Relevant Information in Optimized Persistence vs. Progeny Strategies

Daniel Polani[1], Chrystopher L. Nehaniv[1], Thomas Martinetz[2] and Jan T. Kim[3]
[1]Algorithms and Adaptive Systems Research Group, University of Hertfordshire, Hatfield, UK
[2]Institut für Neuro- und Bioinformatik, Universität Lübeck, Lübeck, Germany
[3]School of Computing Sciences, University of East Anglia, Norwich UK
d.polani@herts.ac.uk

Abstract

Identifying and utilizing information is central to reproductive success. We study a scenario where a multicellular colony has to trade-off between *utility* of strategies for investment in persistence or progeny and the (Shannon-type) *relevant information* necessary to realize these strategies. We develop a general approach to treat such problems that involve iterated games where utility is determined by iterated play of a strategy and where, in turn, informational processing constraints limit the possible strategies.

Introduction

Organisms have limited information processing capabilities. For instance, in the scenarios for the advent and evolution of multicellularity it is unclear how much relevant information a multicellular aggregate or its component cells would have and be able to capture concerning the state of the whole. This information is likely to be severely limited. In this paper we study the trade-off between information and reproductive success of early multicellular aggregates.

The concept of information as introduced by Shannon is a central resource of communication, computation, physics and life (Shannon, 1949; Bennett and Landauer, 1985; Wheeler, 1990; Laughlin et al., 1998; Adami et al., 2000; Schneider, 2000; Kim et al., 2003). It derives its power and mathematical malleability from the strict rejection of semantical elements in the formalism. Nevertheless, it has long been felt that it would be desirable to treat semantics in a framework similar to information theory (Bar-Hillel, 1964), but for a long time it has been strongly doubted whether this could be achieved at all (Gibson, 1979).

One of the possible reasons for this is that the conventional view of Shannon information favours the picture that information is something that is transported from one point to another as a “bulk” quantity. However, increasingly it has become clear that looking at the intrinsic dynamics of information can provide insight into inner structure of information. Perhaps one of the clearest illustrations for this is provided by the *information bottleneck method* (Tishby et al., 1999) which is a framework for the extraction of a specific part of information out of an information “bulk”. This picture of intrinsic structure of information is increasingly being used to study complex systems (Ay and Wennekers, 2003; Klyubin et al., 2004).

For studies of Artificial Life, however, one particularly important aspect of information is *relevance*. The classic view of information explicitly excludes semantics from the discourse. On the other hand, it turns out that aforementioned picture of structure in information provides a conceptually consistent avenue to introduce semantics into the framework of information theory.

Relevant Information

The opportunity of including relevance into information theory is provided by using an appropriate selection of substructures. In the framework of the information bottleneck method, to study the relevant information in a random variable X, a jointly distributed random variable Y is considered (for more details on notation of random variables, probability and information theory, we refer the reader to the Appendix). Any information of relevance in X is being distinguished only with respect to Y, the *relevance indicator variable*. In their model (Tishby et al., 1999), Y is given externally, like a label in supervised learning. They then define the *relevant information* of X as the mutual information $I(X;Y)$.

To formulate a model of relevant information suitable for Artificial Life studies, however, we would prefer a more intrinsically constructed indicator Y that determines which is the relevant information that should be extracted for a given system. Uexküll’s theory of selective perception argues that environmental stimuli act as triggers for certain behaviour patterns (von Uexküll, 1956a; von Uexküll, 1956b). On a more abstract level, world information is relevant for an agent only if it can serve to guide its actions (Nehaniv, 1999).

Thus, if we model the world state S as a random variable from which relevant information is to be extracted and the agent actions A as the relevance indicator variable, then the relevant information becomes $I(S;A)$, as above; however, at this point, we still need to specify the form of A. For this purpose, (Polani et al., 2001) use utility functions.

