

Meaningful Information, Sensor Evolution, and the Temporal Horizon of Embodied Organisms

Christopher L. Nehaniv^{1,2}, Daniel Polani^{1,2}, Kerstin Dautenhahn¹,
René te Boekhorst¹, and Lola Cañamero¹

¹Adaptive Systems Research Group and ²Algorithms Research Group
Faculty of Engineering & Information Sciences, University of Hertfordshire, Hatfield Herts AL10 9AB, U.K.

Abstract

We survey and outline how an *agent-centered, information-theoretic approach to meaningful information* extending classical Shannon information theory by means of utility measures relevant for the goals of particular agents can be applied to sensor evolution for real and constructed organisms. Furthermore, we discuss the relationship of this approach to the programme of freeing artificial life and robotic systems from reactivity, by describing useful types of *information with broader temporal horizon*, for signaling, communication, affective grounding, two-process learning, individual learning, imitation and social learning, and episodic experiential information (memories, narrative, and culturally transmitted information).

Meaningful Information & Sensor Evolution

This paper is a review and programmatic statement on the notion of meaningful information for organisms as temporally grounded entities. Our focus is information (in the sense of Shannon) as it is or might be employed by embodied organisms. We draw a circle from the level of reactively processed sensory information between organism and environment up through many different time-scales. The fact that an organism is able to perceive or be influenced by a given signal reveals the potential for that signal to be important to the agent. The fact that a biological agent has sensors that can detect any of a particular class of signal, such as eyes that respond to light in a certain range of wavelengths, is evidence that it is useful to do so. Similarly, effector channels, carrying information for acting on the world and for producing signals, have evolved. Sensoric and actuating channels are means for an organism to interact with and manipulate its world. Information can only be said to be *meaningful* with respect to a particular embodied organism; it then traverses such channels and is statistically helpful to the organism (Nehaniv & Dautenhahn 1998a; Nehaniv 1999a). It can be studied as such with tools of mathematical information theory (Nehaniv 1999a; Nehaniv, Dautenhahn, & Loomes 1999; Polani, Martinetz, & Kim 2001a; Tishby, Pereira, & Bialek 1999).

Organisms exist, live, act and reproduce in the field of time. Details of the fine temporal structure of an organism's behavior and enaction in the world, via its body, its sensing, and its acting, are crucial for an embodied artificial life perspective (Varela, Thompson, & Rosch 1991). The usefulness of such information to an organism may arise at various levels of increasing temporal scope: reactive, affective and/or learning, and episodic experiential information (Nehaniv, Dautenhahn, & Loomes 1999). See Table 1. As the *temporal horizon* of the relevant information that organisms use increases, this meaningful information may more profoundly affect the internal milieu, learning, the development and architecture of the organism. Acquisition of such information, increasingly temporally removed from the immediate present ('the *now*'),

also requires the use of sensors and impacts on communication channels between organisms, and, in some organisms, possibly on social learning and cultural traditions. The maintenance and application of this information generally requires a corresponding internal elaboration of mechanisms, 'cognitive capacity' for learning, remembering, and representation of meaningful episodic information (whether symbolically, in internal dynamics, or encoded externally). One of the most impressive incarnations of the power of evolution is the adaptation of sensors. Information potentially useful to a biological (or artificial) organism in attaining its goals (e.g. homeostasis, survival, reproductive success, or maximizing a utility function) may play a crucial role in the *evolution of sensors* including the *channels* on which they focus; this holds as well for actuators by which the organism acts in the world, as well as for the internal architecture of the organism and how it exploits information (Nehaniv 1999a; Nehaniv, Dautenhahn, & Loomes 1999; Dautenhahn, Polani, & Uthmann 2001; Polani, Martinetz, & Kim 2001a; Liese, Polani, & Uthmann 2001). From chemical, olfactory, tactile, auditory to active (sonar), photoreceptive and general electrical and magnetic sensors, many sensoric channels are exploited to provide useful information to living agents.¹ Central to all these observations is the universality of exaptation in evolution in the sense that anatomical features and sensoric units are continually reused and reassigned for different sensoric tasks.² Equally important is also the observation that sensing is intimately connected with the appropriate processing of sensed information. Sensors are temporally immediate means providing access to meaningful information, but there exist other types of meaningful information of broader temporal scope (Table 1). The existence of channels of meaningful information is not given *a priori* as in the case of information theory. On the contrary, a formal approach to meaningful information allows one to address how such channels come into existence and evolve. By iden-

¹The development of photoreceptors ("eyes") shows a variety of approaches to solving different perceptual problems with 40-60 independent lines of descent and a variety of developmental levels (v. Salvini-Plawen & Mayr 1977; Liese, Polani, & Uthmann 2001). For example, the chameleon eye determines distances by active focusing (Ott & Schaeffel 1995), while pseudo-optical infrared sensors of pit vipers evolved from skin innervation common to most vertebrates (Noble & Schmidt 1937). There is a significant selection gradient towards complex eye structures (Nilsson & Pelger 1994).

