

The Evolution of Sensorimotor Functionality

Mark A. Bedau

Reed College
3213 SE Woodstock Blvd.
Portland OR 97202, USA
E-mail: mab@reed.edu

Abstract

One can study the evolution of sensorimotor functionality by synthesizing this process in an abstract artificial life model, specifically, a population of agents that interact with each other and with their environment in a way that allows natural selection implicitly to shape their sensorimotor couplings. The present paper defines very general measures of environmental and sensory uncertainty, and of action's direct and indirect effects on perception, and reports a series of observations of these quantities in the context of the model.

Keywords

EVOLUTION – ADAPTATION – SENSORIMOTOR FUNCTIONALITY – SENSORY MECHANISMS
PERCEPTION – ACTION

1 The Project

Agents perceive their environment and act on it. It is obvious that action is influenced by perception; what an agent does depends on its experience. But it is equally true that perception is influenced by action. An agent does not just passively receive sensory information; the agent is embedded within its local environment and its interactions with its environment affect its own sensory experience. Rather than being an input-output device, this reciprocal relationship between perception and action gives an agent the logical form of a sensorimotor “loop” [VARELA91, PARISI92, NOLFI93].

The survival and flourishing of agents, especially those that must act autonomously in an unpredictably changing environment, depends on the *functionality* of these sensorimotor couplings—i.e., the non-accidental beneficial effects they provide the agents [BEDAU92C, BEDAU92A, BEDAU93A, BEDAU91]. On one view, intelligence consists of just this sort of sensorimotor functionality; intelligent agents are those that have the capacity to flourish by means of suitably adjusting their interactions with an environment on which they depend for resources, even though the agent (and the agent's designer, if any) cannot safely make *a priori* assumption about the environment [VARELA91, PARISI92, NOLFI93, CLIFF93A, CLIFF93B, STEELS94]. Adaptive evolution and kindred processes are especially salient sources of this kind of sensorimotor functionality [BEDAU91, VARELA91, PARISI92, NOLFI93, CLIFF93A, CLIFF93B, STEELS94]—indeed, they may be its *only* source. Thus, a fundamental research goal is to understand how the process of adaptive evolution creates and shapes sensorimotor functionality.

Those effects of action on perception involved in sensorimotor functionality take at least two forms. An agent's actions can change its position and orientation in its environment, which affects which local environment it will perceive; I call this a *direct* effect of action on perception. Walking into the kitchen directly affects what you perceive. In addition, an agent's actions can

change the environment itself, and in this way affect its perception of that environment; call this an *indirect* effect of action on perception. Opening the refrigerator door indirectly affects what you perceive.

If the sensorimotor couplings are allowed to *evolve*, the direct and indirect effects of action on perception can simultaneously both increase and decrease functionality. On the one hand, the evolution of actions' effects on perception can *generate* environmental change and unpredictability that makes survival more of a challenge. But, on the other hand, one of actions' effect on perception can be to beneficially *constrain* the environment's unpredictability.

Methodologies for studying evolving sensorimotor functionality in autonomous agents range from building real robots that interact with real environments [STEELS94], to simulating maximally realistic models of robots [CLIFF93A, CLIFF93B], to simulating more abstract models [PARISI92, NOLFI93]. My approach is to simulate maximally abstract models, with the aim of discerning maximally general and fundamental principles governing this phenomena [BEDAU92B, BEDAU91, BEDAU92D, BEDAU93B, BEDAU94A, BEDAU94B, BEDAU94C, BEDAU94D]. Pursuing this project involves addressing questions like these:

- How can we define general measures of the uncertainty of the environment, the discriminatory capacity of sensory mechanisms, and the direct and indirect affects of action on perception?
- In general, what is the nature of action's effect on the environment, and what is the nature of action's direct and indirect effects on perception?
- How do these effects depend on fundamental system parameters, such as the strength of natural selection and the sensitivity of sensory mechanisms?

In what follows I present a very abstract model of the evolution of sensorimotor couplings, then define general measures of various quantities central to understanding sensorimotor functionality, and finally report a series measurements in the model.

