

Towards Closed Loop Information: Predictive Information

Bernd Porr ◇ University of Glasgow, b.porr@elec.gla.ac.uk

Alice Egerton ◇ Imperial College London, aegerton@imperial.ac.uk

Florentin Wörgötter ◇ University of Göttingen and University of Stirling, worgott@chaos.gwdg.de

Motivation: Classical definitions of information, such as the Shannon information, are designed for open loop systems because they define information on a channel which has an input and an output. The main motivation of this paper is to present a closed loop information measure which is compatible with constructivist thinking. **Design:** Our information measure for a closed loop system reflects how additional sensor inputs are utilised to establish additional sensor-motor loops during learning. Our information measure is based on the assumption that it is not optimal to stay reactive and that it is beneficial to become proactive through increased learning about the environment. Consequently our information measure gauges the utilisation of new sensor inputs to generate anticipatory actions. We call this information measure “predictive information” (PI). **Findings:** Our PI is zero if the organism uses only its reflex reactions. It grows when the organism is able to use other sensor inputs to preempt reflex reactions and is able to replace reflexes by anticipatory reactions. This has been demonstrated with a real robot that had to learn to avoid obstacles. **Conclusion:** PI is a new measure which is able to quantify anticipatory learning and, in contrast to the Shannon information, is calculated only at the inputs of an agent. This information measure has been successfully applied to a simple robot task but its application is neither limited to a certain task nor to a certain learning rule.

Keywords: Closed loop system, information measure, differential Hebbian learning, reactive vs proactive systems.

Introduction

Since its introduction by Shannon and Weaver (1949), a quantitative information measure has been used in many disciplines. For example, in telecommunications information is used to quantify the quality of a transmission. This definition of information has also been applied to neurons where the signal transmission from one neuron to the next one can be expressed in terms of information (Rieke, Warland, de Ruyter van Stevenick & Bialek 1997). Such an information measure treats the neuron as an input/output system which is appropriate in this case. However, when an organism acts in its environment its actions feed back to its sensors so that a closed loop perception-action system is established. Consequently the distinction between input and output becomes

fuzzy (Atmanspacher & Dalenoort 1994; Porr & Wörgötter 2005b) so that standard information measures which assess information transmission can no longer be used. In this article we will present an information measure for a closed loop system which measures how additional sensor inputs and therefore additional sensor-motor loops are acquired during learning.

Before we can define closed loop information it is necessary to introduce closed loop systems. Why do we need closed loop control? In the real world our knowledge about our environment is incomplete. We will always encounter situations that we cannot predict or in other words: the environment provides “surprises” which are called disturbances (Palm 2000). Such disturbances will cause a deviation from the desired state of the organism. For example, we might become hungry or an enemy may attack us unexpectedly. A

proper closed loop system is designed so that it can cope with these disturbances and thereby restore the desired state of the organism.

We will now define the point of view of the closed loop system. In contrast to approaches used in engineering, we are going to describe a closed loop system from its own perspective (Atmanspacher & Dalenoort 1994; Porr & Wörgötter 2005b). This internal perspective of the organism was first radically employed by (Foerster 1960). The crucial difference between the internal perspective and the outer perspective is how closed loop systems “observe” the environment and themselves. Forester claims that closed loop systems can only observe by using their own closed loops. Therefore, the only aspect which can be observed is the closed loop which establishes the feedback from the motor output to the sensor input (see Fig. 1). To make it clearer: an organism can only use its own senses to judge if an action has been successful or not. It cannot perceive the action *itself*. It can only perceive the *consequences* of its actions. This leads directly to the question of how an organism *evaluates* its actions. The answer is simple: only by its own sensor inputs. This leads directly to the statement that organisms can only control their inputs and not their outputs (Glaserfeld 1995). The contradiction between input- and output-control can be made clearer by an example, which we call the second chicken/egg problem (Porr & Wörgötter 2005b): Let us interpret the chicken as a closed loop system. The chicken wants to *keep* the egg and acts in a way designed to increase the sit-on-the-egg-time to improve the probability of successful hatching. The farmer, however, wants to *have* the egg. The farmer perceives the hen as an input-output system: food in and egg out.




Figure 1: The organism as a closed loop system. h transforms sensor events into motor outputs. p transforms motor outputs into sensor events. The only thing that can be perceived by the organism are actions fed back via (a). Anything that does not feed back cannot be perceived (b) and cannot be used to establish goals.

The hen, however, operates as a closed loop system to which the farmer is just a disturbance. As soon as the farmer removes the egg the chicken will produce a new egg thereby restoring its desired state. This example shows that external and internal perspectives are fundamentally different and illustrates Foerster's theory of organisms as closed loop systems, acting only according to their internal perspective.

