

An Information-Theoretic Approach for the Quantification of Relevance

Daniel Polani, Thomas Martinetz, and Jan Kim

Institute for Neuro- and Bioinformatics, University of Lübeck, Germany
{polani,martinetz,kim}@inb.mu-luebeck.de

Abstract. We propose a concept for a Shannon-type quantification of information relevant to a decision unit or agent. The proposed measure is operational, can – at least in principle – be calculated for a given system and has an immediate interpretation as an information quantity. Its use as a natural framework for the study of sensor evolution is discussed.

1 Introduction

1.1 Shannon Information and Semantics

Shannon’s seminal study of the transmission capacity of a (possibly noisy) channel introduced one of the most important notions of modern science, namely a quantification of information and of information transmission capability. Shannon’s notion of *entropy* of a memory-less source has deep connections to the notion of entropy known from physics and to a generalized principle of insufficient reason which can be extended to a Bayesian framework (Jaynes 1957a,b; Reichl 1980).

One of the important steps taken by Shannon was to do away with the notion of semantics when dealing with transmitted data. In the Shannon view, when a sequence of symbols is transmitted, every symbol carries the same importance; in particular, there is no additional structure on the set of symbols that can be transmitted. Only their probability of occurrence plays a role in the Shannon model.

The non-interpretation of symbols is one of the strengths of the Shannon model and allows it to introduce a powerful and universal notion of information. However, it has long been felt that, for the purposes of understanding the more complex levels of information processing in natural and artificial computation, and of better understanding natural and artificial language, a variant of Shannon’s information would be in place that would take into account the semantic content of the symbol sequences transmitted.

As Shannon’s approach is not straightforward to generalize in this direction, in the last decades approaches for a formal characterization or quantification of semantics and relevance have been the topic of many discussions. It has been attempted to treat this problem e.g. on the level of philosophy (for a concept of relevance based upon a philosophy of human communication, see e.g. Sperber

and Wilson 1995), or on the level of conceptual foundations of mathematics (e.g. predicate logic, see Bar-Hillel 1964). However, these approaches lack the interpretability, the conceptual crispness and the universality of the original Shannon approach. What is frequently sought is just a simple operational quantification of “semantic” information, conceptually similar to Shannon’s information, but differing in its treatment of semantics and, in particular, being able to quantify *relevance*.

1.2 Adaptive Sensors, POMDPs and Relevant Information

A new incentive to develop a quantification of *semantic*, or as we will prefer to say here, *relevant* information has been brought forward due to the efforts to study adaptive and evolving sensors. The study of sensor evolution is a branch of Artificial Life that lies at the focal point between models to understand biological sensors (Shaaban et al. 1998; Liese et al. 2000), to construct evolving hardware (Lee et al. 1996; Lund et al. 1997) and between fundamental questions concerning the capabilities of evolution regarding the adaptation to the environment (Pask 1959; Cariani 1993; Menczer and Belew 1994; Mark et al. 1998) or the development of selective perception of environmental stimuli acting as triggers for certain behavior patterns (von Uexküll 1956a,b). It is no surprise that this topic has obtained increased interest in recent work (Dautenhahn, Polani, Uthmann 2001).

In the field of sensor evolution, one is interested in agents that adapt or evolve to solve a certain task. This task can be e.g. maximizing survival time or the number of offspring. For this purpose, the agents are expected to evolve appropriate sensors. Such sensor optimization procedures can be driven by explicit sensor quality measures (Nilsson and Pelger 1994). Often, however, the task is implicitly formulated and there is no a priori knowledge about which sensor properties are going to be relevant for the agents to be able to solve it.

The question important in this context is now when sensors will evolve that attempt to capture as much information about the agent environment as possible and when the sensors will be limited to capture only a very restricted bandwidth of information channels. This intuitively depends, among other factors, on what amount of information in the environment is *relevant* for the agent. What we are interested in is to give a precise quantification of this amount of information.

This view is also closely related to the study of partially observable Markovian decision processes (POMDPs, Kaelbling et al. 1996), as also in those it is of particular interest to quantify how much information from the environment is needed to form an adequate decision.