2 The Model

The model used here is designed to be simple yet able to capture the essential features of agents with sensory-motor couplings shaped by the open-ended forces of adaptive evolution [PACKARD89, BEDAU91, BEDAU92D, BEDAU94A, BEDAU94B, BEDAU94D, BEDAU94C]. The evolutionary system consists of many agents that could be called organisms, on analogy with biological systems. The agents exist together in an environment, in this case a toroidal lattice. The lattice has a real-valued field defined on it, $\mathcal{E}^t(x)$, which may be interpreted as an energy field, or food field. The energy field is driven by constantly adding energy (from an external source), $\mathcal{E}^{t+\Delta t}(x) = \mathcal{E}^t(x) + \Delta(x - \xi^t)$, where Δ is a function over the plane and ξ^t is a position in the plane. The agents are constantly gathering energy, extracting it at their location in the field and then expending it through their behavior. Thus, the agents function as the system's energy sinks, and the whole system is dissipative.

In analogy with biological systems, the dynamics of the population as a whole is comprised of all the birth-life-death cycles of the agents. Births occurs when agents accumulate enough energy to reproduce (see below), deaths occur when agents run out of energy, and the lives of agents consist of their interactions the environment. We label each agent with the index i and denote the population level at t by N_i^t . Time is discrete. One unit of time t is marked by each agent interacting with the environment.

During its lifetime, each agent exchanges information with the environment by sensing and affecting the energy field in its local neighborhood. We assume that there is a discrete set of different possible sensory states, $s \in S = \{s_1, \dots, s_{N_S}\}$. The agents exchange no information with each other directly (although this would be an easy generalization). Each agent (labeled with the index i) has certain information associated with it: (i) a current location, x_i^t , (ii) a current sensory state, s_i^t , (iii) a current reservoir of energy, E_i^t , and (iv) a *strategy map*, φ_i , that yields a behavior given sensory information as input. In this model, an agent's behavior is a vector \vec{b} denoting the agent's movement in the environment, $x_i^t \rightarrow x_i^{t+1} = x_i^t + \vec{b}_i^t$; in general, $\vec{b}_i^t = \varphi_i(s_i^t)$ (an exception is explained below). We assume that \vec{b}_i^t is a member of a discrete set of different possible behaviors, $\vec{b}_i^t \in B = \{\vec{b}_1, \dots, \vec{b}_{N_B}\}$. (Another easy generalization would be to include other kinds of possible behaviors.) Each behavior causes a change in the agent's energy supply: an augmentation from extracting energy at the agent's new location, a constant-sized reduction, and a reduction proportional to the distance moved, $E_i^t \rightarrow E_i^{t+1} = E_i^t + \alpha(\mathcal{E}^{t+1}(x_i^{t+1})) - \beta - \gamma(\vec{b}_i^t)$. (Here, $\alpha(z)$ may be interpreted as the energy extraction function, β as the constant metabolic cost of surviving, and $\gamma(z)$ as a function for the variable metabolic cost of a specific behavior).

The strategy map φ_i operating on a sensory input s_i^t has a particularly simple form because the sensory input is discrete. Since $s_i^t \in \{s_1, \dots, s_{N_S}\}$, we may identify the function φ_i with its graph, a set of N_S behavior values, $\{\varphi_{is}\}$. Pursuing the biological analogy, I will consider the strategy elements $\{\varphi_{is}\}$ as i 's *genome*, and each particular element $\vec{b} = \varphi_{is}$ as the trait (or allele) at the s^{th} locus of i 's genome. Reproduction occurs when an agent's energy E_i^t exceeds a threshold, E_r . During reproduction, a parent's energy is split evenly among its children, and a parent's strategy elements $\{\varphi_{is}\}$ are copied over to its offspring, each element with some probability μ of being altered. Altered strategy elements are chosen randomly from the set of possible strategy elements B . We call μ the *mutation rate*.

In order to investigate how natural selection affects the evolutionary dynamics of strategies $\{\varphi_i\}$, I introduce a *behavioral noise* parameter, P_0 , defined as the probability that \vec{b}_i^t is chosen at random from B rather than determined by $\varphi_i(s_i^t)$. If $P_0 = 1$, then agents survive and reproduce differentially, and children inherit their parents' strategy elements (except for mutations), but the inherited strategies $\{\varphi_{is}\}$ reflect only random genetic drift.

A sensory map, $S_\theta: N_x \rightarrow S$, associates a sensory state with each local environment. The sensory map generally produces an agent's sensory state, $s_i^t = S_\theta(x_i^t)$ (but this rule can be violated by sensory noise, in which case s_i^t is chosen randomly from S .) The precise nature of the mapping depends on the sensory apparatus, θ , which passes some of the information available at local environments. All agents have the same kind of sensory apparatus, unless it is put under genetic control and allowed to evolve.