In the preceding paragraph we have argued that an organism evaluates its own actions via its sensor inputs. These inputs can be used to improve future behaviour. Although many sensor events are associated with a real life situation, only a few inputs will probably be able to improve the behaviour of the organism. For example, with a hot surface, it is beneficial to react to the early heat radiation signal rather than a later pain signal elicited by touching the heat source. Many similar sequences of sensor events are encountered during the lifetime of an organism as the consequence of existing far-senses, e.g., vision, hearing, smell, and near-senses such as touch, taste, etc. Generally one observes that the trigger of a near-sense is preceded by that of a far sense, i.e., smell precedes taste, vision precedes touch, etc. Therefore far-senses are often predictive of corresponding near-senses (Verschure & Coolen 1991). Here, we will focus on the view that it is advantageous to react to the earliest of sensor events rather than to wait for later sensations.

We may now introduce our information measure in this closed loop system. In the paragraph above, we have argued that it is beneficial for organisms to generate anticipatory actions. These actions are generated by a

sensor input which predicts the imminent trigger of a late reaction. Consequently our information measure is the utilisation of new sensor inputs to generate an anticipatory reaction. We call this information measure "predictive information." If the organism only uses its late reflex reaction, the predictive information is equal to zero. When the organism is able to use other sensor inputs to preempt the late reflex reaction, thus enabling replacement of the reflex by an anticipatory reaction, this predictive information increases.

Furthermore, the anticipatory reaction is a learned behaviour, rather than being innately available to the organism. Our information measure therefore evaluates the learning process, for which we use the ISO-learning model (Porr & Wörgötter 2003) which is able to replace a reflex reaction by an anticipatory reaction.

Formalising closed loop control

Reactive control

Every closed loop control situation with negative feedback has a so called *desired state* and the goal of the control mechanism is to maintain (or reach) this state as quickly as possible (D'Azzo 1988). In our model we assume that the desired state of the reflex feedback loop is unchanging and defined by the properties of the reflex loop. We define it as $x_0 = 0$. First we will discuss this system in the absence of learning. Fig. 2a shows the situation of a non learning organism embedded into a very simple but generic (i.e. unspecified) formal environment which has a transfer function p_0 . This organism is able to react to an input only by means of a reflex.

Fig. 2b illustrates a possible set of signals which can occur in such a system. When a disturbance occurs, first the disturbance signal d deviates from zero, and then the input x_0 senses this change $x_0 \neq 0$ and only finally the




Figure 2: a Fixed reflex loop: the organism transfers a sensor event x_0 into a motor response v with the help of the transfer function h_0 . The environment turns the motor response v again into a sensor event x_0 with the help of the transfer function p_0 . In the environment there exists the disturbance d which adds its signal at \oplus to the reflex loop. b Possible temporal signal shapes occurring in the reflex loop when a disturbance $d \neq 0$ happens. The desired state is $x_0 := 0$. The disturbance d is filtered by p_0 and appears at x_0 and is then transferred into a compensation signal at v which eliminates the disturbance at \oplus .

motor output v can generate a reaction in order to restore the desired state $x_0 = 0$. No matter how fast the controller is, there will always be a reaction delay in such a system which means that there will always be an unwanted transient at x_0 . A system which operates purely reactive will always experience this transient. However, we will show in the next section that this transient can be used to learn an anticipatory reaction which then finally eliminates the transient at the input x_0 .

Predictive control

The reflex defined above may be prevented by an earlier sensor signal that informs the organism of an imminent disturbance. Fig. 3 shows how the disturbance d elicits a sequence of sensor events: first it enters the outer loop arriving at x_p filtered by the environment (p_1), while it arrives at x_0 only after a delay T . The trigger of the reflex can be avoided if the transfer function h_v generates a signal which compensates for the disturbance d . In an ideal case the inner loop (h_0 and p_0) will be completely eliminated so that the sensor input x_0 is always zero and the organism uses the input x_p instead. This means that the motor reaction v is now solely generated by x_p and no longer by x_0 . However, as the weight of input x_0 is fixed, input x_0 is still available: if anything in the environment changes and input x_p is no longer triggered input x_0 will always serve as a “back-up.”

Learning

The filter h_v in Fig. 3 can be implemented by ISO learning (Porr & Wörgötter 2003) which is able to learn temporal relations between input signals and uses these relations to turn reactive behaviour into proactive behaviour.