These considerations show that quantifying *relevant* information is not an abstract or exotic requirement, but promises to provide a central measure for the quality of information processing in a system, which is an essential aspect in the studies of Artificial Life systems.

2 Quantification Approaches

To quantify the relevance of information or, alternatively, the amount of semantics or relevance in a given data channel, it has become clear that the semantics has to be brought in by recipient of the data (Wittgenstein 1958; Nehaniv 1999). Nehaniv considers the meaning to arise from the “decoding” of information by given agents. An agent is here seen as an automaton where equivalent states with respect to the behavior of the agent or with respect to its goals are collapsed into a single state (see e.g. Crutchfield 1994). In this model, the external data influence the agent automaton (for instance, to attain certain goal states).

A framework for the modeling of strategies to achieve goals is provided by decision theory (Bertsekas 1976) which also underlies the theory of Markovian Decision Problems (Sutton and Barto 1998). In the context of decision theory, a theory of information value has been developed (Howard 1966; Poh and Horvitz 1996). A utility function is a suitably normalized measure that quantifies the reward obtained by taking a certain action in a given state. Solving the decision problem means finding that action that gives the highest expected reward. If the exact state is not perfectly known to the decision system, the attained reward will typically drop below that of the fully informed decision system. In general, given a piece of information (e.g. in form of a random variable that is correlated with the true state variable of the system), it will be possible to attain a higher expected utility than without it. The difference between the two attained values is the information value of the variable according to information value theory. Note that this quantifies the *value* of the information which, in turn, is measured in the units the utility function is given in, be it monetary value, fitness, score or other.

For a quantified notion of *relevant* information in the vein of the Shannon sense, however, different properties would be desirable. In particular, one would like to have a measure that is able to be interpreted in a sense similar to the classical Shannon information; loosely speaking, its value should be expressed in *bits*, the “universal unit of information” and not in the arbitrary units of some utility measure.

Nehaniv (1999) grounds the meaning of Shannon-type information in the usefulness of signals detected by sensors and of actuators manipulating the world and producing signals. This usefulness is there put into the context of fitness (e.g. reproductive success) or also to more short-term goals like homeostasis. In particular, Nehaniv considers this concept in relation to sensor and actuator evolution; as example, he discusses the trade-off of cost and benefit versus the signal length for information transmission evolved for the signaling behavior by squids. In addition, different meanings and, thus, different degrees of informativity can arise for different observers in the same signal channel, even in the same context, depending on the respective goals perceived (von Uexküll 1956b; Nehaniv et al. 1999).

Like Nehaniv, we will relate the amount of relevant information present in a channel to the usefulness to the agent that can be derived from its knowledge. To do so, we here formalize the concept of usefulness via the framework of decision

theory. In addition, we develop a formalism that allows to model the actuator channel as a random variable and to compute the mutual information between sensor and actuator channel which then will serve as measure for the amount of relevant information.

3 A Measure for Relevant Information

First, we introduce some notation. We will denote random variables by uppercase letters (e.g. X), the values they can assume by lowercase letters (e.g. x), the set of these values by calligraphic letters (e.g. \mathcal{X}). The probability that a random variable X assumes a value $x \in \mathcal{X}$ will be denoted by $P(X = x)$ or, for simplicity by $p(x)$ by abuse of notation. Likewise, we will write $p(y|x)$ instead of $P(Y = y|X = x)$ for the conditional probability for Y assuming the value y , given that X assumes the value x .