The important connection between utility and information has been explored already by (Howard, 1966; Poh and Horvitz, 1996) who quantify the *value* of information by measuring the utility difference attained when some information is available vs. when it is not. However, Howard's value of information is not a Shannon-type quantity; it is given in the particular problem-specific units of the particular utility function of the individual scenario. Relevant information, however, is interpreted here as its conjugate quantity, namely the *amount* instead of the *value* of information, and it is universally measurable in *bits*. To be able to do this, it is necessary to translate the given utility into an instantiation of the random variable A (modeling the actions) in a suitable way as to reflect the utility. How to achieve this will form a major thread in the technical sections further below.

In the present paper, we do not consider any cost involved with obtaining relevant information; our primary focus is how much (relevant) information needs to be acquired (and processed) to achieve a certain utility, ignoring the possible cost for the acquisition process[1]. The inclusion of costs and other constraints governing the information acquisition process in the framework to be presented here is beyond the scope of this paper and will have to be addressed in future work.

Life as an Existential Game

Before turning to the process of translating utility into actions, though, we need to turn to the other important ingredient in the framework of relevant information, namely the origin of utility. In AI or economic scenarios, reward can be explicitly given to achieve a certain goal or obtained by a proper scaling of financial benefits. In Artificial Life systems, however, the formulation of utility is a more intricate problem. Even the emergence of fitness in biology still requires elucidation (Michod, 1999). An intrinsically defined utility is introduced in (Klyubin et al., 2005).

A complementary approach is taken by some evolutionary biologists inspired by game-theory. They formulate life as an "existential game", in which there is no way of using pay-offs for any purpose other than continuing the game for as long as possible, and have analyzed its optimal strategy – minimizing the stakes played — in regard to adaptive responses to environmental perturbations (Slobodkin and Rapoport, 1974).

Considering this strategy at the level of a single (generally multicellular) individual faced with problem of how to allocate resources, a problem arises if continued persistence is attained at the price of not producing progeny. Due to simple thermodynamic considerations, one can expect the individual eventually to die despite continued investment in its own self-production (Varela et al., 1974). Investment of some resource into the production of progeny offers an alternative to mere persistence of the individual (becoming, effectively, an existential game for its lineage). The balance between progeny, which might first require growth of the individual, and persistence presents a problem of crucial importance for an organism maximizing the longevity of its lineage.

Persistence vs. Progeny in Multicellular Organization

The trade-off concerning investment of resources into persistence or into production of progeny is a fundamental issue for living systems. Different organisms solve the problem with different, possibly mixed, strategies (e.g. under r-/K-selection, see (Roughgarden, 1979)). The evolutionary dynamics of differentiated multicellularity has been modelled mathematically by R. E. Michod (Michod, 1999; Michod et al., 2003) to analyse the evolutionary stability and the heritability of fitness under different mutation rates and with different modifier (policing) regimes, but leaves the question of resource conversion into persistence vs. progeny open. Such conversion can be taken as a defining property of living systems (Nehaniv, 2005, Appendix). Here we take a first step to unify consideration of this balance with Michod's and other work on multicellularity.

We consider a scenario of a growing multicellular entity (colony) which has to solve the problem of allocating a number of cells for growth (towards producing progeny) and maintaining the rest of the cells (persistence), based on a resource. We also assume that a threshold size is required for the multicellular individual to be able to become an "adult" and replicate itself (this model e.g. applies to sporulation or dispersal scenarios such as in the cellular slime mold *Dictyostelium discoideum*, or minimal number of cells required for reproduction in viable adult *Volvox* algae).

For simplicity, the number of cells present is identified with the colony's current amount of *resource*; we do not model cells of different types, and moreover (following Michod), we abstract completely away from the organizational topology. Maintenance of existing cells incurs a loss of resource at a given rate (the cost of persistence), while production of additional cells by division increases the number of cells but with the risk that they may not be viable. If the colony grows to a certain threshold size accumulating sufficient resources, then it "succeeds" and is able to reproduce. This success provides a natural measure of utility.

Constructing Strategy from Utility

In the previous sections we developed the general framework of relevant information and its connection to utility, and, on the other hand, introduced a particularly important scenario. In the following, we will develop the details of the techniques necessary to bring them together. The present section will introduce a first motivational approach and formulate two remaining problems; these problems will then be addressed and resolved in the remaining sections.