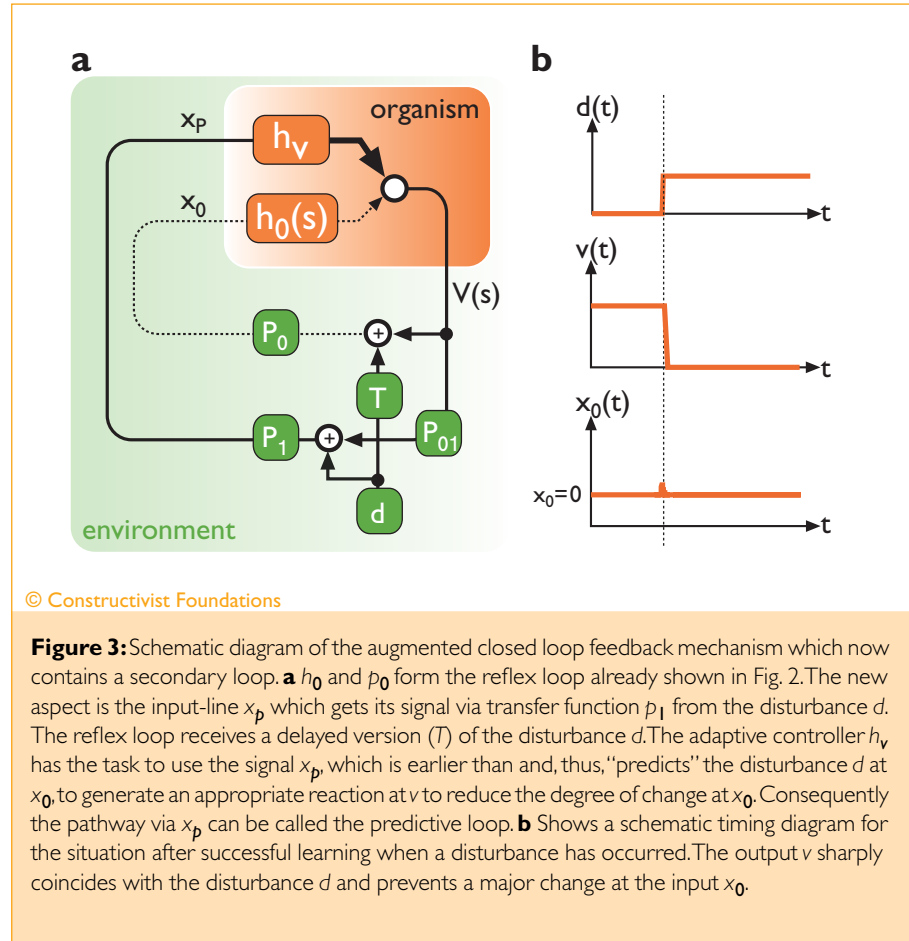
We consider a system of $N + 1$ linear filters h receiving inputs x and producing outputs u . The filters connect with corresponding weights ρ to one output unit v (Fig. 4a).

We will use x_0 to denote the reflex pathway with its corresponding weight $\rho_0 = \text{const}$. The output v is then given as:

$$v = \sum_{k=0}^N \rho_k u_k \quad u_k = x_k * h_k \quad (1)$$

where weight change is performed by the following learning rule:

$$\frac{d}{dt} \rho_j = \mu u_j v' \quad (2)$$



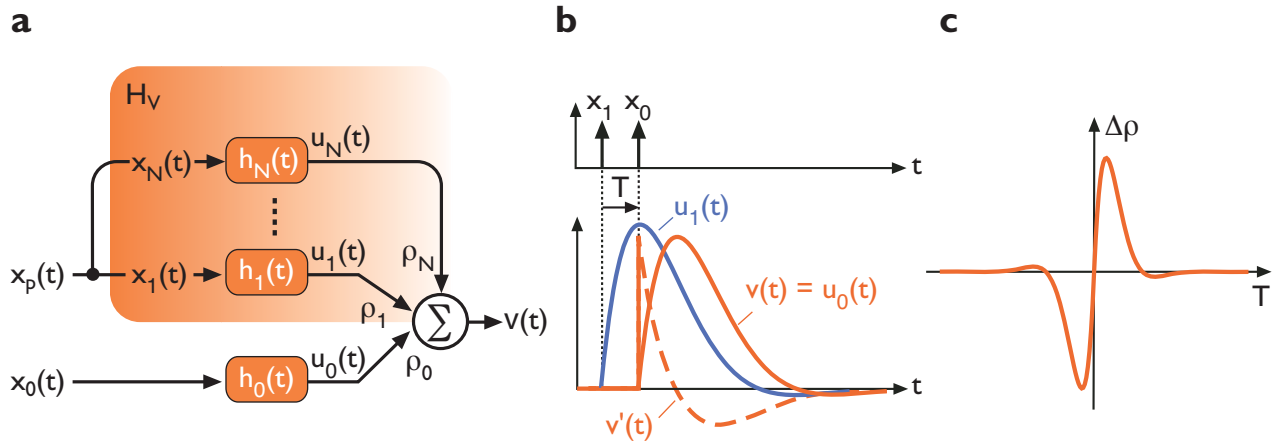
The constant μ is adjusted such that all weight changes occur on a long time scale (i.e. very slowly) compared to the decay of the responses u which ensures that the system operates in the steady state condition. The transfer function h is that of a lowpass or bandpass which transforms a δ -pulse input into a damped oscillation. The bandpass has two parameters: The frequency f and the quality Q which determines the damping where a low Q represents a high damping. In general we use as an initial condition for the weights $\rho_0 = 1$ and $\rho_j = 0, j > 0$.