To keep the mathematics simple, we will always assume the random variables assume a finite number of (therefore discrete) values. The *entropy* of a random variable X is defined as

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x) , \quad (1)$$

where for convenience we assume the binary logarithm. The conditional entropy $H(Y|X)$ of a random variable Y given X is defined as

$$H(Y|X) := \sum_{x \in \mathcal{X}} p(x) H(Y|X = x) \quad (2)$$

$$:= - \sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log p(y|x) . \quad (3)$$

To construct a measure for the relevant information present in a random variable X , Tishby et al. (1999); Kaski and Sinkkonen (2000) assume the existence of some relevance indicator variable Y that is jointly distributed with X . Then they define the relevant information in X via the *mutual information*

$$I(X; Y) = H(Y) - H(Y|X) = H(X) - H(X|Y) \quad (4)$$

between the variables X and Y . The mutual information quantifies (in a Shannon sense) how many bits of uncertainty about Y can be resolved by knowing X . In their work, the relevance indicator variable has to be given a priori, e.g. by appropriate labeling of training samples for X or by human intervention. The relevance is determined by Y . In other words, the model is to determine how many relevant information (about Y) can be extracted from knowing X .

We introduce now our notion of *relevant information* by instantiating the relevance indicator variable Y from the models of Tishby et al.; Kaski and Sinkkonen with a specific, decision-theoretic model (Polani et al. 2001), grounding, as Nehaniv (1999), the relevance of an action in its utility to the agent.

We assume a decision system with (for simplicity of exposition) finite state and action spaces. Our key assumption is that for a given state x , there exists a set $\mathcal{Y}^*(x)$ of equivalently optimal actions $\mathcal{Y}^*(x) := \{y_{i_1}^* \dots y_{i_r}^*\} \subseteq \mathcal{Y} = \{y_1 \dots y_k\}$, where \mathcal{Y} is the set of possible actions. At this point we make no assumptions about the model that selects the optimal actions; one could, for instance, imagine the optimal actions to be maximizing the expectation value of some suitable utility function. Given that state x , we define a conditional probability $p(y|x)$ of selecting an action y by $p(y|x) := 1/|\mathcal{Y}^*(x)|$ if $y \in \mathcal{Y}^*(x)$, else setting $p(y|x) := 0$, where $|\cdot|$ denotes the number of elements in a set. In other words, given a state x , then the conditional probability $p(y|x)$ that an optimal action y will be selected arises from the principle of insufficient reason (or *maximum entropy*, Jaynes 1957a,b) if several optimal actions exist. We call $p(y|x)$ the *action selection model*.

Let now $p(x)$ be the probability that the system is in state x ; denote the respective random variable by X . Define now the random variable Y^* in such a way that it is jointly distributed with X via $p(x, y) = p(y|x)p(x)$, where the $p(y|x)$ is the action selection model given above. Then Y^* denotes a random variable modeling the selection of an optimal action. Then, using Y^* as relevance indicator variable, we define in the sense of Tishby et al.; Kaski and Sinkkonen the relevant information in X for the decision system by the mutual information $I(X; Y^*)$ between the state variable X and the action selection variable Y^* .

4 Some Simple Examples

In this section, some very simple prototypical examples are given that illustrate how the notion of relevant information operates. For this purpose, a couple of instructive constellations are investigated. We generally consider a simple scenario where the decision maker can assume 8 possible states x and can select between 8 possible actions y in each state x . Given a set $\mathcal{Y}^*(x)$ of optimal actions for state x , the action selection model from Sec. 3 becomes

$$p(y|x) = P(Y^* = y | X = x) = \frac{1}{|\mathcal{Y}^*(x)|} \mathbb{I}_{\mathcal{Y}^*(x)}(y), \quad (5)$$

where $\mathbb{I}_{\mathcal{Y}^*(x)}$ is the characteristic function of $\mathcal{Y}^*(x)$ which is 1 for $y \in \mathcal{Y}^*(x)$ and 0 else.

Consider first the situation in Fig. 1(a). The vertical axis denotes the state of the (decision making) agent, the horizontal axis the possible actions. A given state x defines a row, in which darkened squares indicate the action selection model. In the example of Fig. 1(a), for each distinct state there is a unique distinct optimal action to take. One obtains $p(y|x) = \delta_{x, f(x)}$ with a suitably chosen function f . It is easy to see that, in this case, $H(Y^*|X) = 0$ since the optimal action variable Y^* is perfectly determined by knowing the value of X . The relevant information becomes $I(X; Y^*) = H(Y^*) - H(Y^*|X) = H(Y^*) = H(X)$ (the last equation holds because f is injective, i.e. $f(x') \neq f(x'')$ for




Fig. 1. Illustration of the operation of relevant information. The darkened squares indicate the optimal actions y for a state x . The depth of the shading represents the value $p(y|x)$ of the action selection model, where black denotes a value of 1.