²Some amphibians use their lungs to hear (Hetherington & Lindquist 1999). Among flies (which usually cannot hear) the parasitic fly *Ormia ochracea* uses sensory organs which originally served to determine the head position as auditory sensors by which it determines the direction of the chirping of its host (Lakes-Harlan & Heller 1992).

	DISTANT PAST	OWN PAST	RECENT PAST	IMMEDIATE PAST	NOW	IMMINENT FUTURE	NEAR FUTURE	OWN FUTURE	FAR FUTURE
SELF		AUTOBIOGRAPHICAL REMEMBERING	MOODS & SHORT-TERM MEMORY	EMOTIONS & DRIVES	HOMEOSTASIS & AUTOPOIESIS	EMOTIONS & DRIVES	INDIVIDUAL LEARNING	INDIVIDUAL PLANS	
OTHERS	CULTURE & STORIES	BIOGRAPHICAL RECONSTRUCTION	EXPRESSION READING	EMOTIONS & DRIVES	COMMUNICATION & SIGNALING	EMOTIONS & DRIVES	INTENTION READING, IMITATION & SOCIAL LEARNING	GROUP PLANS	IDEOLOGIES
OBJECTS / GOALS		LONG-TERM MEMORY	TRACKING	EMOTIONS & DRIVES	SENSORS & EFFECTORS	EMOTIONS & DRIVES	PREDICTIVE ANTICIPATION	REASON	

Table 1. Possible Temporal Horizons for Information Useful to an Organism vs. Focus of Information.

Illustration of forms of information possibly useful for particular embodied organisms at various levels of temporal scope, directed towards self, others, or objects/goals.

Sensory and effector information in organismal pathways generally has small temporal scope, close to the 'now'.

Affective information (emotions and drives) modulates behavioral control within a limited temporal scope.

Simple forms of *learning*, such as reinforcement learning, and immediate imitation appear to have a slightly broader temporal scope.

(However, some affective phenomena, related to learning and episodic information, may have extremely broad temporal scope.)

Episodic information may have still larger temporal scope, focusing on experiences of the self (memories, remembering, autobiographic information) or the experience of others (imitation, social learning, episodic memory, stories, narrative, culturally transmitted information); *plans* represent future-directed episodic information.

DISTANT PAST	OWN PAST	RECENT PAST	IMMEDIATE PAST	NOW	IMMINENT FUTURE	NEAR FUTURE	OWN FUTURE	FAR FUTURE
POST-REACTIVE	POST-REACTIVE	AFFECTIVE / LEARNING	AFFECTIVE	REACTIVE	AFFECTIVE	AFFECTIVE / LEARNING	POST-REACTIVE	POST-REACTIVE

Table 2. Possible Temporal Horizons for Organisms Employing Useful Information of Differing Scope.

Reactive Organisms: information in sensory and effector pathways generally has small temporal scope, yielding immediate responses.

Affective and Learning Organisms: emotions and drives modulate behavioral control within a small temporal scope, reinforcement learning, two-process learning (see text).

Post-reactive Organisms: Experiential episodic information with larger temporal scope, focusing on experiences (memories, remembering, auto-/biographic information, cultural knowledge, plans)

tifying sources and targets of information useful to agents, one has candidate endpoints of a channel of meaningful information. Over time, especially evolutionary time, channels of useful information need not remain static; however, sensory adaptation is limited by various trade-offs.³