ISO learning is illustrated in Fig. 4b: Derivatives of low pass filtered signals have a *phase lead* so that they precede the low pass filtered signal. In our case (namely ISO-learning) the derivative is taken from the output v of the neuron which is, before learning, identical to the low-pass filtered signal of the reflex u_0 . This derivative u'_0 is then correlated with the filtered predictive input u_1 . Because of the

phase lead of the derivative u'_0 , weights will grow when the signal at x_1 precedes the signal at x_0 . However, if the timing between x_1 and x_0 is reversed the weights will shrink. The corresponding weight change curve is plotted in Fig. 4c.

Closed loop information: Predictive information

At this point we can introduce an information measure which reflects the performance of predictive learning as outlined above. As the measure should reflect the success of predictive learning, it is reasonable to start with zero information before learning. In our model this means that we start with a late reflex via x_0 which is associated with zero information. Predictive actions via $x_k, k > 0$ are elicited



© Constructivist Foundations

Figure 4: **a** The neuronal circuit of ISO-learning. The shaded area marks the connections of the weights $\rho_k, k \geq 1$ onto the neuron which represent h_v with $x_k = x_p, k \geq 1$. **b** shows the inputs x_0, x_1 , the impulse responses u for a choice of two different resonators h and the derivative of the output v' . **c** shows the initial weight change $\rho_1(T)_{t=0}$ for $h_1 = h_0, Q = 1, f = 0.01$ after having stimulated the two filters h_0 and h_1 with delta pulses $x_1(t) = \delta(t)$ and $x_0(t) = \delta(t - T)$. This pulse-sequence was repeated every 2000 time steps. After 400,000 time steps the weight ρ was measured and plotted against the temporal difference T . The learning rate was set to $\mu = 0.001$.

when other sensor inputs are used in addition to the reflex input. Consequently, our information measurement should grow if these sensor inputs are used by ISO-learning. If there is more than one predictive sensor input all additional sensor inputs should contribute to the information measure.

We will now define our information value by the weights of ISO learning:

$$PI = - \sum_{j=0}^N \frac{|\rho_j|}{\sum_{k=0}^N |\rho_k|} \ln \left[\frac{|\rho_j|}{\sum_{k=0}^N |\rho_k|} \right] \quad (3)$$

where N is the number of inputs to ISO learning and ρ_j are the corresponding weights. We call this information "Predictive Information" (PI). Now, we have to test if our information measure behaves according to our theory. In the reflex-only situation, the weight ρ_0 is the only weight which is non zero. With $\rho_0 \neq 0, \rho_k = 0, k > 0$ Eq. 3 becomes zero which is the desired outcome. Note that the predictive information is always zero as long only one weight is non-zero.

To gain a better understanding of Eq. 3 we have to recall when the weights $\rho_k, k > 0$ start to grow. ISO learning implements differential Hebbian learning (Kosco 1986) which means that the weights grow when the correspond-

ing inputs $x_j, j > 0$ precede the signal v and, because of the fixed weight ρ_0 , precede the signal x_0 . Only those weights grow which are able to generate an earlier reaction in relation to the signal at x_0 (or v). Consequently the weights directly reflect the predictive power of their corresponding inputs. The more weights grow the more inputs are used for the anticipatory actions. Our measure PI reflects this: The larger the weights, the higher the value of PI .

For example, a naive person might only eat when she/he gets hungry. This represents basic reactive behaviour which utilises in our case just the pathway through x_0 . Because only the weight ρ_0 is non-zero the predictive information is zero. Eventually, the person learns to anticipate when he/she will get hungry and will eat earlier. In our formal model this predicative information enters the learning circuit through x_1 . Learning will increase the weight ρ_1 which is associated with the predictive input x_1 which leads to a non-zero predictive information because the two weights ρ_0 and ρ_1 are now non-zero. The more predictive signals x_2, \dots are available and used via nonzero weights ρ_2, \dots the more the predictive information will grow.

