The Mathematics of Emotions

Zippora Arzi–Gonczarowski

Typographics Ltd, Jerusalem 96222, Israel

2000 draft

Abstract

An agent may have mental perceptual and cognitive habilitations, and also physical sensory motor neural skills, yet it also needs to be impelled to actually perform, in a selective manner that is relevant to its relationship with its environment. The pre-theoretical intuition is that only independent and autonmous interaction could really demonstrate its intelligence. A mathematical schema of perceptual cognitive transitions is described, and it is shown how autonomous action tendencies, namely emotions, are the natural engines that may impel actual performance of the transitions that are described by the schema. The formalism enables a rigorous distinction between quite a few types of autonomous action tendencies, yet they are all based on the same unified standard and on rather minimalistic premises. This provides an integrated view of perception, cognition, affect, and autonomous action.

1 Introduction

An AI realization of agents could be compared, in a certain sense, to some prevalent gadgets: The mechanical and electrical circuits for the intended function of the device are installed inside, yet it stays inert unless somebody plugs it in and turns it on. In simpler gadgets there are no knobs or dials: they are started by mere connection. This means that they have a single, predetermined, function. In complex devices there are quite a few switches, meaning that one could choose from a few modes of operation. Recently, the concept of *intelligent*, or *smart*, gadgets has emerged. Whenever the concept is not totally abused, it typically means that the device features a certain *autonomy*: it has sensors that respond to stimuli (that are not mediated by knobs and dials), selects one out of several modes of operation, and requires little or no decision making from the person whom it serves. When using these gadgets, one enjoys a moment of relief by 'playing lazy' (i.e. delegating one's autonomy) or by 'playing stupid' (i.e. delegating one's intelligence).

Hence, intelligence is tightly coupled with autonomy. It is neither attributed to the complexity of the circuits, nor to the variety of functions, but rather to the autonomous aspect of behavior: the ability to respond independently to changing stimuli. The essential role that interaction with an environment plays in intelligence has been recognized since Turing. [1] says: *'a prerequisite for something to be intelligent is that it has some way of sensing the environment and then selecting and performing actions.'*

The AI analogs have the hard-, firm-, or soft-,ware to perform various intelligent tasks (planning, problem solving, etc.). Questions asked: How they should be 'started'? How they should select the right course of action? If we drive them, then they are not autonomous and they do not seem very intelligent either.

This extended abstract is based on a mathematical schema of perceptual-cognitive transitions that offers a high-level blueprint[1] for the 'cognitive circuits' of an AI artifact. It will be shown how: (*i*) Affect and emotions naturally replace the knobs, dials, and switches on the 'control board' of an intelligent artifact. (*ii*) The proposed formalization enables a rigorous systematization and classification of a broad spectrum of sensible and sensitive autonomous activity.

The proposal is based on a mathematical category of *perceptions*, and *perception morphisms* formalize cognitive processes as transitions between perceptions. A categorical approach is *prescriptive* in the sense of [11]: it deals with the nature and the structure of processes and not with particular instantiations. Specific sensory-motor-neural appartuses, coupled with specific environments, should provide substitution instances of the schema. Hence, biological agents become inspirating paradigms rather than the objective of the description.

The proposed schema offers 'circuit blueprints' for various perceptual cognitive and affective processes. Interpretations and joins between various perceptions of the same environment were formalized in [7], providing 'circuit blueprints' for adaptation, communications, and learning. Generation of a representation on top of basic capabilities was formalized in [6], hence the schema has a 'circuit' to organize a grounded representation of the environment by itself. Transitions between perceptions of analogical environments were formalized in [5], providing 'circuit blueprints' for analogy making. Transitions from perceptions of actual environments to internal conceptions of designed environments were formalized in [3], providing 'circuit blueprints' for creative planning. Problems, conjectures, and results were meticulously stated using the formal premises. Quite a few justifications have been accumulated for the use of mathematical category theory in general, and the proposed premises in particular. In [2], the formalism was incremented by binding perception with emotive reactions. The import of the present work is in showing how this increment naturally generalizes to mental behavior and to autonomy control that could activate the 'circuits' in a sensible and sensitive manner.