An Information-Theoretic Notion of Relevant Information. An organism has to capture information from its environment via its sensors. The resources, information acquisition, as well as information processing of organisms are limited, in principle. It is therefore central to introduce an information measure that only relates to information *relevant* to the agent. The formalization of such a concept of relevant information is a central necessity to be able to develop an information-theoretical approach to understand the sensory processing of embodied organisms. In contrast to (Howard 1966)’s “information value”, our notion is inherently information-theoretical and can be measured in *bits*; moreover, going further than standard practice in the field of animal communication, which also takes an information-theoretic approach, e.g. (Bradury & Vehrencamp 1998), we systematically relate *information* to *utility* for an organism. *Meaningful information* is defined here as 1) *information in interaction games between an organism and its environment or between organisms mediated with respect to their own sensors and actuators and as 2) useful for satisfying homeostatic and other drives, needs, goals or intentions* (Nehaniv 1999a).⁴ In particular, meaningful information need not be linguistically nor even symbolically mediated. It may or may not involve representations, but must arise in the dynamics realizing the agent’s functioning and interaction in its environment (cf. the notion of ‘structural coupling’ of (Maturana & Varela 1992)), supporting adaptive or self-maintaining or reproductive behaviors, goals, or possibly plans.⁵ Under evolu-

³Perhaps the most directly “Heisenbergian” one is that between temporal and spatial resolution. For instance, flies have a higher temporal, but a lower spatial resolution than humans (Kortmann, Postma, & van den Herik 2001). Such trade-offs can arise due to aspects of light sensitivity, the necessity of depth perception, but interestingly also due to energetic causes. For example, information processing requires a certain degree of energy dissipation, of which the theoretical fundamental limit is given by thermodynamics on the quantum level (Bennett & Landauer 1985); in living systems the underlying “machinery” has a basal metabolic rate which determines the minimal rate of heat dissipation. (Kortmann, Postma, & van den Herik 2001) show that the necessity to dissipate excess heat created during the capture of visual information is a significant limiting factor.

Moreover, many interaction partners may be present (e.g. both predators on the organisms and potential mates), and signals may be perceived by many interaction partners. Selection pressures thus result from trade-offs in the signals and behaviors the agent employs, on the one hand, and, on the other, the costs and benefits incurred since various other agents might detect this behavior. Trade-offs may also arise due to costs involved in development, maintenance and processing required for the use of particular sensory and effector channels against the benefits these might provide.

⁴See (Nehaniv 1999a; Nehaniv, Dautenhahn, & Loomes 1999) for a discussion of the relationship of *interaction games* to *language games* (in the sense of Wittgenstein), and notions of *semiosis* (in the sense of Peirce) as a more appropriate model for an agent-based perspective on meaningful information than naive, “objectivist” notions of semantics.

⁵Similarly, these considerations apply to software agents, which might in a sense be considered embodied with respect to their particular environments as long as mutually perturbing channels exist between the agent and its environment (this ontology-independent definition of *embodiment* is due to (Quick et al. 1999)), with *degree of embodiment* measurable according to the complexity of the dynamics occurring between the two. The systematized, dynamic behavior of the system of agents

tion, sensor and actuator channels used in recurring types of interaction games will over generations to some degree be optimized in order to better achieve survival and reproduction, cf. (Adami, Ofria, & Collier 2000).

The simplest way to introduce usefulness to an agent is via a utility measure. Measuring the information-theoretic value of mutual information between the utility function and sensory information can guide behavior selection (Polani, Martinetz, & Kim 2001a). Moreover, it can provide feedback on the efficacy of the sensory information channel and thus guide the direction of modification, adaptation, or evolution of the channel itself (Nehaniv 1999a; Polani, Martinetz, & Kim 2001a). In the information bottleneck principle of (Tishby, Pereira, & Bialek 1999), relevance related to a random variable X (which one can interpret as the state of a system) is modeled by a *relevance indicator variable* Y , also a random variable, which is jointly distributed with X . Y can be seen as a (supervised) soft labeling of the states X . In this model, the relevance is completely modelled by Y . In other words, any feature in X is made relevant or irrelevant by being able to reflect the state of Y . (Tishby, Pereira, & Bialek 1999) therefore call the mutual information $I(X; Y) := H(X) - H(X|Y)$ (with $H(X)$ the entropy of X and $H(X|Y)$ the conditional entropy of X given Y) the *relevant information* in X . The information bottleneck principle now searches for a compressed random variable \hat{X} , obtained by a probabilistic mapping from X via a conditional mapping $p(\hat{x}|x)$. The mapping is chosen so that it compresses the information in X as far as possible (i.e. it minimizes $I(X; \hat{X})$) while keeping the relevant information in \hat{X} , i.e. $I(Y; \hat{X})$ at a constant level. Thus, it attempts to derive the most compact representation of relevant information (represented by Y) from the state X . If $I(Y; \hat{X}) = I(Y; X)$, compressing X into \hat{X} does not lose any relevant information, and \hat{X} is fully informative with respect to the relevance indicator variable Y .

