From predictive coding to percept – a computational approach

Felix Creutzig

Consciousness remains a difficult problem for philosophers and neuroscientist alike and, even more, there is no agreement on the definition of the problem. It is most difficult to imagine why the physical world generates qualia, the subjectiveness of personal experience. Theoretical neuroscientists argue that answering more tractable problems first may also help to approach such a hard question. Relatively easy problems relate to the substance of consciousness: What are the neural correlates of consciousness and where are they located? What could be the function of consciousness? Analyzing these problems, it is reasonable to investigate the neuronal basis of brain function and see how far we can get. Hence, the hope is that we unravel hidden concepts that give us the language to tackle even the question, why phenomenological states feel like anything at all. Indeed, substantial progress has been made on these questions in the preceding decades (Koch, 2004).

Here, we take a basic approach and try to comment on the relation between substance and function of consciousness. From an evolutionary point of view, the nervous system has evolved to relate environmental stimuli to the organism's action. Or, in other words, it has evolved to give an interpration of sensory input such that this interpretation can guide motor output. There is no reason to believe that consciousness is not part of this framework. But what should be the precise function of consciousness? Could a zombie do the same job?

In this text, we try to transfer the concept of information into a wider functional view of consciousness and emphasize the dimension of time. First, we review some evidence that sensory processing may follow the guiding principle of filtering temporal regularities. Second, we rephrase experimental results showing that neural correlates of consciousness are localized in cortical brain regions. Third, we try to distinguish between unconscious and conscious sensory-motor interactions. Forth, we integrate some speculations on the evolutionary function of having perceptions and take a viewpoint that regards perceptions as a bottleneck between sensory processing and complex motor response functions, emphasizing the role of time.

Function of sensory processing

How do neurons in the periphery of the sensory system process and transmit signals from the organism's natural environment? Photoreceptors encode visual input pixel by pixel, auditory neurons encode sounds frequency by frequency in the cochlea and the receptors of the olfactory system are specialized on specific molecules each. In subsequent stages of the neural pathway, this input is combined in particular ways. This information processing is subject to intensive research as problems are relatively well defined, accessible and can be studied in simple animal systems such as insects and rats.

Experimental results and theoretical studies have pushed forward the hypothesis that inputs are processed according to their probabilistic structure (Attneave, 1954; Barlow, 1961). More precisely, the hypothesis states that environmental stimuli are efficiently coded by reducing redundancy. Reduncancy reduction may be used for data compression. However, such an efficiently compressed code may be difficult to decode (e.g., Barlow 2001), a point that has to be clarified by further research. But, in general, finding redundancies also means extracting statistical regularities of external stimuli. In fact, all interesting things that we experience consist of statistical regularities - in contrast to the uninteresting flickering TV screen. We can use these statistical regularities to fill in missing pixels in images automatically.

Based on these ideas of finding redundancies, statistical methods and algorithms such as principal and independent component analysis have been developped. Applying these methods on sets of natural images, one finds “independent components” similar to receptive fields in the visual part of the neocortex.

The living organism is not contrasted with a set of images but with a flow of temporally correlated sensory input. Animals also act in time: Making use of current and previous input, animals try to figure out an appropriate response. This relationship between input past and output future can be observed on all timescales. Flies have to identify and evade obstacles within milliseconds, whereas human career planning deals with decades. Hence, extracting temporal statistical regularities is an important task of the nervous system.

Some evidence indicates that the visual systems exploits temporal correlations. In the retina, predictive coding is used for enhanced surround inhibition (Srinivasan, 1982). In the visual cortex, feedback connections predict incoming events and can explain extra-classical receptive-field properties (Rao, 1999). From the computational viewpoint, an information-theoretic ansatz of predictive coding demonstrates that another established algorithm, slow feature analysis, can be understood in terms of extracting temporal regularities (Creutzig, 2007). Slow feature analysis extracts the slow components of signals varying in time. Slow feature analysis itself cannot only explain properties of visual receptive fields (Berkes, 2003) but also higher-order phenomenons, such as so-called place cells in the hippocampus which encode the location of the animal but not the visual input (Franzius, 2006).