Weights stop growing when the reflex has been successfully avoided. In this case the

value of PI should also be at its maximum value. Let us assume for now that all inputs x_j are normalised to one. Having normalised inputs makes it easier to compare the loop gains in our feedback system. In case of pure reactive control only the weight ρ_0 is non-zero and therefore defines the gain of the feedback loop. During learning the weights $\rho_j, j > 0$ start to grow until the learning goal $x_0 = 0$ has been reached. It makes sense to assume that the feedback gain established by the weights $\rho_j, j > 0$ is similar to the gain of the feedback loop (ρ_0). Thus, it can be assumed that similar values for $\rho_j \approx \rho_0, j > 0$ for the predictive pathway will evolve as long as the different inputs are independent or form a sequence of events. Then every input $x_j, j > 0$ establishes a feedback loop on its own. In such a case the value of PI is at its maximum value because of an equal distribution of the weights. In our example (see Fig. 3) the input x_p feeds into a filterbank so that the corresponding weights are not independent. In such a scenario the sum of the weights $\sum_{j=0}^N \rho_j$ will probably generate a similar closed loop gain than the reflex gain ρ_0 . Consequently PI will be lower than for independent inputs $x_j, j > 0$. If one wants to compare the predictive information from two different sensors then Eq. 3 can be simply split up into the contributions from


one sensor and from another sensor. Summarising, the measure PI reflects the utilisation of additional inputs $x_j, j > 0$ to preempt the reflex at input x_0 . The more inputs are used the higher the PI but only under conditions where the closed loop gain of the predictive pathway(s) is comparable to the gain of the reflex pathway.

Robot experiment

The task of the robot experiment is collision avoidance. For a more detailed description of the experiment we refer the reader to Porr & Wörgötter (2003) and Porr, Ferber & Wörgötter (2003). The built-in reflex behaviour of the robot is a retraction reaction after the robot has collided with an obstacle. This represents a typical feedback mechanism; the desired state is that the signal at the collision sensor should remain zero. In order to avoid deviation from the desired state (i.e. collision) the robot is fitted with range finders which predict imminent collisions. ISO-learning is employed to learn the existing temporal correlation between the range-finder and the collision sensor signals. After learning the robot can generate a motor reaction in response to the range finder signals and thereby avoid collision and the retraction reflex.

The detailed circuit diagram of the robot is given in Fig. 5. The robot has two ISO learners, one for the steering angle and one for the speed of the robot. Here we focus on the ISO learner and its corresponding weights which control the steering angle. The predictive information can be calculated using Eq. 3. To be able to compare the predictive information associated with the two range-finders we calculated the predictive information for the left and for the right sensor separately.

We conducted two experiments. In the first experiment both range-finders were fully operational so that both of them were able to predict collisions. Comparing the traces of the range finders (RF, lower trace) with the signals from the collision sensors (CS, upper trace) it is apparent that there is a strong temporal correlation between both signals. Specifically, a peak at the collision sensor (CS) is usually preceded by a peak coming from the range finder (RF). In the presence of a strong temporal correlation the weights from the corresponding range-finders grow (see



© Constructivist Foundations

Figure 5: The robot has three collision sensors (CS), two range finders (RF) and two output neurons (speed ds and steering angle $d\phi$). The reflex-behaviour is triggered by the collision sensors (dotted lines). The corresponding weights of the reflex are adjusted in such a way that the robot performs an appropriate retraction reaction ($\rho_0^{ds} = 0.15$ and $\rho_0^{d\phi} = -0.5$). The two signals from the left and the right range finder are fed into two filter-banks with $N = 10$ resonators with frequencies of $f_k = 1 \text{ Hz}/k; k \geq 1$ and $Q = 1$ throughout. The 20 signals from the two filter banks converge on both the speed neuron (ds) and on the neuron responsible for the steering angle ($d\phi$).

Fig. 4c). Therefore, the robot successfully learns to avoid obstacles by using both range finder signals. The corresponding predictive information grows for both ranger finders.

In the second experiment the right range-finder was “blindfolded” by a small cap over its sensor. Analysis of the traces in Fig. 4d revealed that there was no longer a temporal correlation between the signal from the right range-finder (RF) and the signal from the collision sensor (CS) and that only the left range-finder was able to predict a collision (Fig. 4b). As expected the predictive information is still high for the left sensor but stays low for the right sensor.