---

[1] Thus, our scenario should not be confused with the "two-armed bandit" problem (Dubins and Savage, 1965; Holland, 1975).

Above, we viewed relevant information as given by the mutual information $I(S;A)$ between world state variable S and agent action A. To compute it, it is necessary to specify its joint distribution $p(s,a)$. From now on, assume that the probability $p(s)$ for a state s is given a priori. Given $p(s)$ the joint probability $p(s,a) = p(a|s)p(s)$ is fully determined by the *strategy* $\pi = p(a|s)$, the conditional probability distribution on actions a given state s.

To make sense in an agent context, $p(a|s)$ must reflect the needs or preferences of the agent. These preferences can be formalized as the utility $U(s,a)$ of taking an action a in a state s. The utility is a real value, and the agent prefers high to low utility values. In (Polani et al., 2001), from the utility U the authors derive a strategy $p(a^*|s)$ by assuming for each given state s that $p(a^*|s)$ is an equidistribution over the *optimal* actions a^* with respect to U. This model for relevant information has many desirable features. For instance, if the same unique action is optimal in every state, then the relevant information vanishes: one can select the same (optimal) action in each state, and there is no need to elicit any information about the state. If, on the other hand, each state s requires a different optimal action, all states have to be distinguished: the relevant information has the value $H(S)$, as it must cover the full entropy (uncertainty) about the current state to select the right action.

The above choice of the strategy has still two problems: first, it only considers actions which are strictly optimal with respect to the utility. This is not always sensible, as can be seen by the following example: consider a quiz show candidate that can select between two doors, one of which contains a prize. Opening one door results in a payoff of 100,000 units, the other in a payoff of 0 units. Assuming the prize is equally likely to be placed behind either of the two doors, there is 1 bit of relevant information with respect to the choice of the door. On the other hand, if one door reveals 100,000 units, the other door reveals 99,999 units, this formalism will still see exactly 1 bit, because it selects the strictly optimal action. Often, though, the small difference of 1 unit will be deemed irrelevant and thus both doors are almost equally desirable. In that case, relevant information should be close to 0 bit. Although not immediately evident in this toy example, such a consideration plays a role when, for example, information processing capacity is a scarcer resource than utility. As organisms often operate close to the limits of information processing (de Ruyter van Steveninck and Laughlin, 1996), such a trade-off may be necessary for an appropriate biologically relevant model.

The second problem is the following: for different states the choice of actions is not coordinated. For instance, consider a system in two different states with two different actions to select in each. Assume that in state 1 one has optimal action 1 (which is selected with probability 1.0). In state 2, however, let both actions be optimal. According to above strategy, each of them is selected with probability 0.5. If the states are equiprobable, the relevant information required to employ this strategy is nonzero since one has to distinguish the states to identify the action distribution available. Since both states share action 1 as optimal action, one could instead use a strategy with 1 as optimal action for *both* states. While this does not reduce utility, it is more parsimonious in terms of relevant information, as it requires no information to distinguish states at all. Both remaining problems will be resolved in the following sections.

Trading Off Utility and Information

The second of above problems, namely the problem of *informational parsimony* can be formulated as search for a strategy $p(a|s)$ minimizing required relevant information $I(S;A)$, but under the condition that the actions realized by A still have to be optimal for a given state s, formally:

$$I(S;A^*) = \min_{p(a|s):p(a|s)p(s)>0 \Rightarrow a \text{ optimal for } s} I(S;A)\ . \qquad (1)$$

Consider without loss of generality only systems whose states s satisfy $p(s) > 0$. Then a simple argument shows that the strategies $\pi^* = p(a^*|s)$ that are restricted to optimal actions are exactly those for which the expected utility $\mathbf{E}_\pi[U(S,A)] = \sum_{s,a} p(a|s)p(s) \cdot U(s,a)$ attains its maximal value $U_{\max}$. Equation (1) constitutes an informationally parsimonious formulation of required relevant information. Moreover, it forms the basis to relax the restriction of optimality addressing the first problem mentioned above. To this end, in the original constraint considering only strategies π maximizing expected utility, i.e. $\mathbf{E}_\pi[U(S,A)] \overset{!}{=} U_{\max}$, we will replace $U_{\max}$ by some smaller value $\hat{U}$. This corresponds to trading in utility for a reduction in required relevant information, as in the second quiz show example.