In mathematical theories, generalizations and principles are typically described by equations. If the concepts and measurement units of several equations match, then they may be embedded in one another, forming an integrated whole. In place of equations, the proposed formalization employs commutative diagrams, that are *'the categorist's way of expressing equations'* [8, p.83]. Like equations, the diagrams can be composed into an integrated compound whole because they are all based on the same categorical ontology and premises.

The perceptual-cognitive 'circuits' will be summarized first. The limited length of this paper does not permit a totally self contained digest, but a rather skeletal synopsis, and readers are refered to other works for details. However, the essence of the formal constructs will be indicated, with emphasis on emotions and affective content that are modeled at each stage. Following the discussion

---

[1] The terminology is borrowed from [10], who suggest that categorical constructs provide *blueprints* for the design of cognitive activities.

Figure 1: A perception schema

of the integrated 'circuit' diagram, additional types of affective behavior and emotions will be proposed as the natural 'control mechanism' that activates the transitions, providing for autonomous mental behavior.

2   Basic Perceptual 'Circuits'

The pre-theoretical intuition is that high level perception is based on a classification of environmental chunks. A *Perception* is defined as a 3-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 perceived environment, *world elements* (*w-elements*) that could perhaps be discerned by a perceiving artifact. The set $\mathcal{I}$ stands for the internal labels of regularities in w-elements, *connotations* that have a subjective existence that is dependent on the perceiving artifact. The 3-valued *Perception Predicate* (*p-predicate*) ϱ relates w-elements and connotations. Actual sets $\mathcal{E}$ and $\mathcal{I}$, and the values of ϱ, once given, provide a substitution instance, capturing the intuition that perceptions and sensations are innate to artifacts, and develop relative to their environments. The $\mathcal{P}$'s stand for embodied perceptual states. They are high-level, happening at and above the level of recognition of cohesive wholes, where conscious cognizance begins to play a role. Perceptions vary across agents, situations, etc. The diagrammatic description of a perception will be based on figure 1: An oval designates a set of w-elements, a circle designates a set of connotations, and the connecting thin line represents some predicative connection ϱ between the two.

Emotions at the level of this definition were introduced in [2]. Typically, they have to do with survivability. They consist of reactions that are conjured by perception, providing for agents that can not only passively perceive, but also respond and interact with the environment. In (object oriented) programming terminology: for every connotation $\alpha \in \mathcal{I}$ and w-element $w \in \mathcal{E}$, the combination of α, w, and $\varrho(w, \alpha)$ could send a message to an *object*. *Methods that are activated* by these messages are the reactions that are associated with perception, and they are part of the definition of $\mathcal{P}$. This formalism is most typically required to capture wired reactive physical behavior of survival (like self defence, food consumption, reproduction). All agree that this is the evolutionary origin of the emotions of biological agents. A basic conjecture of this study is that once the 'circuitry' for reactions is provided, it may evolve to serve action tendencies beyond fundamental urges. This could be compared to technologies (e.g. the internet) that have been conceived for certain ends, and are then being applied for more objectives, that have not necessarily been anticipated at the outset. Researchers of biological evolution use the term *exaptations* for novelties that arise as features acquired in one context before being coopted in a different one [14]. Possible examples: (*i*) A wiring could 'exapt' by being connected to *anything that the artifact is capable of doing*. For example, in agents that feature memory or rational capabilities, reactions could be wired to update or retrieval of *data members*, delegation of tasks to higher-level rational procedures, etc. 'Think!' or 'Remember!' could be wired, just like 'Fight!'