In summary, the predictive information reflects the utilisation of the different inputs to develop anticipatory responses in relation to the reflex response. If a sensor signal is not

correlated to the reflex behaviour it does not generate behaviour and therefore the predictive information stays low. If a sensor signal is able to predict the reflex it is used by ISO-learning to generate an appropriate anticipatory reaction which results in higher predictive information.

Discussion

We have shown the successful application of predictive information to a real world robotic paradigm where the predictive information reflects the utilisation of different sensor signals to generate an anticipatory action.

Many other information measures have been defined (for a review in the context of constructivism see Porr 2002) but the infor-

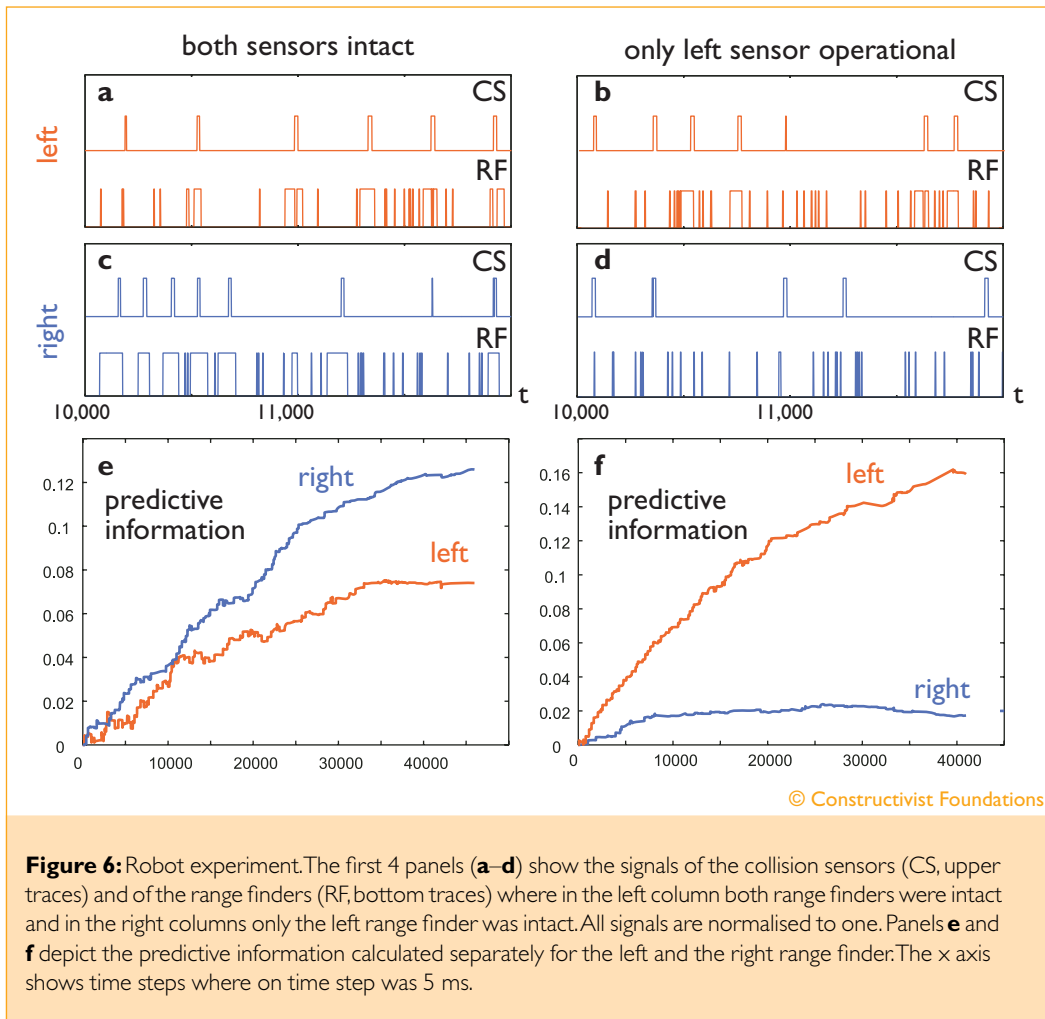


Figure 6: Robot experiment. The first 4 panels (a–d) show the signals of the collision sensors (CS, upper traces) and of the range finders (RF, bottom traces) where in the left column both range finders were intact and in the right column only the left range finder was intact. All signals are normalised to one. Panels e and f depict the predictive information calculated separately for the left and the right range finder. The x axis shows time steps where on time step was 5 ms.