$x' \neq x''$). Thus, full information about the state X is relevant since each state x requires its individual optimal response.

In Figure 1(b), this is different. It is still possible to consider the (unique) optimal action as function $y = f(x)$ of the present state x , but f is no longer injective. The relevant information is still given by $H(Y^*)$, but now $H(Y^*) < H(X)$. For equiprobabilistic X , one obtains $I(X; Y^*) = H(Y^*) - 0 \text{ bit} = 1 \text{ bit}$ as relevant information. The full uncertainty of the present state is $H(X) = 3 \text{ bit}$. This means that to find a relevant action, the current state has to be determined only to the precision of 1 bit and not to its full precision of 3 bit.

Fig. 1(c) somewhat reminds of Fig. 1(a), but now for each state x , there is a selection of four different actions that are equivalently optimal. In this case, the action selection model becomes $p(y|x) = 1/4$ for an action y that is optimal for x and $p(y|x) = 0$ else. Easy calculation shows that the relevant information becomes $I(X; Y^*) = H(Y^*) - H(Y^*|X) = 3 \text{ bit} - 2 \text{ bit} = 1 \text{ bit}$, assuming X is equiprobabilistic. This fits with intuition, as one has more freedom to select an action and less specific state information is required to decide which actions to take. Thus, less relevant information is therefore attributed to the system.

5 Discussion

Above definition of relevant information is natural in many senses. It has an immediate interpretation as an information quantity of Shannon type. Via the decision mechanism it addresses and incorporates the subjectivity of relevance (i.e. it takes into account the question “*relevant for what purpose?*”). When there is no available choice in actions, i.e. always the same unique action is optimal, the relevant information vanishes. In this case, it is of no value knowing the state of the system.

The same holds if there is no correlation between the state and the possible actions that can be chosen, e.g. if the action selection model $p(y|x)$ is identical for




Fig. 2. Relevant information and information acquisition via a sensoric channel. See text for details.

all x . This fits intuition that, if, by knowing X , the choice of Y^* is not affected, there is no relevant information in X with respect to Y^* .

The potential of the approach lies in the fact that, given an action selection model, it does not only provide a very natural quantification of the amount of relevant information given a certain state, but it allows also the study of the effect of incomplete state information. For illustration, assume the (true) state variable X of the agent to be perceived by the agent not directly, but through a sensorial channel which is modeled by a random variable \hat{X} , obtained from X via some probabilistic mapping $p(\hat{x}|x)$, analogously to the scenario of Tishby et al.’s information bottleneck method (Fig. 2). For the selection of an action, the total amount of relevant information is given by $I(X; Y^*)$. However, the “useful” information available to the agent via the sensoric state variable \hat{X} is only $I(\hat{X}; Y^*)$. Elementary information theory allows to show that $I(\hat{X}; Y^*) \leq I(X; Y^*)$. The relevant information $I(X; Y^*)$ therefore serves to gauge the maximum amount of useful information that the sensorial channel can possibly extract from the present state of the agent.

A paradigmatic instance of such a scenario is the adaption or evolution of sensors. In this context, $I(\hat{X}; Y^*)$ as compared to $I(X; Y^*)$ provides a measure for the usefulness of a sensory variable for an agent to solve the task at hand. If $I(X; Y^*)$ is already small (which can e.g. be the case even for large state spaces if the set of optimal actions is small) then even a simple and thus cheap sensor may capture most of the relevant information in the environment and this amount can instantly be quantified by $I(\hat{X}; Y^*)$. The trade-off between the bandwidth of the sensor channel and the overall relevant information can be used to fine-tune the necessary resources. Here, the information bottleneck framework by Tishby et al. (1999) provides powerful tools for theoretical studies of the setting. In addition, we emphasize that the concept of relevant information developed here is immediately applicable to the study of POMDPs (Kaelbling et al. 1996). POMDPs, in turn, we believe to have deep and intricate connections to the field of sensor evolution whose exploration is still to be undertaken.