3

Figure 2: A transition between two perceptions

or 'Flight!'[2]. (ii) Goals and purposeful behavior may be embedded by 'always' drives that are unconditionally 'wired to everything'[3]. Conflicts between wirings will be discussed shortly.

The flow between perceptions is formalized by *perception morphisms (p-morphisms)*: If $\mathcal{P}_1 = \langle \mathcal{E}_1, \mathcal{I}_1, \varrho_1 \rangle$ and $\mathcal{P}_2 = \langle \mathcal{E}_1, \mathcal{I}_2, \varrho_2 \rangle$ are perceptions, then a p-morphism $h : \mathcal{P}_1 \to \mathcal{P}_2$ defines the set mappings: $h : \mathcal{E}_1 \to \mathcal{E}_2$, $h : \mathcal{I}_1 \to \mathcal{I}_2$, and *No–Blur* is the structure preservation condition: for all w in $\mathcal{E}$, α in $\mathcal{I}$, whenever $\varrho_1(w, \alpha) \neq u$ then $\varrho_2(h(w), h(\alpha)) = \varrho_1(w, \alpha)$. The diagrammatic description of p-morphism transitions consists of thick lines between sets of w-elements and between sets of connotations as in figure 2. By [5], every such transition can be factorized into an *interpretation*, which consists of the mapping of connotations, and a *literal analogy*, which consists of the mapping of environments. They can be composed in any order. That is why they are shown as parallels in the figure. Whether the interpretation (or the literal analogy) is the first or the second factor effects the *metaphorical perception* that is generated in between. The dotted diagonals in figure 2 designate the metaphorical perceptions that blend connotations from one perception with w-elements from another.

The import of p-morphisms to affective behavior has to do with motivation and adaptivity. It is twofold: (i) Emotive reactions are part of the definition of a perception, as was just described, hence perceptual states are also affective states. A transition from $\mathcal{P}_1$ to $\mathcal{P}_2$ may involve a change in some, or all, reactions, featuring a change of mood or attitude. Example: expressions that are perceived as *sneering grin*s by $\mathcal{P}_1$, and urge a tense, defensive, reaction, may become *confiding smile*s for $\mathcal{P}_2$, urging relaxation and trust. (ii) As argued above, a reaction could be wired to anything that the artifact is capable of doing, and hence also to *the activation of a p-morphism*: perceiving something in the environment can be wired to an internal transition. It is noted that no additional 'circuitry' is needed either for (i) or for (ii): The basic reactive circuitry can be wired to mental transitions that either cause emotional change (case i), or the change is caused by emotions (case ii). Example: An agent that perceives how the environment responds to one of its behaviors may be impelled to undergo an internal transition to a modified perceptual state that features that behavior toned up (reinforcement) or down, according to the perceived response.

A combination of (i) and (ii) means that perception of something in the environment could be wired to an internal transition to a new affective state with modified, better tuned, attitudes and reactions. The intriguing property of this combination is that the described behavior is not overt. The change

---

[2] In beings with 'high-level' wired tendencies, deprivation from consuming the 'high-level' drive may cause emotional agony not unlike that which is experienced when an organism is deprived of consuming basic survival urges.

[3] Unpleasant examples are fanatic doctrines that ride piggy back on the life saving circuitry.

Figure 3: Boolean Representation Generation

of affective state may be actually observed from the outside only if a modified overt reaction is really conjured, which may happen (if at all) after a long delay, when the external catalyst that caused the transition is no longer perceptible. Trying to trace the course of change is somewhat like psychoanalysis. Returning to the analogy from the beginning, what if the smart gadget that one is using responded to some external stimulus (unobserved by the user) by changing its own internal workings, and the outcome could be noticed only long after that stimulus is gone...

Composition and the identity are defined by those of set mappings. A theorem shows that perceptions with p-morphisms make a mathematical category, designated $\mathcal{P}rc$, providing a well developed infrastructure for a mathematical theory.