The minimization task for $I(S;A)$ under this constraint can be elegantly transformed into an unconstrained minimization problem via the use of Lagrange multipliers (similar to the approach by (Tishby et al., 1999) for the information bottleneck):

$$\min_{p(a|s)} \Big(I(S;A) - \beta \cdot \mathbf{E}[U(S,A)] \Big)\ . \qquad (2)$$

The “inverse temperature” β implicitly encodes the constraint $U_{\max}$. For $\beta \to \infty$, the optimization limits actions to the optimal ones. A reduction of β, on the other hand, makes the optimization less sensitive to differences in utility.

The expression in (2) is formally almost equivalent to the rate-distortion functional known from classical information theory (Cover and Thomas, 1991). Except for the sign, U corresponds to the distortion function of rate-distortion theory comparing a signal at the input of a noisy channel to the distorted result at the output of the channel. We remark that the classical Blahut-Arimoto algorithm (Cover and Thomas, 1991; Tishby et al., 1999) trivially generalizes to our case with its convergence and uniqueness guarantees; we use it to calculate the optimized strategies $p(a|s)$ and the relevant information required for their realization $I(S;A)$ in (2). Without being able to go into details, we mention that the algorithm consists of an alternating self-consistent iteration that incrementally improves the strategy $\pi = p(a|s)$ until convergence.

Persistence vs. Progeny: Strategies

Return now to the original problem of studying strategies in scenarios with a trade-off between persistence and progeny. We use the following model (a generalization of the Gambler's problem (Dubins and Savage, 1965; Sutton and Barto, 1998)).

Consider a colony of s identical cells which we will consider as a superorganism. To sporulate and thus to reproduce, its size needs to reach or exceed a critical number $\bar{s}$. As long as $s < \bar{s}$, the colony is faced with a dilemma: it can "invest" a of its cells towards *progeny*. This part a of the colony has the probability $\bar{p} < 0.5$ of doubling its size successfully in the next time unit, but with probability $1 - \bar{p}$ it will die off entirely. The rest of the colony, $s - a$, is *persisting* and is guaranteed to lose a proportion δ of cells. If δ is smaller than the risk of dying for the progeny cells, one could be tempted to stick with persistence as strategy, in the short run. However, in the long run, this will lead to death of all cells. For larger δ investment into progeny will become increasingly aggressive, in a "desperate" attempt to reach the threshold $\bar{s}$. But even for vanishing decay rates $\delta \to 0$ it is not possible to reach $\bar{s}$ without investment in progeny.

This scenario defines an iterated game. Beginning with an initial colony size of s_0 at time $t = 0$, the colony selects an action a_0 (i.e. invest an amount of colony cells for progeny). With probability $\bar{p}$, the investment will be doubled and the colony will have size $s_{t+1} = (1-\delta)(s_t - a_t) + 2 \cdot a_t$ at time $t + 1$. With counterprobability $1 - \bar{p}$, however, its new size will be $s_{t+1} = (1-\delta)(s_t - a_t)$. The game is iterated until either the colony dies, i.e. $s = 0$, or the colony achieves sporulation, i.e. $s \geq \bar{s}$.