An especially simple kind of sensory apparatus is a set of sensory thresholds, $\theta = \{\theta_1, \dots, \theta_t\}$, which partition the possible environmental states at each site in a local environment into discrete "bins" and associate each bin with a unique signal. In this case, the sensory state associated with the sites in a local environment is a set of signals from those sites. The thresholds determine at which energy levels a site's signal changes. In this way, the sensory thresholds determine what raw sensory information the agents receive from their environment.

In the simplest case, one threshold θ assigns the energy level at each site x in a local environment into one of two "bins," to which signals are then assigned, $\text{sig}(x^t) = T(\mathcal{E}^t(x))$, where $T(z)$ is a function which is zero for $z < \theta$ and one otherwise. θ may take on values within the interval defined by the minimum possible value of $\mathcal{E}(x)$ and one greater than the maximum possible value of $\mathcal{E}(x)$. Note that if θ is set to either its maximum or minimum value, then all sites at all times send the same signal; the agents are "blind."

This model provides an abstract and general way to study the evolution of sensorimotor functionality. In this setting, the sensorimotor coupling goes from the sensory states the agent experiences in response to its environment, to the action triggered in response to that sensory state by the agent’s corresponding strategy element, then to the environmental changes produced by the agent’s actions, at which point the cycle repeats. The evolving link in this coupling is the strategy elements. Sensorimotor functionality is crucial for these agents, since their survival and reproduction depends directly on their ability to sustain themselves by finding continually locating energy in their environment. At the same time, the collective effect of the agents’ actions continually causes unpredictable effects on the environment’s energy field. The agents possess no *a priori* model of the environment, but the adaptive evolution of their strategy elements builds an *a posteriori* model of it—a collection of strategy elements well suited for flourishing in that environment. This *a posteriori* model is continually adapting in an open-ended manner to new environmental contingencies. The agents are autonomous in the sense that their internal action-generating models are created through an evolving history of contingent interactions with their environment.

3 The Measures

As a necessary prelude to a general study of the evolution of sensorimotor functionality, this section presents objective, general measures of the uncertainty of the environment (two forms), of the discriminatory capacity of the agents’s strategy map (two versions), and of the direct and indirect effects of action on perception. All these measures are based on the Shannon entropy [SHANNON49], which quantifies the information-theoretic uncertainty of a probability distribution. But first I describe how I quantify sensorimotor functionality.

3.1 Sensorimotor Functionality

Since our model is energy driven, and the survival and reproduction of the agents depends directly on their success at extracting energy from the environment, we can judge the overall performance of the population by its efficiency at extracting and using energy from its environment. A crude measure of this efficiency is reflected by the *residual energy* in the environment \mathcal{E}_{resid}^t , i.e., the energy that is not contained in the organisms:

$$\mathcal{E}_{resid}^t = \sum_x \mathcal{E}^t(x). \tag{1}$$

(Recall that $\mathcal{E}^t(x)$ is the residual energy at site x .) Population performance is inversely related to residual energy. Residual energy can be directly compared across different systems if all background system parameters affecting performance are equivalent. This is the case in all observations reported here.

3.2 Uncertainty Associated with the Environment

Possible environmental states are distributed over the world sites, and this distribution changes over time. In the present model, the the world contains only an energy field (besides the organisms), so possible environmental states correspond to possible levels of energy. Different aspects of the uncertainty—or, conversely, structure—associated with the environment can be distinguished and measured on the basis of probability distributions concerning different aspects of the energy field.

Site uncertainty of the environmental states, $H_{x,\mathcal{E}}^t$, reflects how uniformly the residual energy is distributed among those single sites in the world that contain some energy:

$$H_{x,\mathcal{E}}^t = - \sum_{x=1}^{N_x} p(x) \log_2 p(x), \quad (2)$$

where N_x is the number of sites in the world, and $p(x)$ is the probability that a randomly chosen unit of residual energy is located at world site x , which is measured as the proportion of the residual energy at x :

$$p(x) = \frac{\mathcal{E}^t(x)}{\mathcal{E}_{resid}^t}. \quad (3)$$

(Recall that s_i^t is i 's sensory state and N_S is the number of possible sensory states.)

Environmental state uncertainty of the sites, $H_{\mathcal{E},x}^t$, reflects how uniformly the energy values found at sites in the world are distributed over the possible energy levels:

$$H_{\mathcal{E},x}^t = - \sum_{l=0}^{N_{\mathcal{E}}} p(l) \log_2 p(l), \quad (4)$$

where $p(l)$ is the probability that a randomly chosen world site x contains energy level l , and $N_{\mathcal{E}}$ is the maximal possible energy level at a site (in energy level units). That is, $p(l)$ is the proportion of sites that have level l :

$$p(l) = \frac{\sum_x \delta(\mathcal{E}^t(x) - l)}{N_x}, \quad (5)$$

where $\delta(x)$ is the Dirac delta function, equal to one if $x = 0$ and zero otherwise.