In the information bottleneck model, Y is given externally, e.g. by manual labeling. How can we obtain a relevance indicator variable in the setting of agent modeling? For an organism, the choice of the right action is relevant. It is therefore natural to consider the choice of actions as relevance indicator variable (Polani, Martinetz, & Kim 2001b). Thus, for an agent model, Y is chosen as a random variable over the set of actions. To construct the probabilities for this random variable, one assumes a utility function $U(x, y)$ to be associated with an action y taken in a state x . From this utility, a joint conditional distribution $p(y|x)$ modeling the selection of an action y in a state x can be directly constructed by equiprobable selection of an optimal action in a given state (Polani, Martinetz, & Kim 2001a). Given an a priori distribution $p(x)$ for the states, one obtains a joint distribution for X and Y via $p(x, y) = p(y|x)p(x)$, from which e.g. the relevant information for an agent $I(X; Y)$ can be computed.⁶ In the context of sensor evolution, X plays the role of the environment and \hat{X} plays the role of the information transmitted by the sensors. The sensoric state \hat{X} attempts to carry only *relevant* information and to discard the rest. This is – from a mathematical point of view – the principle that we believe

and environment in such a case is referred to as an *interaction game*. Generalizing the ideas of Wittgenstein we say *meaning of the signals can be and can only be defined in terms of their usage in interaction games*.

⁶Unlike in the original information bottleneck principle, in the agent modeling view, the joint distribution of X and Y is not assumed a priori, but has to be constructed from the utility function.

underlies sensor evolution and governs its direction.⁷

The Temporal Horizon of Organisms

The common feature of learning and memory is that they provide ‘extrasensory’ meaningful information by which an organism may modulate or guide its immediate or future behavior. With generally smaller temporal scope, this also occurs with moods and emotions. Remembering involves simple or complex episodic structure, and so communication of narrative (state-histories, memories, or stories) between agents can provide ‘extrasensory’ channels of meaningful information.⁸ (See also Table 1). Martin Heidegger (Heidegger 1972) saw the state of man as being as situated in the Now, being here in the imminence of the Future in relation to the impinging Past. This *temporal horizon* seems extremely broad in humans compared to other animals, as is evidenced by our emotions such as hope and regret, concern with planning for future actions and story-telling about past or imagined events. Affect may combine with learning to provide further flexibility and widening of the temporal horizon. Extrasensory data from social learning and from narrative and historical temporal grounding help an organism escape from the present in its perception-action cycle.

Reactivity is the control of action based immediately on information from stimuli present in the surrounding environment with minimal use of internal state information. This works well for very simple behaviors such as obstacle avoidance. It even appears to be of primary importance in living systems. But for more complex behaviors a wider temporal scope in order to better contextualize actions is needed. The simplest behaviors are grounded in homeostasis, the property of biological systems to maintain key aspects of the internal milieu and interaction with the external environment within narrow ranges of important parameters. Here, *drives* realize internal informational signals (e.g. hormones), to modulate behavior. *Emotions* carry information that may be useful in behavior selection and *learning*, further freeing a biological agent from the lowest levels of blind reactivity. Information arising from experiences directly of the robot, animal or other agent itself are called “first person”; the experience of another is “second person” when this related to oneself; and “third person” information is information about objects and events from a completely external standpoint (Nehaniv 1999b). Attribution of intentions and extracting useful information from the behavior of a second person proceeds via some or all of: 0) recognition of the other as a potential interaction partner (rather than a non-agentive object); 1) the recognition of another agent as similar to the first person partner; 2) relating first person meaningful information in the present (current state) to the second person, which may make its actions, displays and signals (behavior) usefully interpretable; 3) relating first person past or temporally removed experience (narrative, stories, or history) to the second, attributing to the second person a historical ground-

⁷NB: Y models an action space. This space is usually much smaller than the typical sensor state space. In principle, with the present concept of relevant information, one would expect a much smaller sensor state space. However, the present version of the concept does not yet model temporal aspects of the perception-action cycle (but cf. the automata-theoretic approach of (Nehaniv 1999a; Nehaniv & Dautenhahn 1998b)). Such aspects will add complexity into the action and longer term behavior selection mechanisms, and require the consideration of memory and information of broader temporal scope in the calculation of relevant information (see below).