Identify colony size s with the state s of the system. Then, a strategy $\pi = p(a|s)$ gives a probability distribution of actions to take if the colony is in that state. Consider now the probability that a colony starting in state s will succeed in reaching sporulation, following a fixed strategy π throughout its lifetime. We define $V^\pi(s)$, the *value* of the state s under strategy π, to be this probability. Then the utility $U^\pi(s,a)$ for a state-action pair (s,a) is the expected value of the colony state on selecting action a and following π thereafter[2]. The calculation of V and thus U can be solved in the framework of Markov Decision Problems. For V^π, the Bellmann equation holds (Sutton and Barto, 1998):

$$V^\pi(s) = \sum_a p(a|s) \sum_{s'} p(s'|s,a)[r(s,a,s') + V^\pi(s')], \qquad (3)$$

where $p(s'|s,a)$ is the probability of ending in a successor state s' after applying action a in state s in the dynamics of our model, and the reward $r(s,a,s')$ is 0, unless s' is the sporulation goal state, in which case r is given by the full success value 1; the last case corresponds to the fact that in that case the sporulation state is reached with probability 1. Equation (3) is self-consistent, i.e. the same V appears both on the left and the right side of the equation. This can be turned into the well-known *value iteration* algorithm, starting with an arbitrary V, inserting it on the right side, and obtaining a new V on the left, and iterating. It is long known in the theory of dynamic programming that this iteration converges to a unique V^π for given π, providing a numerical method for its computation. From V^π, one can compute the utility U^π via

$$U^\pi(s,a) = \sum_{s'} p(s'|s,a)[r(s,a,s') + V^\pi(s')] \, . \qquad (4)$$

Relevant Information in Iterated Games

In dynamic programming, one then proceeds to compute the optimal strategy by improving π with respect to V^π and recalculating V^π under the improved strategy until convergence. However, here we are not merely interested in an unconstrained optimal strategies, but we have to trade off some utility against the amount of information required to implement a strategy. We wish to find the optimal strategy $\pi^* = p(a^*|s)$ under the constraint that no more than a given amount $I(S;A^*)$ of state information can be used to select actions. As we increasingly constrain this amount of information we allow π to utilize from state s to select the actions a, the strategy will increasingly lose utility. However, at the same time, less relevant information will be needed.

As mentioned, for a given strategy π one can compute V^π through the iteration of (3), and, from this, utility U^π. On the other hand, given a utility U, one computes the strategy π' minimizing the relevant information for a given utility level. Thus, one has the following process:

$$\pi \xrightarrow{(3)} V^\pi \xrightarrow{(4)} U^\pi \xrightarrow{(2)} \pi' \qquad (5)$$

As long as one only considers optimal strategies ($\beta \to \infty$), this is not a problem: the optimal utility U^* can be calculated separately and then used to optimize (2). In general, however, one is interested also in less-than-optimal utilities. Here the solution is more intricate: we calculate consistent strategy-utility pairs for general β by iterating (5), i.e. interlacing (3) and (2) to obtain a self-consistent π-U pair. The resulting double iteration appears to converge to a unique solution in all examples considered, though at this point we

---

[2] In the literature for Markov Decision Problems, U is instead called the Q function. Here we use the notation U in agreement with (Bertsekas, 1976) in emphasis of the utility function aspect.

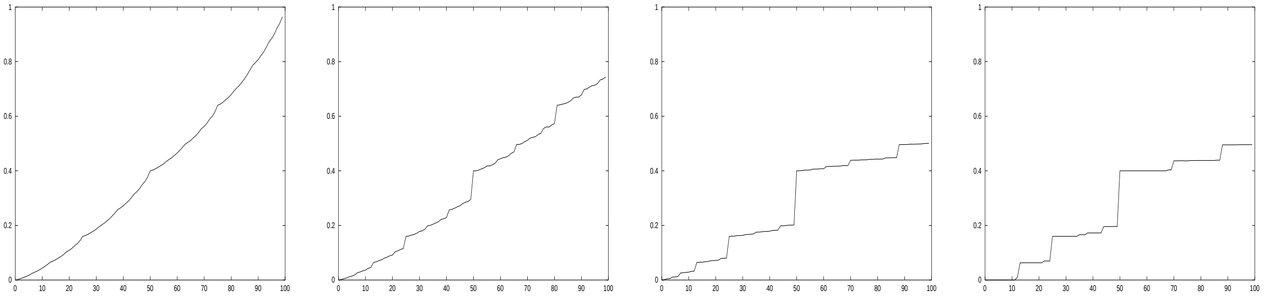

(a) $\mathbf{E}V^{\pi}(S) \approx 0.40$ (b) $\mathbf{E}V^{\pi}(S) \approx 0.35$ (c) $\mathbf{E}V^{\pi}(S) \approx 0.28$ (d) $\mathbf{E}V^{\pi}(S) \approx 0.27$