The values of site uncertainty and environmental state uncertainty generally behave quite differently. Note that if the environmental state uncertainty is zero, then the energy field is uniformly distributed across all sites in the world, so the site uncertainty is maximal. In addition, note that if the environmental state uncertainty is maximal, then all possible energy levels are uniformly distributed over world sites, so site uncertainty will be intermediate. Finally, note that if the environmental uncertainty is zero, then all the world's energy is located at one site, so the site uncertainty will be positive but nearly zero. Since site and environmental uncertainty measure different quantities, the pair provide a useful composite view of the uncertainty associated with the environment. (Additional measures of uncertainty associated with the environment could also be defined.)

3.3 Uncertainty Associated with the Sensory Apparatus

The agents experience sensory states produced by contact with their environment. The frequency with which the environment generates different sensory states in effect determines the sensitivity, i.e., discriminatory capacity, of the agents' sensory mechanisms. The frequency of sensory states can be measured with respect to either the subset of local environments that the agents actually encounter or all local environments in the world, each of which yields a measure of sensory sensitivity.

Sensory uncertainty of the agents, $H_{s,i}^t$, reflects how uniformly distributed are the sensory states s detected by the agents i :

$$H_{s,i}^t = - \sum_{s=1}^{N_s} p_i(s) \log_2 p_i(s), \quad (6)$$

where $p_i(s)$ is the probability that the sensory state that a randomly chosen agent i experienced is s , measured as follows:

$$p_i(s) = \frac{\sum_i \delta(s_i^t - s)}{N_i^t}. \quad (7)$$

(Recall that s_i^t is i 's sensory state and N_i^t is the population level.)

The agents' sensory uncertainty is determined by three factors: the uncertainty of the world the agents sense (discussed in section 3.2), the sensory mapping S_θ , and the possibility of sensory noise (refer to section 2 for the latter two). Aside from possible sensory noise, the agent's sensory uncertainty reflects how much raw information about the environment the agents have detected.

Sensory uncertainty of the sites, $H_{s,x}^t$, reflects how uniformly the sensory states s are distributed over the world sites x , i.e., how much uncertainty there is about which sensory state s an agent located at a random site x would experience:

$$H_{s,x}^t = - \sum_{s=1}^{N_s} p_x(s) \log_2 p_x(s), \quad (8)$$

where $p_x(s)$ is the probability that the sensory state that an agent located at site x would sense is s , measured thus:

$$p_x(s) = \frac{\sum_x \delta(S_\theta(x^t) - s)}{N_x}. \quad (9)$$

(Recall that $S_\theta(x^t)$ is the sensory state an agent located at site x at time t would experience.)

Sensory uncertainty of the sites indirectly measures the uncertainty of the environment, specifically, the amount of raw information in the environment that could be discerned by the agents given their sensory apparatus θ . This quantity is sensitive to changes in θ , of course.

3.4 The Effects of Action on Perception

In selection is absent (i.e., $P_0 = 1$), the agents actions will be randomly chosen from the available set of possible actions (see section 2). In this case, aside from random sampling effects, the agents' actions will place them in a subset of local environments that is representative of those available in the world. Thus, the agents' sensory uncertainty will be approximately the same as the sites' sensory uncertainty. This entails that the absence of selection prevents the possibility of there being any significant direct effect of action on perception, i.e., any effect on perception due to action placing agents in a distinctive subset of available local environments. Thus, an approximate measure of the magnitude of *action's direct effect on perception*, V_D^t , is the difference between the agents' and the sites' sensory uncertainties,

$$V_D^t = |H_{s,i}^t - H_{s,x}^t|. \quad (10)$$

V_D^t reflects the direct effect of action on perception to the extent that V_D^t that is not due to sampling effects. The value of V_D^t without selection is a rough measure of the magnitude of the difference due just to sampling.