Hence, filtering the past for information about the future is a suitable principle for the sensory system, consistent with experimental findings and successful algorithms.

Percepts are localized

A large part of the brain capacities is reserved for visual processing. Hence, it seems natural to investigate the conscious percepts of visual experience. The flow of visual information goes from the retina to the thalamus, then to the occipital cortex. In the cortex, information is processed first in V1 before being forwarded to more specialized regions, e.g., the region MT for motion detection. Cells in V1 encode bars directed into specific orientations. Experiments show that cells can be active in response to presented stimuli even if this stimulus is not part of a conscious percept (He, 1996). On the other hand, the activity of neurons in MT is an explicit representation of a conscious sensation. In one setup monkeys had to track moving dots on a computer screen under noisy conditions and to report the perceived direction of movement. The response of single neurons corresponded to the monkey's perception but not to the actual movement of the dots (Britten, 1992). This is an example of how neural activity correlates to conciousness. In binocular rivalry experiments, two different pictures are given to right and left eye. You perceive just one of them. But after some seconds, this pictures begins to fade and the other one appears as a percept. Unlike to cells in the visual areas V1, V4 and MT, activity in single cells in the inferior temporal cortex is correlated with the perceived picture (Logothetis, 1998). We cannot conclude whether this neural activity covaries with or constitutes the percept. However, we can state that perceptions are localized in specific cortical areas, probably in populations of neurons.