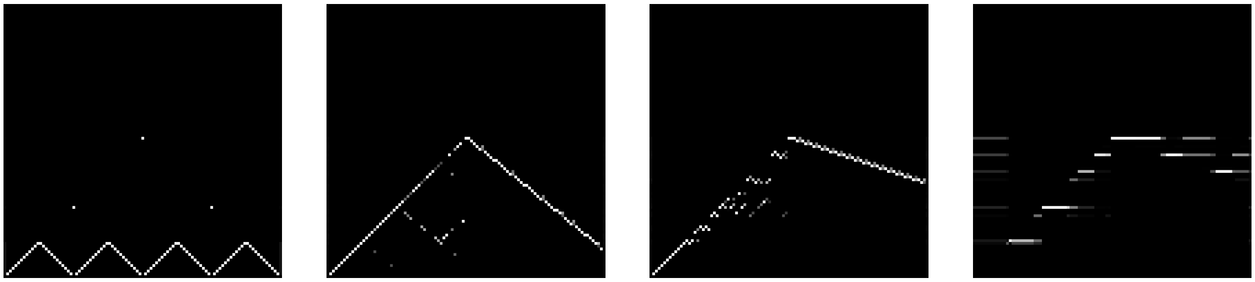

(e) Optimal Strategy, $\delta = 0$, $I(S;A^*) \approx 3.70\,\mathrm{bit}$

(f) Optimal Strategy, $\delta = 0.1$, $I(S;A^*) \approx 5.36\,\mathrm{bit}$

(g) Optimal Strategy, $\delta = 0.5$, $I(S;A^*) \approx 4.84\,\mathrm{bit}$

(h) "Soft" Strategy, $\beta = 100$, $\delta = 0.5$, $I(S;A^*) \approx 1.73\,\mathrm{bit}$

Figure 1: Values and persistence vs. progeny strategies for different parameter settings. The top line shows the value functions $V^{\pi}(s)$ and its expectation value for different scenarios and parameters, the bottom line shows the corresponding strategies π. $V^{\pi}(s)$, the probability that the colony will achieve sporulation under the given strategy, is plotted vs. the state $s = 0,\dots,99$. The strategy π is visualized as a matrix of brightness values, where the horizontal axis denotes the state s, the vertical denotes the actions a and the brightness indicates the probability $p(a|s)$, where black is probability 0, and white is probability 1.

do not possess an analytical proof. The importance of this approach is that it allows us to compute optimized strategies and their required relevant information for general iterated single-player games, while trading off utility and relevant information.

Results

We now have the full formalism to treat the persistence vs. progeny scenario with respect to relevant information. We consider a number of scenarios, with s an integer, $\bar{s} = 100$, $\bar{p} = 0.4$ and varying decay rates δ for the persisting part of the colony. For vanishing decay rates $\delta = 0$, the optimal strategy (corresponding to $\beta \to \infty$) can be calculated analytically (Dubins and Savage, 1965). For nonvanishing δ, optimal strategies can be calculated using the value iteration procedure described above. Once β is small, the interlacing algorithm introduced in last section has to be used. We will also use it consistently to find (near-)optimal strategies by setting $\beta \to \infty$ (in practice, $\beta = 100{,}000$) and over a wide range of smaller β, iterating until $\sum_s |V_{\mathrm{new}}(s) - V_{\mathrm{old}}(s)| < 10^{-5}$ for the value functions of two successive iterations.

Figure 1 shows some value functions and strategies for the persistence vs. progeny scenario. In Figure 1(e), one can see the hierarchical structure of the optimal strategy for $\delta = 0$. For $\delta > 0$, the optimal strategy needs to distinguish more states (Fig. 1(f)- 1(g)), as reflected by the increase of $I(S;A^*)$. Trading off even only 0.01 of expected success probability $\mathbf{E}V^{\pi}$ reduces the required information by more than 60% (Fig. 1(g) and 1(h)).