Actions could have a significant effect on the environment without having any indirect effect on perception. Agents are able to detect only a fraction of the raw information that exists in the environment, even that information created through their own affects on the environments. Thus, there is no guarantee that agents can perceive their effects on the environment. However, the sites' sensory uncertainty measures precisely the environmental information that they *can* perceive. We have noted that changing the sensory mechanism θ can change the sites' sensory

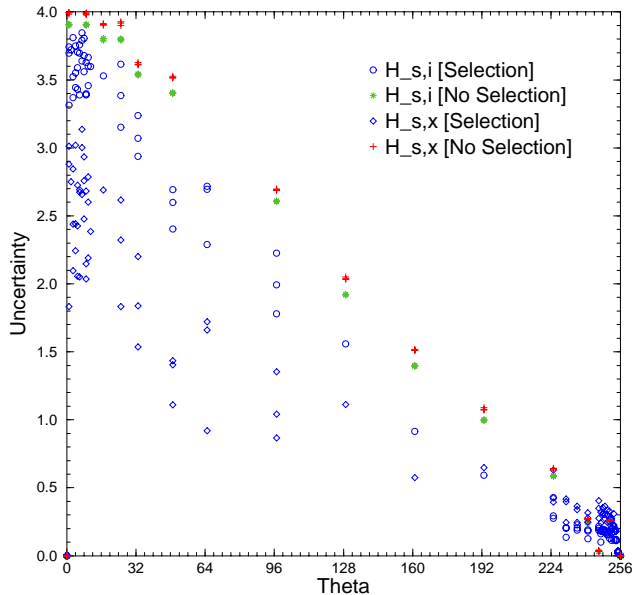


Figure 1. Time averaged sensory uncertainty of the world sites, $H_{s,x} = \langle H_{s,x}^t \rangle_t$, and of the agents, $H_{s,i} = \langle H_{s,i}^t \rangle_t$, as a function of θ , with and without selection.

uncertainty. In fact, if we can assume that changing θ does not change the no-selection agents' effect on their environment, then the no-selection sites' sensory uncertainty is a measure of how changing θ changes the sites' sensory uncertainty. In this case, a rough measure of the *indirect effect of action on perception*, V_I^t , is the difference between the sites' sensory uncertainty with and without selection:

$$V_I^t = |H_{s,x}^t(P_0 = 0) - H_{s,x}^t(P_0 = 1)|. \quad (11)$$

We are entitled to interpret V_I^t as the indirect effect of action on perception only if we can assume that changing the sensory mechanism does not change the kind of effect that no-selection agents have on their environment.

4 The Experiment

I observed the quantities defined in section 3 in an extensive series of simulations of the model. I used an especially simple form of sensory map, determined by just one sensory threshold θ . Previous studies of this model have shown that agents can achieve significantly greater sensorimotor functionality at mutation rates in the range $10^{-4} < \mu < 10^{-3}$, and that at these mutation rates population diversity displays complex dynamics [BEDAU92D, BEDAU93B, BEDAU94A, BEDAU94C]. Endogenous determination of an individual's mutation rate also evolves towards these same mutation rates [BEDAU94B, BEDAU94C]. Thus, in the simulations reported here the mutation rate was set at $\mu = 10^{-3}$. Each simulation consisted of 10^5 time steps. All parameters in all simulations were the same except for the presence or absence of selection ($P_0 = 0$ or $P_0 = 1$) and the setting of the sensory threshold θ . (A full specification of the parameters used




Figure 2. Left: Time averaged residual energy, $\mathcal{E}_{resid} = \langle \mathcal{E}_{resid}^t \rangle_t$ as a function of θ , with and without selection. Right: Time averaged residual energy as a function of time averaged sensory uncertainty of agents, $H_{s,i} = \langle H_{s,i}^t \rangle_t$, with and without selection.

in the simulations is provided in Appendix A.)

The notable observations from these simulations are the following:

Maximal Sensory Discriminability. Figure 1 shows that sensory uncertainty of both agents and sites strongly depends on the sensory threshold θ . (The differences between selection and no selection, and between agents and sites, will be discussed below.) Blind agents ($\theta = 0$ or $\theta = 256$), extract no information from the environment and so always have zero sensory uncertainty. When agents can make sensory discriminations, environmental uncertainty generally monotonically decreases with θ . The lower θ is, the more sensory discriminations agents can ($H_{s,x}$) and do ($H_{s,i}$) make; very high θ make the agents almost blind. Thus, in the present class of environments maximal sensory discriminability tends to be provided by θ settings that distinguish between the presence and absence of energy. (This is believable, given that these environment consists of energy-free “deserts” sparsely scattered with pyramidal energy heaps.)