6 Summary and Future Work

A framework for the quantification of relevant information has been presented. Combining decision- and information-theoretic aspects, the concept incorporates many properties that are intuitively expected from such a quantity. It quantifies relevance with respect to a given agent or decision system, it yields a measure for the usefulness of sensors or about an agent’s state of knowledge with respect to a POMDP. The concept introduces an operational notion of a Shannon-type quantification of relevant information that can, in principle, always be computed.

We intend to apply the notion of relevant information that has been introduced here on framework level to a couple of sensor evolution scenarios in the near future. Abstracted sensor evolution models can be easily interpreted in the terms developed in the present paper. They provides a suitable framework to better understand the mechanisms that govern sensor evolution and to pursue methods of improving it that may be generalizable to hardware. Jung et al. (2001b) developed a simple but instructive sensor evolution model in whose analysis our present concept of relevant information was used (Jung et al. 2001a).

On a conceptual level, there are several interesting aspects by which we plan to extend our model. First, the notion of relevant information can be strengthened to quantify the *minimum* required information about the system state X to choose a right action Y^* . This view slightly changes the method how the action selection model is constructed and the consequences of this (slightly more complicated) modification are to be explored.

A second important aspect is the fact that often there may be few exactly optimal actions, but a large selection of suboptimal actions which lead almost to the same reward as the optimal one. Thirdly, it is in general not easy to identify precisely the optimal action to take in a given state. The decision system will usually have only a limited model about which action may be the best.

Interestingly, the last two problems can be solved at once by attacking them simultaneously. One can regard the decision system or agent as an estimator for the optimal action to take. The estimator, e.g. being based on some learning or evolution system, will be trained by experimental data. It will typically give imperfect estimations about which actions are optimal, depending on the history of its training. The action selection model $p(y|x)$ then reflects this imperfection. As optimal and suboptimal rewards converge, by virtue of the continuity of the estimator will cause also the action model to converge in a way that is consistent with the specialized definition from Sec. 3. The imperfect model \hat{Y}^* for the action to be taken is then be regarded as an imperfect “sensor” for the best possible action, in analogy to the sensor part. A forthcoming paper will develop this generalization of the model in detail.

Acknowledgements

We wish to express our gratitude to the anonymous reviewers for their detailed and helpful comments.

Bibliography

- Bar-Hillel, Y., (1964). *Language and Information*. Addison-Wesley.
- Bertsekas, D. P., (1976). *Dynamic Programming and Stochastic Control*, vol. 125 of *Mathematics in Science and Engineering*. New York: Academic Press, Inc. First edition edition.
- Cariani, P., (1993). To evolve an ear: epistemological implications of Gordon Pask's electrochemical devices. *Systems Research*, 10(3):19–33.
- Crutchfield, J. P., (1994). Observing Complexity and the Complexity of Observation. In Atmanspacher, H., editor, *Inside versus Outside*, 234–272. Springer.
- Dautenhahn, Polani, Uthmann, (2001). Artificial Life Journal — Special Issue on Sensor Evolution. *Artificial Life Journal — Special Issue on Sensor Evolution*, 7(2). (in press).
- Howard, R. A., (1966). Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, SSC-2:22–26.
- Jaynes, E. T., (1957a). Information theory and statistical mechanics. *Phys. Rev.*, 106(4):620–630.
- Jaynes, E. T., (1957b). Information theory and statistical mechanics II. *Phys. Rev.*, 108(2).
- Jung, T., Dauscher, P., and Uthmann, T., (2001a). On Individual Learning, Evolution of Sensors and Relevant Information. In Polani, D., Uthmann, T., and Dautenhahn, K., editors, *Proc. GECCO-2001 Workshop on Evolution of Sensors in Nature, Hardware and Simulation*.
- Jung, T., Dauscher, P., and Uthmann, T., (2001b). Some Effects of Individual Learning on the Evolution of Sensors. In *Proc. European Conference on Artificial Life (ECAL 2001, Prague)*.
- Kaelbling, L. P., Littman, M. L., and Moore, A. W., (1996). Reinforcement learning: a survey. *Journal of Artificial Intelligence Research*, 4:237–285.
- Kaski, S., and Sinkkonen, J., (2000). Metrics that learn relevance. In *Proc. IJCNN-2000*, vol. V, 547–552. IEEE Service Center, Piscataway, NJ.
- Lee, W.-P., Hallam, J., and Lund, H. H., (1996). A Hybrid GP/GA Approach for Co-evolving Controllers and Robot Bodies to Achieve Fitness-Specified Tasks. In *Proc. IEEE 3rd Int. Conference on Evolutionary Computation*. NJ: IEEE Pr.
- Liese, A., Polani, D., and Uthmann, T., (2000). On the development of spectral properties of visual agent receptors through evolution. In Whitley, D., Goldberg, D., Cantú-Paz, E., Spector, L., Parmee, I., and Beyer, H.-G., editors, *Proc. Genetic and Evolutionary Computation Conference (GECCO), Las Vegas, Nevada*, 857–864. Morgan Kaufmann.
- Lund, H. H., Hallam, J., and Lee, W.-P., (1997). Evolving Robot Morphology. In *Proc. IEEE 4th Int. Conference on Evolutionary Computation*. NJ: IEEE Press.