3   Higher Level Perceptual Circuits

Boolean constructs were applied to further develop a theory. Figuratively, the 'plane' that is shown in figure 2 is going to serve as a 'base' for a diagram that looks like a 'box'. The construction of two 'supporting walls' will be summarized now.

Analytic organizations of grounded representations were formalized in [6] by *Boolean generations*, that close sets of connotations under Boolean operations, transforming the $\mathcal{I}$'s into Boolean algebras. (With an adequate restriction of the 3 valued p-predicate for these perceptions.) P-morphism transitions are then based on Boolean homomorphisms between connotations, capturing acute, structure aligning, interpretations. Category theoretical natural transformations systematized the transitions into perceptions that feature the Boolean property. A Boolean combination of connotations is interpretable as a logical formula, so that higher-level reasoning moduls could take this representation as input. It follows that the concreteness of the basic perceptual apparatus is married with the powers of abstraction and the rational capabilities of the higher-level apparatus. The transition is schematized in figure 3, where the Boolean set of connotations is topped with a diamond. The arrow marked *analyze* designates the natural transformation[4]. Everything here pertains to a single environment (the oval). The discussion will be extended later to more environments.

The import of the Boolean construct to affective behavior are more advanced levels of autonomy and control. Boolean combinations of perceptual constituents are themselves perceptual constituents, that can be wired to reactions, as in simple perceptions. A possibility follows to couple arbitrarily

---

[4] Two canonical Boolean closures are studied in [6].

Figure 4: Analytic representations with interpretation

complex combinations of perceptual constituents with whatever actions, overt as well as introvert, that the agent can effect. The most obvious application is an autonomous regulatory control of emotional conflicts: A complex combination of perceptual constituents may eventually be wired to a complex combination of conflicting reactions. This calls for regulatory control that should be collocated at the 'junction' of the wirings. The lattice structure of Boolean closures provides natural junction collocations for such wired control. In the biological context, the need to deal with conflicting action tendencies could have been a significant pressure behind the evolution of an entangled net: It is likely that social agents needed to regulate their natural impulses well before complex internal connections 'exapted' to representations for abstract thought.

As with simple reactions, an integrated reaction could consist of anything that the artifact is capable of doing: physical reactions (e.g. run away to avoid the conflict), mental transitions (e.g. transit to a state of mind with no conflict: 'sour grapes', 'so what!', etc.), memory functions (e.g. 'What did I do with similar conflicts before?'), message transmissions (e.g. a helpless facial expression), rational deliberation (e.g. try to deduce an infered logical solution), etc. A Boolean perception is still a perception in the sense of the basic definition: complex as well as simple connotations could be wired to immediate life saving procedures that override everything else, similar, perhaps, to the 'alarms' in [13].

Once the Boolean construct offers a mechanism to arbitrate an integrated solution to conflicting action tendencies, it can 'exapt' to arbitrate in *conflicts that have nothing to do with emotions.* Conflicting actions could be solicited also by reasoning procedures. If a rational conflict cannot be rationally resolved, then a 'gut feeling' resolution is to be arbitrated by the emotional circuitry: the infrastructure for that is already there (...not for Elliot from [9]).

Being a non-trivial possession, integrative regulatory control could be lost: (*i*) Some types of emotions (e.g. romantic attraction in humans) may resist being subdued to regulatory control. The result would be a derailing of attention and control, perhaps similar to the tertiary emotions in [13]. (*ii*) Boolean closures have exponential computational complexity. Even fast computers with enormous memories (also bigger brains and evolved nervous systems) cannot cope with arbitrarily complex Boolean combinations of perceptual constituents. Hence, being overwhelmed by too many rousing stimuli could also result in a (partial or total) derailing of attention and control.