Functionally Optimal Sensory Discriminability. Figure 2, left, shows that the presence or absence of selection affect on sensorimotor couplings makes a significant difference to population performance. When sensorimotor couplings cannot adapt to their environment, performance is flat across the θ spectrum (as could be expected) and, for given θ values, worse than when sensorimotor couplings do adapt (again, as could be expected). Although the amount of sensory information extracted from the environment varies with θ (figure 1), figure 2, right, shows that this has no effect on performance in conditions of no selection.

On the other hand, when the presence of selection permits sensorimotor couplings to undergo adaptive evolution, performance depends strongly on θ , being optimized at low values, approximately in the range $1 \leq \theta \leq 50$. Comparison of figure 2, left, with figure 1 shows that, when selection is present, level of performance correlates extremely well with both how much sensory information about the environment the agents *do* extract (the

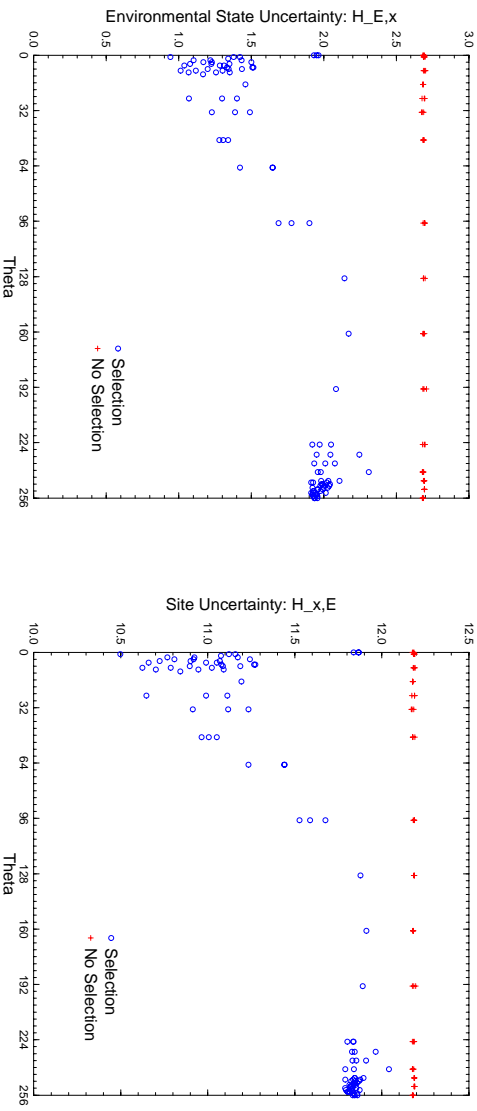


Figure 3. Left: Time averaged environmental state uncertainty of the world sites, $H_{\mathcal{E},x} = \langle H_{\mathcal{E},x}^t \rangle_t$, as a function of θ , with and without selection. Right: Time averaged site uncertainty of the environmental states, $H_{x,\mathcal{E}} = \langle H_{x,\mathcal{E}}^t \rangle_t$, as a function of θ , with and without selection.

agents’ sensory uncertainty, $H_{s,i}$) and how much they *could* extract (the sites’ sensory uncertainty, $H_{s,x}$). As an extreme case, note that blind bugs ($\theta = 0$ or $\theta = 256$) perform quite poorly. In general, populations that extract more information from their environment also extract more energy from it (figure 2, right).

There is an intriguing detail in these observations: Residual energy increases at the highest values of agents’ sensory uncertainty under conditions of selection. Thus, in these simulations, maximal agents’ sensory uncertainty does *not* correspond to maximal functionality. (For example, the simulation which produced the best performance—the lowest point in figure 2 left and right—achieves far from the highest sensory uncertainty of the agents.) Whether or not this effect is significant can be determined only with more extensive simulations.

These results clearly indicate that agents with evolving sensorimotor couplings benefit when their sensory mechanisms provide more information about their environment. Although there is some question whether maximal sensory discriminability yields optimal functionality (see the far right of figure 2 right), the general trend is clear: the more information, the better. (This could be expected in the present context. It should be noted that, in other contexts, more information is not necessarily better; for example, subjecting agents to lots of irrelevant information might make them perform worse.) The general pattern of these results clearly shows that θ , the sensory discriminability of the sensory mechanism, indirectly controls the degree to which sensorimotor couplings can adapt and become functional. If the sensory discriminability is low enough, then the sensorimotor couplings cannot adapt at all and sensorimotor functionality never improves over that of blind agents. In the present class of environments, it turns out that the greatest sensorimotor functionality tends to come from sensory mechanisms that discriminate between something approximating the difference between some energy and no energy.