We will now compare our predictive information with the Shannon information. Our information measure uses certain properties of ISO learning which can be also found in other learning rules like Hebbian learning (Hebb 1949), which has been extensively investigated regarding information processing. According to Linsker (1988) Hebbian learning implements information maximisation which means that a neuron transmits as much information as possible from its inputs to its output. This so-called “infomax principle” is equivalent to the detection of the first principle component in the input space (Oja 1982). In terms of weights this means that those weights grow whose inputs are highly correlated. As a result we preserve maximum variance at the output of the neuron which is equivalent to maximum information transmission through a channel. Linsker’s model is therefore directly concerned with information transmission as described by the work of Shannon and Weaver. Information in this context is only

mation measure of Shannon & Weaver (1949) is most pertinent to the present discussion. The total average information which is caused by M discrete events is given as:

$$I = - \sum_{k=1}^M p_k \log(p_k) \quad (4)$$

This is the average information generated by all M events which have the probabilities p_k . It can be proven that if all events have the probability $p_k = 1/M$ the information is at its maximum which means that all events have maximum uncertainty. The probabilities p_k can be interpreted in different ways, leading to different interpretations of the information measure: The probabilities p_k could be associated

with the amount of *surprise* the events have caused which is reflected by the logarithm in Eq. 4. On the other hand the probabilities could be also interpreted as the amount of bits which are needed to describe the probabilities. Even more interpretations arise when not only the information content but also the transmission of information is considered. Such an information measure based on the transmission of events operates on the conditional probabilities $p(y_k|x_k)$ of the channel. Information is zero in such a case when an event y_k has been triggered by any event x_k at the input of the channel. The information is highest if one even y_k has been caused exactly by one event at the input x_k . In such a scenario it is also possible to introduce redundancy by increasing the channel capacity or by adding an auxiliary channel.

limited by the signal to noise ratio and the cross correlations between different channels. In our model, however information is defined as a relevance measure. Information is measured against the reflex reaction and not against the whole input space. In our case the input space might be large (i.e. many sensors). However, if these signals cannot be used to preempt the feedback reaction the predictive information remains at zero. Therefore, in contrast to our measure the information defined by Shannon and Weaver might be very high. For example, the Shannon-information of an eye or a video-camera is probably very high. Also the optical nerve might transmit a large amount of information in form of Shannon information. However, the predictive information might still be zero if this information does not lead to learning and ultimately to behaviour.

There is a second important difference between Hebbian learning and ISO learning. While the weights in Hebbian learning undergo unlimited growth ISO learning stabilises the weights at the moment when the reflex has been successfully avoided. Our predictive information takes this into account. The predictive information is only high when the gain of the learned predictive loop is comparable to the gain of the reflex loop. For example, if one weight grows in an unlimited manner the predictive information goes to zero because the system has basically only one input left.

The third difference between Linsker's infomax principle and our predictive information is that our measure is designed for closed loop control whilst Linsker's measure relates to open loop scenarios. Our predictive information demands feedback in order to evaluate predictive actions. To our knowledge only Touchette & Lloyd (2004) and Klyubin, Polani & Nehaniv (2004) have presented recently closed loop information measures. However, in contrast to our approach they define closed loop information by information transmission from the sensors of the agent to its motor output. This definition makes it possible to compare open with closed loop control but makes the assumption that an organism can observe its output which contradicts the constructivist's view.

So far our predictive information uses properties of ISO learning, in particular it performs differential Hebbian learning and it stabilises as soon as the reflex has been avoided. Our ICO learning model may also be used in situations where nested predictive loops are not needed (Porr & Wörgötter 2005a, 2006). The application to reinforcement learning (Sutton 1988) should be also possible but a major difficulty arises from the fact that the feedback from the environment is more indirect because the critic first modifies the actor which then generates motor actions.

More generally it would be interesting to develop a measure which is completely independent of the learning rule. Such measure must be calculated by just looking at the different input signals. This would make our predictive information applicable to other closed loop systems such as fuzzy control systems or systems which employ symbolic control such as classical AI systems. The work by

Klyubin, Polani & Nehaniv (2004) promises interesting opportunities in this sense, as their information measure does not depend on a specific learning rule.

