, a 'positive' emotive reaction when a 'successful' design is internally conceived (anticipated success), or a 'negative' emotive reaction when a 'failed' design is internally conceived (anticipated failure). Internally conceived perceptions of other affective agents enable an internal prediction of their reactions at something, at further reactions at their reactions, etc. Such processes of creative design of conceived environments provide for a broad spectrum of internal mental activity that is rich both in cognitive and in affective contents.

8 Conclusions

We propose a mathematical categorical formalization of artificial perceptions as infrastructure for a unified virtual architecture of artificial affective-cognitive processes. It features a few levels of emotional behavior:

- Basic sensations are coupled with basic emotive reactions to environmental elements.
- Boolean perceptions support management of combinations of basic sensations and the (possibly conflicting) associated emotive reactions.
- Observation of emotive reactions of other agents, and of lawlike patterns of these emotive reactions, provides basis for affective interaction in a society of artifacts with emotions.
- Boolean internal representations that are schemed and organized, individually for every perception and environment, where emotive reactions are labeled and collocated, provide a natural setting for 'emotional development'.
- 'Internal' perceptions of a conceived environments may conjure emotive reactions as well. This is similar to activation of emotive reactions by thought processes, a familiar human experience.

References

[1] Z. Arzi-Gonczarowski. Categorical tools for perceptive design - formalizing the artificial inner eye. In *Proceedings of the 4th International Conference on Computational Models of Creative Design*, Heron Island, December 1998. Forthcoming.

[2] Z. Arzi-Gonczarowski. Percieve this as that – analogical cognitive transitions with categorical tools. In *Electronic Proceedings of the Fifth International Symposium on AI and Mathematics*, Fort Lauderdale, Florida, January 1998. http://rutcor.rutgers.edu/ amai.

[3] Z. Arzi-Gonczarowski and D. Lehmann. From environments to representations–a mathematical theory of artificial perceptions. *Artificial Intelligence*, 102(2), July 1998.

[4] Z. Arzi-Gonczarowski and D. Lehmann. Introducing the mathematical category of artificial perceptions. *Annals of Mathematics and Artificial Intelligence*, 23, Summer 1998.

[5] J. Bates. The role of emotions in believable agents. *Communications of the ACM*, 37(7):122–125, 1994.

[6] B. Blumberg. Action-selection in hamsterdam: Lessons from ethology. In *Third International Conference on the Simulation of Adaptive Behavior*, pages 108–117, Brighton, England, 1994. MIT Press.

[7] R.A. Brooks. Intelligence without representation. *Artificial Intelligence, special volume on Foundations of AI*, 47:139–159, 1991.

[8] M.A. Croon and F.J.R. Van de Vijver, editors. *Viability of Mathematical Models in the Social and Behavioral Sciences*. Swets and Zeitlinger B.V., Lisse, 1994.

[9] A.R. Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain*. Gosset/Putnam Press, New York, NY, 1994.

[10] D. Goleman. *Emotional Intelligence*. Bantam Books, 1995.

[11] S. Harnad. The symbol grounding problem. *Physica*, D 42:335–346, 1990.

[12] H. Kitano. A model for hormonal modulation of learning. In *Proceedings of IJCAI-95*, volume 1, pages 532–538, Montreal, Canada, 1995.

[13] F.W. Lawvere. Tools for the advancement of objective logic: Closed categories and toposes. In J. Macnamara and G.E. Reyes, editors, *The Logical Foundations of Cognition*, pages 43–55. Oxford University Press, 1994.

[14] P. Maes. Artificial life meets entertainment: Lifelike autonomous agents. *Communications of the ACM*, Special Issue on Novel Applications of AI, 1995.

[15] F. Magnan and G.E. Reyes. Category theory as a conceptual tool in the study of cognition. In J. Macnamara and G.E. Reyes, editors, *The Logical Foundations of Cognition*, pages 57–90. Oxford University Press, 1994.

[16] N.J. Nilsson. Eye on the prize. *AI Magazine*, 16(2):9–17, Summer 1995.

[17] R. Pfeifer. Artificial intelligence models of emotion. In V. Hamilton, G.H. Bower, and N. Frijda, editors, *Cognitive Perspectives on Emotion and Motivation*, pages 287–320. Kluwer, 1988.

[18] R.W. Picard. *Affective Computing*. The MIT Press, 1997.

[19] S. Reilly. Believable social and emotional agents. Technical Report CMU-CS-96-138, School of Computer Science, Carnegie Mellon University, 1996.

[20] A. Sloman, L.P. Beaudoin, and I.P. Wright. Computational modeling of motive-management processes. In N. Frijda, editor, *Proceedings of the Conference of the International Society for Research in Emotions*, pages 344–348, Cambridge, July 1994. ISRE Publications.

[21] J.D. Velásquez. Modeling emotions and other motivations in synthetic agents. In *Proceedings of AAAI-97*, pages 10–15, Providence, Rhode Island, 1997.