Figure 4. Left: Indirect effect of action on perception, $V_I = |H_{s,x}^{Po=0} - H_{s,x}^{Po=1}|$, as a function of θ . Right: Direct effect of action on perception, $V_D = |H_{s,i} - H_{s,x}|$ as a function of θ , with and without selection.

Action Affects Environment. Figure 3 shows that action can have an effect on the environment, and that this effect depends on whether or not sensorimotor couplings are adapting to their environment. When selection is absent, the site and environmental state uncertainty are each flat across the spectrum of sensory thresholds. Thus, the details of the agents’ sensory map have no effect on the structure of the environment. Thus, in the no-selection condition, agents’ actions do not affect the structure in the environment. On the other hand, when sensorimotor couplings are able to undergo adaptive evolution successfully, then the agents’ actions do affect the environment. In particular, both environmental state uncertainty and site uncertainty decline in the presence of selection; agents cause both the possible energy levels and the proportion of energy to be less uniformly distributed. In addition, the magnitude of this effect on the environment depends on the value of θ . (It is interesting to note that both site and environmental state uncertainty seem basically flat for the upper half of the θ spectrum, roughly $128 < \theta < 256$, since every other quantity measured varies somewhat with θ in this range of values.) Since we have seen that θ indirectly controls the degree to which sensorimotor couplings can adapt, the effect of action on the environment is proportional to the degree to which sensorimotor couplings adapt.

Action Indirectly Affects Perception. Action must have an effect on the environment if it is to have an indirect effect on perception. The effect of action on the environment (discussed above) shows that action *might* indirectly affect perception, but this is uncertain since actions’ effect on the environment is not sufficient for its indirect effect on perception. Close scrutiny of figure 1 resolves this uncertainty. Notice that the sites’ sensory uncertainty is much lower when selection is present than when selection is absent. This is an indirect effect of action on perception, more clearly revealed in figure 4, left. (We are entitled to interpret V_I as a measure of the indirect effect of action on perception since we saw above that changes in sensory discriminability, θ , do not change actions’ effect on

the environment under conditions of no selection.) Notice also that this effect is generally much stronger around those θ values that make the sites' sensory uncertainty much larger, roughly $1 \leq \theta \leq 50$. Action indirectly lowers the uncertainty of the environment to the extent that the sensory discriminability allows sensorimotor functionality to be shaped by adaptive evolution.

Action Directly Affects Perception. When selection is absent, agents' and sites' sensory uncertainty is roughly the same (there is a slight difference due to sampling effects in the agents' sensory uncertainty). On the other hand, when selection is present sensory uncertainty of the sites is significantly lower than sensory uncertainty of the agents, so V_D is significantly positive (figure 4, right) over most of the θ spectrum, especially those θ values at which performance is most enhanced (roughly, $1 \leq \theta \leq 50$). Action is clearly affecting perception directly, by putting agents into a detectably different subset of the local environments available in the world. Furthermore, the strength of this effect depends on θ , being maximized when the sensory discriminability is greatest, roughly, roughly $1 \leq \theta \leq 50$. Thus, action directly lowers the uncertainty of the environment to the extent that sensorimotor functionality is subject to adaptive evolution.

5 Conclusion

I have presented very general methods for objectively measuring the direct and indirect effects of action on perception and their connection with sensorimotor functionality. Both the direct and indirect effects can be observed in the very simple model explored here. Their existence and strength depends on the discriminability of the sensory mechanism, and thus is correlated with the degree to which adaptive evolution can shape sensorimotor functionality. The robustness of these effects in this very simple model suggests that they are very general and fundamental processes.

Future work will extend the present project in various respects. First, I am studying the correlations among the uncertainty measures defined here and the other macroscopic quantities I have examined in other work [BEDAU91, BEDAU92D, BEDAU93B, BEDAU94A, BEDAU94B, BEDAU94D, BEDAU94C]. Second, I am letting θ evolve, and thus adding the possibility for a new kind of sensorimotor functionality. Like the adaptive evolution of mutation rates [BEDAU94B], this constitutes a simple form of the evolution of evolvability—the second-order adaptation of the mechanisms that govern the first-order adaptation of strategy elements. Preliminary results show that θ evolves toward those values that yield optimal functionality. Third, I am studying how to measure the amount of *useful* information that is extracted from the environment. The sensory uncertainties defined here measure the extraction only of “raw” information. Sensorimotor functionality might well be set in significantly sharper relief when reflected by the capacity of sensory mechanism to extract useful information from the environment.

Acknowledgements

Thanks to Robert Seymour for valuable help producing and analyzing these data, and for valuable comments on the document.