ABOUT THE AUTHORS

Bernd Porr has a degree in Physics and Communication Science / Journalism (both from the University of Bochum). In 2000 he moved to Scotland and did his PhD in sequence learning and predictive control at the University of Stirling. In 2004 he took up a post as lecturer at the University of Glasgow at the department of electronics and electrical engineering. His main research interests are in the fields of biologically inspired adaptive control, synaptic plasticity, image processing and radical constructivism. *Alice Egerton* has a degree in Pharmacology (University of Glasgow) and a PhD in Psychopharmacology (University of Strathclyde). Her doctoral work has focussed on preclinical models of schizophrenia, with particular reference to behavioural pharmacology and underlying neurobiology at the Yoshitomi Research Institute of Neuroscience in Glasgow (YRING). She now works in Psychiatry Research at Imperial College London, where her primary research interest is the investigation of brain dopamine systems using positron emission tomography (PET) in relation to schizophrenia and drug dependence.

Florentin Wörgötter has studied Biology and Mathematics in Düsseldorf. He received his PhD in 1988 in Essen working experimentally on the visual cortex before he turned to computational issues at the Caltech, USA (1988–1990). After 1990 he was researcher at the University of Bochum concerned with experimental and computational neuroscience of the visual system. Between 2000 and 2005 he had been Professor for Computational Neuroscience at the Psychology Department of the University of Stirling, Scotland where his interests strongly turned towards "Learning in Neurons." Since July 2005 he leads the Department for Computational Neuroscience at the Bernstein Center at the University of Göttingen. His main research interest is information processing in closed-loop perception-action systems, which includes aspects of sensory processing, motor control and learning/plasticity.

Acknowledgements

We thank Vi Romanes and the referees for their helpful feedback.

References

- Atmanspacher, H. & Dalenoort, G. (1994) Introduction. In: Atmanspacher, H. & Dalenoort, G. (eds.) *Inside versus outside*. Springer: Berlin, pp. 1–12.
- D'Azzo, J. J. (1988) *Linear control system analysis and design*. McGraw: New York.
- Foerster, H. von (1960) On self-organizing systems and their environments. In: Yovits, M. & Cameron, S. (eds.) *Self-organizing systems*. Pergamon Press: London, pp. 31–50.
- Glaserfeld, E. von (1995) *Radical constructivism: A way of knowing and learning*. Falmer Press: London.
- Hebb, D. O. (1949) *The organization of behavior: A neuropsychological study*. Wiley-Interscience: New York.
- Klyubin, A. S., Polani, D. & 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*. IEEE Computer Society: Seattle, Washington, pp. 177–180.
- Kosco, B. (1986) Differential hebbian learning. In: Denker, J. S. (ed.) *Neural Networks for computing*. Volume 151 of AIP conference proceedings. American Institute of Physics: New York, pp. 277–282.
- Linsker, R. (1988) Self-organisation in a perceptual network. *Computer* 21(3): 105–117.
- Oja, E. (1982) A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology* 15(3): 267–273.
- Palm, W. J. (2000) *Modeling, analysis and control of dynamic systems*. Wiley: New York.
- Porr, B. (2002) *Systemtheorie und Naturwissenschaft Eine interdisziplinäre Analyse von Niklas Luhmanns Werk*. Deutscher Universitäts-Verlag: Wiesbaden.
- Porr, B., von Ferber, C. & Wörgötter, F. (2003) Iso-learning approximates a solution to the inverse-controller problem in an unsupervised behavioural paradigm. *Computation* 15: 865–884.

- Porr, B. & Wörgötter, F. (2003) Isotropic sequence order learning. *Neural Computation* 15: 831–864.
- Porr, B. & Wörgötter, F. (2005a) Non-hebbian synaptic plasticity allows for the stable implementation of one-shot predictive learning. In: Zimmermann, H. & Kriegelstein, K. (eds.) *Proceedings of the 6th Meeting of the German Neuroscience Society, 30th Göttingen Neurobiology Conference 2005*, p. 1223.
- Porr, B. & Wörgötter, F. (2005b) What means embodiment for radical constructivists? *Kybernetes* 34(1/2): 105–117.
- Porr, B. & Wörgötter, F. (2006) Strongly improved stability and faster convergence of temporal sequence learning by utilising input correlations only. *Neural Computation* 18 (in press).
- Rieke, F., Warland, D., de Ruyter van Stevenick, R. & Bialek, W. (1997) *Spikes: Exploring the neural code*. MIT Press: Cambridge MA.
- Shannon, C. E. & Weaver, W. (1949) *The mathematical theory of communication*. University of Illinois Press: Urbana.
- Sutton, R. (1988) Learning to predict by method of temporal differences. *Machine Learning* 3(1): 9–44.
- Touchette, H. & Lloyd, S. (2004) Information-theoretic approach to the study of control systems. *Physica A* 331: 140–172.
- Verschure, P. & Coolen, A. (1991) Adaptive fields: Distributed representations of classically conditioned associations. *Network* 2: 189–206.

Received: 17 November 2005

Accepted: 07 February 2006