⁸We use the term “story” to refer to meaningful episodic information, although this substantially generalizes the usual (human-centric) notion.

edness and biography. Some cases of (2) are called *empathic resonance*, and are situated in the *now*, while some cases of (3) involve a wider temporal horizon – extending from the past toward the now and from the imminent future toward the now; these cases are called *biographic reconstruction* (Dautenhahn 1997; Nehaniv & Dautenhahn 1998a). Agents capable of dynamically reconstructing the biographies (histories) of the self and/or others during their lifetimes are called *autobiographic agents* (Dautenhahn 1996; 1998). *Interaction* and *communication* between situated agents often carries second person information. *Social learning* and *imitation* can be used in the immediate temporal context to acquire behavioral competencies from others to be usefully applied later (see (Dautenhahn & Nehaniv 2002)). The *communication of experiential episodic information* (in whatever form) can be viewed as a generalization of “story-telling” or narrative intelligence, allowing organisms to benefit from one another’s temporally removed experiences. In the intra-agent case this leads to a notion of *remembering* the agent’s own useful experiences, and in the inter-agent case to story-telling and benefit from ‘listening to’ others’ experiences.

Affect in Embodied Organisms. The reward or cost of experiencing reinforcing stimuli provide a uniform dimension along which an embodied organism may assess the result or desirability of action. Most attempts to introduce ‘emotion-parameters’ into AI systems can be seen as an attempt to solve the well-known contextualization problem in AI and robotics, i.e. to transcend simple reactivity by allowing the settings of parameters to modulate behavior, so as to respond appropriately to the given context. (Darwin 1872) realized the importance of emotions and their expression in animals, and his lead has been followed by recent builders of artificial systems.⁹ The feeling of experience (*qualia*) is to be distinguished from operational notions of taxis/tropisms, drive, emotion: *Taxes* and *tropisms* are ‘hard-wired’ approach or avoidance behaviors in response to stimuli, e.g. turning toward light (in plants), or moving up a gradient of food (*E. coli*) or pheromone concentration. The behaviors are stereotyped and not instrumentally arbitrary, i.e. the agent does not employ and cannot even be trained to employ alternative strategies of behavior in response to the stimulus but reacts in a fixed manner. *Drives* are homeostatic and instinctual mechanisms of internal motivational change or modulators of behavior in response to internal aspects of state: hunger, thirst, sex drive, maintaining temperature and other variables within acceptable ranges while interacting to the environment. *Reinforcing stimuli* are those which an organism will work to obtain or continue (positive reinforcing stimuli), or will work to avoid or terminate (negative reinforcing stimuli). In experimental psychology, this constitutes an operational definition for the identification of stimuli as positively or negatively reinforcing. Some stimuli are reinforcing innately, i.e. by the design, nature, or default structuring of the agent, and are referred to as *primary reinforcers*. A neutral stimulus associated (by Hebbian learning) with a reinforcer becomes a *conditioned reinforcer* if the agent will either work to obtain or avoid it.

Emotion and Learning. *Emotions* for animals can be defined as changes in state in response to reinforcing stimuli (Gray 1975; Rolls 1999). Since the definition of reinforcer is operational, so is this definition of emotion. We observe that this formal definition also makes sense for artificial agents. Studies of emotion reveal that two dimensions are

⁹Emotion systems involved in feedback control of situated agents may serve to provide the grounding for embodied agents in the body/environment coupling (Cañamero 2001; Cañamero & Petta 2001).

extremely relevant in what is understood to be required for emotion: first, emotions are *valenced*, they are either good or bad, pleasant or unpleasant; and second, they have *degree* (level of intensity). Since emotions are changes in state elicited by reinforcing stimuli, their valence and degree can serve as measures of the [un]desirability of pursuing a course of action that leads to further stimuli. In particular, the course of action to take for obtaining or avoiding a reinforcing stimulus is not encoded in either the valence or degree of the emotion, yet the agent can take this valence and degree as a guide to suggest a course of action: to work (*by unspecified actions*) either to obtain or to avoid a stimulus. How the agent works to obtain a stimulus can be to choose to invoke more general strategies and behaviors generically applicable to large classes of situations: e.g. approach, grab, flee, hide. In this way, stimulus and response are de-coupled and the relations for behavior in response to a stimulus are modifiable, dynamically learnable and reconfigurable. The *common currency* of emotion can serve to modulate the control of the agent and to motivate or suppress certain responses in its interaction with the world. This provides a mechanism for a *two-process theory of learning* (Mowrer 1960; Gray 1975; Rolls 1999). A full implementation of two-process learning in constructed agents remains to be achieved. Emotions provide the utility information regarding different reinforcing stimuli and courses of action related to them, i.e. relevance information which might not be immediately accessible to sensors. Affective information is an analogue to sensory information but is pushed further “inside the body” and can have broader temporal scope for directing action and for learning than direct sensory experience.