References

- [BEDAU91] M. A. Bedau and N. H. Packard. Measurement of evolutionary activity, teleology, and life. In C. G. Langton, C. E. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, volume X of *SFI Studies in the Sciences of Complexity*, pages 431–461. Addison-Wesley, Redwood City, CA, 1991.
- [BEDAU92A] M. A. Bedau. Goal-directed systems and the good. *The Monist*, 75:34–51, 1992.
- [BEDAU92B] M. A. Bedau. Philosophical aspects of artificial life. In F. J. Varela and P. Bourguine, editors, *Towards a Practice of Autonomous Systems*, pages 494–503. Bradford Books/MIT Press, Cambridge, MA, 1992.
- [BEDAU92C] M. A. Bedau. Where’s the good in teleology? *Philosophy and Phenomenological Research*, 52:781–806, 1992.
- [BEDAU92D] M. A. Bedau, F. Ronneburg, and M. Zwick. Dynamics of diversity in an evolving population. In R. Maenner and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 95–104. Elsevier, New York, 1992.
- [BEDAU93A] M. A. Bedau. Naturalism and teleology. In S. J. Wagner and R. Warner, editors, *Naturalism: A Critical Appraisal*, pages 23–51. University of Notre Dame Press, Notre Dame, IN, 1993.
- [BEDAU93B] M. A. Bedau and A. Bahm. Order and chaos in the evolution of diversity. In *Proceedings of the Second European Conference on Artificial Life (ECAL93)*, Brussels, Belgium, 1993.
- [BEDAU94A] M. A. Bedau and A. Bahm. Bifurcation structure in diversity dynamics. In R. Brooks and P. Maes, editors, *Artificial Life IV*. Bradford Books/MIT Press, Cambridge, MA, 1994.
- [BEDAU94B] M. A. Bedau and R. Seymour. Adaptation of mutation rates in a simple model of evolution. In R. Stonier and X. H. Yu, editors, *Complex Systems—Mechanisms of Adaptation*. IOS Press, Amsterdam, 1994.
- [BEDAU94C] M. A. Bedau, R. Seymour, M. Zwick, and N. H. Packard. Phases of evolution. Preprint, 1994.
- [BEDAU94D] M. A. Bedau, M. Zwick, and A. Bahm. Measures of diversity in a simple model of evolution. Preprint, 1994.
- [CLIFF93A] D. Cliff, I. Harvey, and P. Husbands. Analysis of evolved sensory-motor controllers. In *Proceedings of the Second European Conference on Artificial Life (ECAL93)*, Brussels, Belgium, 1993.
- [CLIFF93B] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2:73–110, 1993.
- [NOLFI93] S. Nolfi and D. Parisi. Auto-teaching: Networks that develop their own teaching input. In *Proceedings of the Second European Conference on Artificial Life (ECAL93)*, Brussels, Belgium, 1993.

- [PACKARD89] N. H. Packard. Intrinsic adaptation in a simple model for evolution. In C. G. Langton, editor, *Artificial Life*, SFI Studies in the Sciences of Complexity, pages 141–155. Addison-Wesley, Redwood City, CA, 1989.
- [PARISI92] D. Parisi, S. Nolfi, and F. Cecconi. Learning, behavior, and evolution. In F. J. Varela and P. Bourguine, editors, *Towards a Practice of Autonomous Systems*, pages 207–216. Bradford Books/MIT Press, Cambridge, MA, 1992.
- [SHANNON49] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana-Champaign, IL, 1949.
- [STEELS94] L. Steels. The artificial life roots of artificial intelligence. *Artificial Life*, 1:75–110, 1994.
- [VARELA91] F. J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind*. MIT Press, Cambridge, MA, 1991.

A Parameter Values for the Simulations

The world has 128 sites on a side and periodic boundary conditions. Δ produces a pyramidal energy pile; $\Delta(\xi)$ has a peak of 255 at ξ and falls linearly to 0 at a distance of 8 sites from ξ , where ξ^t is chosen at random from the world sites. Sensory information is received from the current site and its four nearest neighbors, with each site yielding one bit of information, so $N_S = 32$. \vec{b} is a displacement from zero to fifteen steps in one of the eight compass directions, so $N_B = 128$. $\alpha(\mathcal{E}(x))$ is a function returning all of $\mathcal{E}(x)$ up to a threshold $\mathcal{E}_m = 100$, $\beta = 20$, and $\gamma(\vec{b})$ is a function returning the magnitude of \vec{b} . $E_r = 500$.