Figure 2 shows the trade-off curve between achievable utility for a given amount of utilized relevant information. Analyzing this trade-off in the scenario with $\delta = 0.5$, the optimal strategy (as in Fig. 1(g)) requires a maximum of $4.84\,\mathrm{bit}$, achieving an expected utility of 0.28. It should be noted that here, even on reduction of available relevant information, the achievable expected utilities do not drop significantly. In particular, it is possible to find policies that reduce the required relevant information to almost $I(S;A^*) = 0\,\mathrm{bit}$

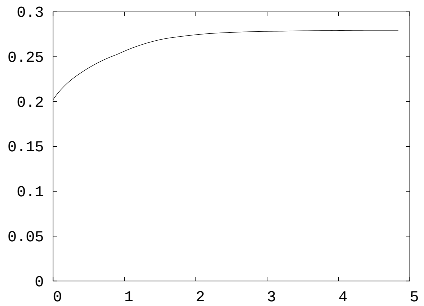

Figure 2: Trade-off between expected utility and relevant information for the persistence vs. progeny game. The expected utility $\mathbf{E}V^{\pi}$ on the vertical axis is plotted against the relevant information $I(S;A^*)$ (in bits) on the horizontal axis; it shows which expected utility can at most be achieved by policies utilizing a given amount of relevant information.

while still achieving a utility of 0.20, i.e. losing less than 30% of the maximum utility. We emphasize again that in this study we are not considering the cost of acquiring information; instead, we consider the highest utility that can be achieved for a given amount of acquired (and processed) relevant information.

Discussion and Conclusions

We presented a general constructive approach to treat information-utility trade-off problems in iterated games. This allows one to calculate optimized strategies in temporally extended problems such as the persistence vs. progeny resource allocation problem, under different balances between informational resolution and utility. The aim is always to achieve an as parsimonious utilization of information as possible for a given level of utility.

Note that, in the persistence vs. progeny scenario, actions are modeled at the multicellular colony level rather than the individual cell level, which is *not* equivalent to strategies played by individual cells independently. In particular, this reflects coordinated information processing within the multicellular entity. Note that in general such processing might be disrupted by adverse mutations in constituent cells.

The conclusion from (Michod, 1999) is that small propagule sizes can mitigate such adverse effects on the heritability of fitness at the higher level. The current model effectively assumes an equidistribution over the possible states s. In future work, we intend to take into account the propagule size after reaching the threshold $\bar{s}$. In particular, propagules consisting of a single cell will require a modification of the a priori distribution $p(s)$ compared to the present model.

Appendix: Notation

Random variables are denoted with capital letters X,Y,Z, assuming a finite set of values x,y,z. Instead of writing $\Pr(X=x)$ we write, by abuse of notation $p(x)$ for the probability that variable X assumes the value x; alternatively, $p(x)$ may denote the entire probability distribution. Notation $p(x,y)$ is used for jointly distributed variables. The conditional probability that Y assumes value y given that X assumes the value x is written $p(y|x)$.

Define *entropy*, a measure of uncertainty of a random variable X, by $H(X) := -\sum_x p(x)\log p(x)$, log taken base 2, and thus $H(X)$ is measured in *bits*. The mutual information between X and Y is given by $I(X;Y) = H(X) + H(Y) - H(X,Y)$ and measures how much uncertainty is lost (information gained) about Y if X is being observed.

References

Adami, C., Ofria, C., and Collier, T. C. (2000). Evolution of biological complexity. *Proc. Natl. Acad. Sci. USA*, 97:4463–4468.

Ay, N. and Wennekers, T. (2003). Dynamical properties of strongly interacting markov chains. *Neural Networks*, 16(10):1483–1497.

Bar-Hillel, Y. (1964). *Language and Information*. Addison-Wesley, Reading, Mass.

Bennett, C. H. and Landauer, R. (1985). The fundamental limits of computation. *Scientific American*, pages 48–56.