Experiential Episodic Information. The problem of communicating episodic information can be approached by the passing of expressions (e.g., historical ‘expanded’ sequences of salient events). Such expressions (‘histories’) may refer to the agent’s own experiences or the experiences of another agent. Passing such an expression corresponds to revealing part of one’s autobiography or, more generally, to telling a story. Algebraic, but not yet information-theoretic, methods have been applied to communicating experiential episodic information (Nehaniv & Dautenhahn 1998b; 1998a). Such ‘story-telling’ is a basic element of remembering and re-construction of experiences of one agent, and is a central element in social communication of experience. Giving artificial social agents such a capacity could result in historically embedded artifacts and could help release them from the lowest levels of reactivity.

Worlds of experience have temporal horizons limited in various ways: 1) *reactive systems* – nearly completely limited to the *now*, respond quickly, with only minimal impact of internal state; 2) *affective systems* – systems whose drives, motivations, and emotions (as formally defined above) modulate their attainment of the goals (such as survival) and help contextualize behavior; 3) *learning / affective systems* – systems like those of type (2) which employ learning (e.g. two-process learning) – of course, learning without affect can and has been added to systems of type (1); 4) *post-reactive temporally grounded systems acting with respect to a broad temporal horizon*, story-telling and remembering systems, autobiographic agents, systems with various higher degrees of social and narrative intelligence. See Table 2. Post-reactive Artificial Life can be approached using encoded experiential episodic information to capture the temporal nature of being. Histories, memories, and shared stories must be grounded in the interaction games and channels of meaningful information for the particular agent. In a constructive biology of post-reactive systems, the sky over a broad temporal horizon

opens for behavior arbitration to take better advantage of prior experience, anticipation, and future goals of the self and others, as well as demonstrated actions, or salient experience remembered, observed, or communicated.