The generating arrow of figure 3 is the basic tool for the 'erection of a wall'. If there exists a simple path (interpretation, communication) between two perceptions, as in figure 2, then this path is preserved also after the respective Boolean generations take place, and can be extended to a Boolean-structure-preserving-

Figure 5: Conceived environments with analogies

path between the generated representations. This property is mathematically warranted by the natural transformation. The underlying commutative diagram is shown in figure 4: A path from the lower left circle to the upper right diamond can be effected in either one of two possible ways: One could first generate a mental representation and then follow with a structured interpretation, or, alternatively, one could first follow a simple interpretation between basic perceptions, and then generate a structured mental representation that is based on the interpretation. This systematizes the interrelation between analytical and interpretive (communicative, learning) capabilities.

A salient property of the premises is the symmetry between $\mathcal{E}$, the environment, and $\mathcal{I}$, the representation. From a purely technical, context free, point of view, the roles that a w-element and a connotation play in the definitions are interchangeable. This *duality* has the technical consequence that any construct or theorem that is established for connotations (w-elements) can automatically be applied to w-elements (connotations), mutatis mutandis. The duality was applied to erect a second wall that faces the wall from figure 4. A formalization of creative–imaginative processes was studied in [3]. It is summarized in figure 5, which is dual to figure 4: it was technically based on mathematical results that were achieved by sweeping the roles of $\mathcal{E}$ and of $\mathcal{I}$. However, the cognitive processes that are formalized here are different. In perceptions with *conceived Boolean environments* the sets of w-elements are Boolean algebras, providing an adequate internal conception of combinations of similes and examples from the actual environment. (Boolean environments are designated here by an oval topped with a diamond.) This sets a basis for a planned perceptual–cognitive manipulation of environments, for the creative imagination and rigorous planning of designs. Transitions between perceptions of conceived Boolean environments are based on Boolean homomorphisms of w-elements. They systematize structure aligning analogies [5].

Similar to the generation of mental representations, natural transformations formalized methodical cognitive transitions from authentic environments to conceived environments. A Boolean combination of w-elments is interpretable as a logical formula that can be further applied for a rigorous effective plan to realize the conceived design, marrying creative–imaginative capabilities with higher-level rational capabilities. The natural transformation warrants that, if there exists a simple analogy path between two environments, then this path is preserved by the respective Boolean generations, and can be extended to a Boolean-structure-preserving-path between the conceived environments. This is the import of the diagram in figure 5, that interrelates between analogies and creative design. A transition from the lower left oval to the upper right oval (with diamond) can be effected in either one of two ways: One could first con-
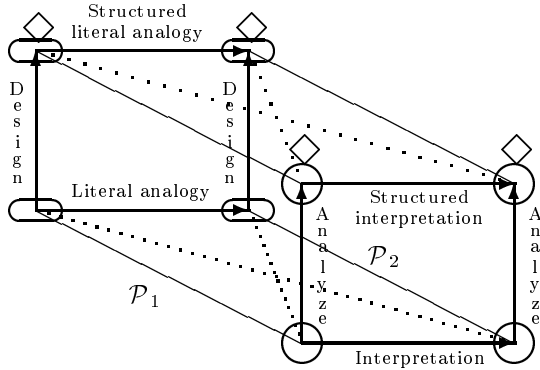
7

Figure 6: The synthesis

ceive of a design and then follow with a structure aligning analogy to another design, or, alternatively, one could first follow a simple analogy between existing environments, and then conceive of a design that is already based on the analogical environment. Everything here pertains to a single set of connotations (the circle), and the discussion will be extended later.

Emotions that are conjured by perceptions of conceived environments systematize 'what if' emotions, perhaps like the deliberative layer in [13]. An agent that perceives an ulterior environment with its inner eye may feature emotive reactions 'as if' the imagined situation was real. Example emotions of this type could be anxieties that are caused by anticipation of failure or success that have not happened yet, but are internally conceived.