Bertsekas, D. P. (1976). *Dynamic Progamming and Stochastic Control*, volume 125 of *Mathematics in Science and Engineering*. Academic Press, Inc., New York, first edition.

Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley, New York.

de Ruyter van Steveninck, R. R. and Laughlin, S. B. (1996). The rate of information transfer at graded-potential synapses. *Nature*, 379:642–645.

Dubins, L. E. and Savage, L. J. (1965). *How To Gamble If You Must*. McGraw-Hill, New York.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin Company, Boston.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor.

Howard, R. A. (1966). Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, SSC-2:22–26.

Kim, J. T., Martinetz, T., and Polani, D. (2003). Bioinformatic principles underlying the information content of transcription factor binding sites. *Journal of Theoretical Biology*, 220:529–544.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2004). Organization of the information flow in the perception-action loop of evolved agents. In *Proceedings of 2004 NASA/DoD Conference on Evolvable Hardware*, pages 177–180. IEEE Computer Society.

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005). Empowerment: A universal agent-centric measure of control. In *IEEE Congress on Evolutionary Computation, 2-5 September 2005, Edinburgh, Scotland*, pages 128–135. IEEE.

Laughlin, S. B., de Ruyter van Steveninck, R. R., and Anderson, J. C. (1998). The metabolic cost of neural information. *Nature Neuroscience*, 1(1):36–41.

Michod, R. E. (1999). *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton University Press.

Michod, R. E., Nedelcu, A. M., and Roze, D. (2003). Cooperation and conflict in the evolution of individuality. IV. Conflict mediation and evolvability in *Volvox carteri*. *BioSystems*, 69(2-3):95–114. Special Issue on Evolvability (Ed.: C. L. Nehaniv).

Nehaniv, C. L. (1999). Meaning for observers and agents. In *Proc. IEEE International Symposium on Intelligent Control / Intelligent Systems and Semiotics, ISIC/ISAS'99, September 15-17, 1999 Cambridge, Mass.*, pages 435–440.

Nehaniv, C. L. (2005). Self-replication, evolvability and asynchronicity in stochastic worlds. In Lupanov, O. B., Kasim-Zade, O. M., Chaskin, A. V., and Steinhöfel, K., editors, *Stochastic Algorithms: Foundations and Applications, LNCS*, volume 3777, pages 126–169. Springer.

Poh, K. L. and Horvitz, E. (1996). A graph-theoretic analysis of information value. In Horvitz, E. and Jensen, F., editors, *Proc. Twelfth Conf. on Uncertainty in Artificial Intelligence*, pages 427–435, San Francisco, CA. Morgan Kaufman.

Polani, D., Martinetz, T., and Kim, J. (2001). An information-theoretic approach for the quantification of relevance. In Kelemen, J. and Sosik, P., editors, *Advances in Artificial Life (Proc. 6th European Conference on Artificial Life)*, volume 2159 of *LNAI*, pages 704–713. Springer.

Roughgarden, J. (1979). *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. Macmillan Publishing Company, New York.

Schneider, T. D. (2000). Evolution of biological information. *Nucleic Acids Research*, 28:2794–2799.

Shannon, C. E. (1949). The mathematical theory of communication. In Shannon, C. E. and Weaver, W., editors, *The Mathematical Theory of Communication*. The University of Illinois Press, Urbana.

Slobodkin, L. B. and Rapoport, A. (1974). An optimal strategy of evolution. *The Quarterly Review of Biology*, 49(3):181–200.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning*. MIT Press, Cambridge, Mass.

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*.

Varela, F., Maturana, H., and Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *BioSystems*, 5:187–196.

von Uexküll, J. (1956a). *Bedeutungslehre*. Rowohlt, Hamburg.

von Uexküll, J. (1956b). *Streifzüge durch die Umwelten von Tieren und Menschen*. Rowohlt, Hamburg.

Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In Zurek, W. H., editor, *Complexity, Entropy and the Physics of Information*, Santa Fe Studies in the Sciences of Complexity, pages 3–28, Reading, Mass. Addison-Wesley.