References

- Adami, C.; Ofria, C.; and Collier, T. C. 2000. Evolution of biological complexity. *Proc. Natl. Acad. Sci. USA* 97:4463–4468.
- Bennett, C. H., and Landauer, R. 1985. The fundamental limits of computation. *Scientific American* 48–56.
- Bradury, J. L., and Vehrencamp, S. 1998. *Principles of Animal Communication*. Sinauer.
- Cañamero, L., and Petta, P., eds. 2001. *Cybernetics & Systems (special issue on Grounding Emotions in Adaptive Systems)*, volume 32(5-6).
- Cañamero, L. 2001. Building emotional artifacts in social worlds: Challenges and perspectives. In *Emotional and Intelligence II: The Tangled Knot of Social Cognition*, 22–30. AAAI.
- Darwin, C. 1872. *The Expression of Emotion in Man and Animals*. London: Albemarle.
- Dautenhahn, K., and Nehaniv, C. L., eds. 2002. *Imitation in Animals and Artifacts*. MIT Press / Bradford Books.
- Dautenhahn, K.; Polani, D.; and Uthmann, T., eds. 2001. *Artificial Life (special issue on Sensor Evolution)*, volume 7(2). MIT Press.
- Dautenhahn, K. 1996. Embodiment in animals and artifacts. In *AAAI’96 Symp. on Embodied Action and Cognition*. AAAI.
- Dautenhahn, K. 1997. I could be you - the phenomenological dimension of social understanding. *Cybernetics and Systems Journal, Special Issue on Epistemological Aspects of Embodied AI* 28(5):417–453.
- Dautenhahn, K. 1998. The art of designing socially intelligent agents - science, fiction, and the human in the loop. *Applied Artificial Intelligence, Special Issue on Socially Intelligent Agents* 12(7-8).
- Gray, J. A. 1975. *Elements of a Two-Process Theory of Learning*. Academic Press.
- Heidegger, M. 1972. *On Time and Being*. Harper Torchlight Books.
- Hetherington, T. E., and Lindquist, E. D. 1999. Lung-based hearing in an earless anuran amphibian. *Journal of Comparative Physiology* 184:395–401.
- Howard, R. A. 1966. Information value theory. *IEEE Transactions on Systems Science and Cybernetics* SSC-2:22–26.
- Kortmann, R.; Postma, E.; and van den Herik, J. 2001. Evolution of visual resolution constrained by a trade-off. (*Dautenhahn, Polani, & Uthmann 2001*).
- Lakes-Harlan, R., and Heller, K.-G. 1992. Ultrasound-sensitive ears in a parasitoid fly. *Naturwissenschaften* 79:224–226.
- Liese, A.; Polani, D.; and Uthmann, T. 2001. Study of the simulated evolution of the spectral sensitivity of visual agent receptors. *Artificial Life (Special Issue on Sensor Evolution)* 7(2):99–124.
- Maturana, H. R., and Varela, F. J. 1992. *The Tree of Knowledge: the Biological Roots of Human Understanding, revised edition*. Shambala Publications, Inc.
- Mowrer, O. H. 1960. *Learning Theory and Behavior*. Wiley.
- Nehaniv, C. L., and Dautenhahn, K. 1998a. Embodiment and memories — algebras of time and history for autobiographic agents. In Trappl, R., ed., *Cybernetics and Systems ’98*, volume 2, 651–656. Austrian Society for Cybernetic Studies.
- Nehaniv, C. L., and Dautenhahn, K. 1998b. Semigroup expansions for autobiographic agents. In Imaoka, T., and Nehaniv, C., eds., *Proc. First Symp. on Algebra, Languages, and Computation (Aizu, Japan - 30 October-1 November 1997)*, 77–84.
- Nehaniv, C. L.; Dautenhahn, K.; and Loomes, M. J. 1999. Constructive biology and approaches to temporal grounding in post-reactive robotics. *Sensor Fusion and Decentralized Control in Robotics Systems II, Proceedings of SPIE* 3839:156–167.
- Nehaniv, C. L. 1999a. Meaning for observers and agents. In *Proc. IEEE International Symp. on Intelligent Control/Intelligent Systems and Semiotics*, 435–440.
- Nehaniv, C. L. 1999b. The second person - meaning and metaphor. In *Computation for Metaphors, Analogy, and Agents*, volume 1562, Lecture Notes in Artificial Intelligence. Springer Verlag. 380–388.
- Nilsson, D.-E., and Pelger, S. 1994. A pessimistic estimate of the time required for an eye to evolve. *Proc. Roy. Soc. Lond.* 256:53–58.
- Noble, G. K., and Schmidt, A. 1937. The structure and function of the facial and labial scales of snakes. *Proc. Am. Phil. Soc.* 77:263.
- Ott, M., and Schaeffel, F. 1995. A negatively powered lens in the chameleon. *Nature* 373(6516):692–694.
- Polani, D.; Martinetz, T.; and Kim, J. 2001a. An information-theoretic approach for the quantification of relevance. In Kelemen, J., and Sosik, P., eds., *Advances in Artificial Life (Proc. 6th European Conference on Artificial Life, Prague, September 10-14, 2001)*, LNCS, 704–713. Springer.
- Polani, D.; Martinetz, T.; and Kim, J. 2001b. On the quantification of relevant information. Presented at SCAI’01 (Scandinavian Conference on Artificial Intelligence), Feb. 19–21, 2001.
- Quick, T.; Dautenhahn, K.; Nehaniv, C.; and Roberts, G. 1999. On bots and bacteria: Ontology-independent embodiment. In *Proc. ECAL’99, Fifth European Conference on Artificial Life (Switzerland - September 1999)*.
- Rolls, E. T. 1999. *The Brain and Emotion*. Oxford.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 1999. The information bottleneck method. In *Proc. 37th Annual Allerton Conference on Communication, Control and Computing, Illinois*.
- v. Salvini-Plawen, L., and Mayr, E. 1977. On the evolution of photoreceptors and eyes. In Hecht, M. K.; Steere, W. C.; and Wallacs, B., eds., *Evolutionary Biology*, volume 10. New York: Plenum Press. 207–263.
- Varela, F. J.; Thompson, E.; and Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.