A composite diagram emerges: a base with two walls define a box, a whole that features more than the some of its parts. By figure 6, a 'top cover', two 'side walls', and two 'diagonal walls' are gained, representing more perceptions and composite transitions, all of which can be interwoven in a single architecture. The category theoretical equational reasoning affirms that the composite box commutes. Various AI cognitive habilitations are interrelated in a wider theoretical framework, with a high-level prescriptive blueprint for an integrated computational framework.

Each one of the new walls describes a transition that takes a basic perception ($\mathcal{P}_1$ and $\mathcal{P}_2$, respectively) and scales it up to a cognitive perception with (i) Analytic mental representation, (ii) A perceptive inner eye that conceives of potential designs and plans, (iii) Integrative emotional capabilities with autonomous regulatory control that are enabled by the Boolean structure as described before. The top cover describes an interpretive and analogical transition that applies Boolean homomorphisms to align the high level capabilities ($i-iii$)) that were just described.

Diagonals and diagonal walls of the diagram have to do with metaphorical perceptions, that were studied in [5] (not all diagonals are shown in the figure). Emotions that are conjured by metaphorical perceptions may feature interesting discrepancies: A perceptual state that associates between an environment from one perception and connotations from another perception, could bring about emotive reactions that have developed relative to another *literal* context. They could be quite unexpected in the borrowed context. Great writers as well as cunning indoctrinators employ the subtle ruse ingeniously.

8

4 Emotions that Ignite the Circuits

Autonomous agents should somehow be motivated to actually do what the blueprint diagram describes that they *may* do. The question is what, if anything at all, happens when the 'circuit' is 'switched on': (*i*) Should the artifact be initialized to some state, and what should be the perceptual and affective constituents of this state. (*ii*) If the artifact is autonomous to decide whether at all it should perform perceptual-cognitive transitions along the arrows of the diagram, then there is also the issue of a selection between many different transitions. How should all this happen?

The only possible answers to these questions involve action tendencies that should be built into the artifact: (*i*) The artifact could be initialized to a 'genetic, inherited' perceptual state that features the necessary reactions for survival (in the spirit of 'Bootstrapping the Controller' from [11]. The pressure of survival requirements probably encouraged the evolution of the reactive circuitry in the first place. (*ii*) Transitions along the arrows, may ride on the same circuitry, and happen as responses to external stimuli.

Some autonomous action tendencies seem not to have definite reasons or purposes. They are characterized by inconnection to the artifact's relationship with its environment and feature a certain inertia. Perseverant explorer types, for example, are often motivated by such drives. There are examples from history where these tendencies seemed to override basic survival urges. A fallout of the formalism is an extension of the spectrum of emotions that it systematizes, to behavior that is driven by internal mental agendas. Internal agendas could be captured as built-in drives towards *attractor states* (although one may never really get to the attractor state)[5]. A formalization of such states is based on 'terminal objects' in mathematical category theory. It was shown in [7] that the *Total Universal Perception of* $\mathcal{E}$, $\mathcal{U}_{\mathcal{E}} = \langle \mathcal{E}, 2^{\mathcal{E}}, \epsilon \rangle$, has a lax terminal property, and it is unique up to isomorphism. A similar construct, with the *Universal Environment of* $\mathcal{I}$, is dual. Very loosely, an artifact with an autonomous 'curious and interpretive inclination' might have a built-in tendency to move along the arrows of the front wall of the box, invariably analyzing and improving its internal representation. Dually, an artifact with an autonomous 'imaginative designer inclination' might have a built-in tendency to move along the arrows of the back wall of the box, inexorably conceiving and synthesizing novel environments. Subtypes can be formalized by a subtle classifications of arrow routes that are selected, because these routes are not unique.