- Mark, A., Polani, D., and Uthmann, T., (1998). A Framework for Sensor Evolution in a Population of Braitenberg Vehicle-like Agents. In Adami, C., Belew, R., Kitano, H., and Taylor, C., editors, *Proc. of Artificial Life VI, Los Angeles, June 26-29*, 428–432. MIT Press.
- Menczer, F., and Belew, R. K., (1994). Evolving Sensors in Environments of Controlled Complexity. In Rodney A. Brooks, P. M., editor, *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, 210–221. MIT Press.
- Nehaniv, C. L., (1999). Meaning for Observers and Agents. In *Proc. IEEE International Symposium on Intelligent Control / Intelligent Systems and Semiotics, ISIC/ISAS'99*, 435–440. Cambridge, Massachusetts, USA.
- Nehaniv, C. L., Dautenhahn, K., and Loomes, M. J., (1999). Constructive Biology and Approaches to Temporal Grounding in Post-Reactive Robotics. In McKee, G. T., and Schenker, P., editors, *Sensor Fusion and Decentralized Control in Robotics Systems II, Boston, Mass. (Proc. SPIE Vol. 3839)*, 156–167.
- 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.
- Pask, G., (1959). Physical Analogues to the Growth of a Concept. In *Mechanisation of Thought Processes: Proceedings of a Symposium held at the National Physical Laboratory (No. 10)*, vol. II, 877–928. London: Her Majesty's Stationary Office.
- 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*, 427–435. San Francisco, CA: Morgan Kaufman.
- Polani, D., Martinetz, T., and Kim, J., (2001). On the Quantification of Relevant Information. Presented at SCAI'01 (Scandinavian Conference on Artificial Intelligence), Feb. 19-21, 2001.
- Reichl, L., (1980). *A Modern Course in Statistical Physics*. Austin: University of Texas Press.
- Shaaban, S. A., Crognale, M. A., Calderone, J. B., Huang, J., Jacobs, G. H., and Deeb, S. S., (1998). Transgenic Mice Expressing a Functional Human Photopigment. *Investigative Ophthalmology & Visual Science*, 39(6):1036–1043.
- Sperber, D., and Wilson, D., (1995). *Relevance: Communication and cognition*. Oxford: Blackwell. Second edition.
- Sutton, R. S., and Barto, A. G., (1998). *Reinforcement Learning*. Bradford.
- 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*.
- von Uexküll, J., (1956a). *Bedeutungslehre*. Hamburg: Rowohlt.
- von Uexküll, J., (1956b). *Streifzüge durch die Umwelten von Tieren und Menschen*. Hamburg: Rowohlt.
- Wittgenstein, L., (1958). *The Blue and Brown Books*. Harper & Brothers.