5 Summary and Future Work

If perception is based on chunking reality and classifying the chunks, then perceptive researchers of behavior naturally practice this procedure with respect to the domain of emotions. Classifications of emotions have been offered, among others, by [9] who distinguishes between primary and secondary emotions, and by [13] who distinguishes between (at least) three levels of emotions. This work proposed a few classifications that are enabled by a mathematical schema (some coincide, perhaps, with classifications offered before): (*i*) Basic survival reactions vs. other behaviors that ride on similar 'circuitry'. (*ii*) Overt behavior

---

[5] A similar idea is offered by the dynamical systems stance in cognitive science [12].

vs. introvert mental transitions to modified perceptual states. ($ııı$) Integration of complex combinations of (emotive or rational) action tendencies vs. loss or lack of integrative control. ($ıııı$) Reactions to conceived situations vs. real situations. ($ııııı$) Reactions to metaphorical vs. literal perceptual constituents. ($ıııııı$) Reactions that need to initialize an agent vs. later, contingent, emotional development. ($ıııııı$) Action tendencies that are responses to external catalysts vs. internal agendas.

Quite a few of the classifications ($ı - ıııııı$) are orthogonal to one another, hence their combinations enable a systematization of a computational catalogue of emotions, registering compound classifications that are grounded in mathematical rigor. Common to all these types of emotions are the minimalistic categorical premises. Like a reduced instruction set for a computer, they conflate the types of building blocks that are required for an architecture, but not necessarily the spectrum of autonomous behavior that is modeled.

As a prescriptive formalism, the schema does not specify which connotations, reactions, and environmental chunking are most suitable for a given application, or how to go about finding these constituents. This is a schema for a high level architecture that could hopefully support a broad spectrum of applications, that should come with their own specifications. A hand written example is studied in [4], but the formalism still waits to be practically implemented in a programmed system.

References

[1] J.F. Allen. Ai growing up – the changes and opportunities. *AI Magazine*, 19(4):13–23, Winter 1998.

[2] Z. Arzi-Gonczarowski. Wisely non rational – a categorical view of emotional cognitive artificial perceptions. In D. Cañamero, editor, *Papers from the 1998 AAAI Fall Symposium: Emotional and Intelligent*, pages 7–12, Orlando, Florida, October 1998.

[3] Z. Arzi-Gonczarowski. Categorical tools for perceptive design: Formalizing the artificial inner eye. In J.S. Gero and M.L. Maher, editors, *Computational Models of Creative Design IV*, pages 321–354. Key Centre of Design Computing and Cognition, University of Sydney, Australia, 1999.

[4] Z. Arzi-Gonczarowski. A blueprint for a mind by a categorical commutative diagram. In *Proceedings of the AISB2000 Symposium - How to Design a Functioning Mind*, Univ. of Birmingham, UK, April 2000.

[5] Z. Arzi-Gonczarowski. Perceive this as that - analogies, artificial perception, and category theory. *Annals of Mathematics and Artificial Intelligence*, 26, 2000.

[6] Z. Arzi-Gonczarowski and D. Lehmann. From environments to representations–a mathematical theory of artificial perceptions. *Artificial Intelligence*, 102(2):187–247, July 1998.

[7] Z. Arzi-Gonczarowski and D. Lehmann. Introducing the mathematical category of artificial perceptions. *Annals of Mathematics and Artificial Intelligence*, 23(3,4):267–298, November 1998.

[8] M. Barr and C. Wells. *Category Theory for Computing Science*. Prentice Hall, 1995.

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

[10] 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.

[11] P. Petta, C. Pinto-Perreira, and R. Ventura. Autonomy control software: Lessons from the emotional. In *Proceedings of AA'99 Workshop on Autonomy Control Software*, 1999.

[12] R. Port and T. VanGelder. *Mind as Motion: Dynamical Perspectives on Behavior and Cognition*. MIT Press, Cambridge, Mass., 1995.

[13] A. Sloman. Architectural requirements for human-like agents, both natural and artificial (what sorts of machines can love?). In K. Dautenhahn, editor, *Human Cognition and Social Agent Technology*. John Benjamins Publishing, 2000.

[14] I. Tattersall. Once we were not alone. *Scientific American*, pages 39 − 44, January 2000.