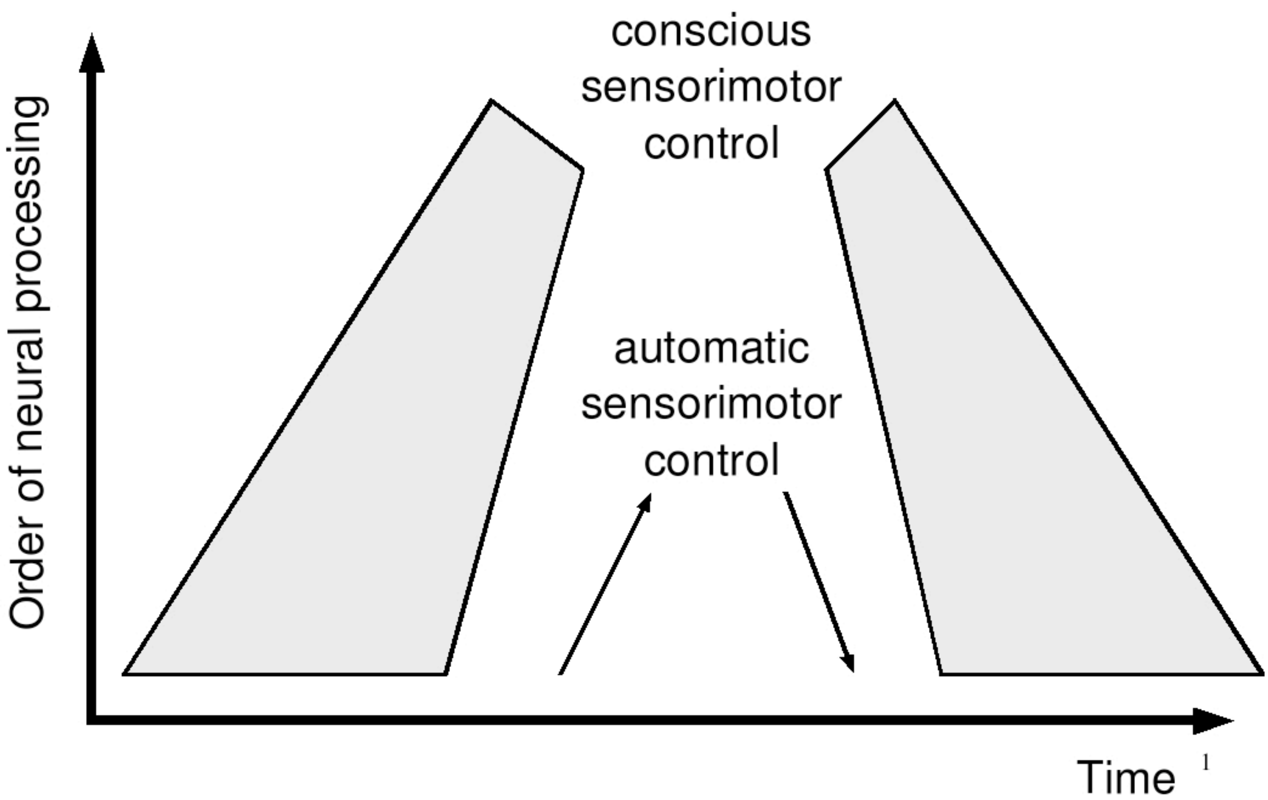

Figure 1: Consciousness is crucial for sensorimotor coordination. Many reactions happen unconsciously and in an automatic fashion. However, in novel and complex

situations, we have to integrate over past experience for future action, usually on long time scales, indicated by the grey-shadowed area. For such situations, we have to rely on consciousness.

Unconscious action

Many of our actions happen unconsciously and in an automatic fashion. After learning, we can cycle through the city and even stop in front of red traffic lights in a completely automatic but very precise way. Though some of these actions need cortical processing of visual input, they not necessarily involve consciousness. Even more, the knowledge of the world hold by the motor system is more precise than the estimations of our consciousness. For example, people were asked to estimate the steepness of hills by visual judgement or by adjusting the tilt of their flat, outstretched hand (Profitt, 1995). Motor output was very accurate but judgement by perception was systematically biased upward.

Milner and Goodale argue that two different cortical parthways exist: (A) the conscious vision-for-perception pathway and (B) the unconscious vision-for-action pathway (Milner, 1995). The latter – as well as other non-cortical circuits - is responsible for immediate and stereotyped actions in response to sensory input, needing only a short processing time. For action, exact position and spatial orientation is crucial, whereas for perception, object identification is relevant. In computational terms, the two pathways need to calculate different invariances. This computational need may explain the split of the two pathways.

In fact, the direct input-output relations may be interpreted in the framework of predictive information: temporal regularities are represented and converted into effective action.

What is the function?

Consciousness seems mainly to be a property of more complex animals. Thus, biologists infer that consciousness must have an evolutionary function. In the following, we will try to integrate some ideas on evolutionary function to a speculative framework. Evolutionary function is never on the passive side but must always be intimately related to the behavior of an organism. Hence, perception cannot be sufficient for itself but must have consequences for action.

In what kind of situation do we need to be aware? Consciousness seems to be required for complex and novel situations where we cannot rely on stereotyped response patterns. In fact, whenever we need short-time memory not to act now but in the near future, we need conscious processes, probably correlated to activity in the frontal cortex (Fuster, 1973). We also seem to need consciousness for future planning and symbolic operations such as thinking. Crucially, thinking and planning itself may not be conscious processes. Indeed, it is proposed that consciousness occures at the interface between sensory processing and planning. Only the observations of our thinking may then be presented as perceptions. This theory is called the intermediate-theory of consciousness (Jackendorf, 1987).

Combining these different ideas, we are now in a position to draw a coherent picture. Perceptions can be seen as a communication platform that summarizes the relevant information about the world into one current neuronal state offering this information to higher planning states. Also, by broadcasting this information over many cortical regions, episodic and semantic memories may be recruited. Planning states and newly recruited memories may then combine with the previous percept to form a new one. Hence, the percept is not a determinant of the input only but also of previous experiences, socialization and additional symbolic operations.

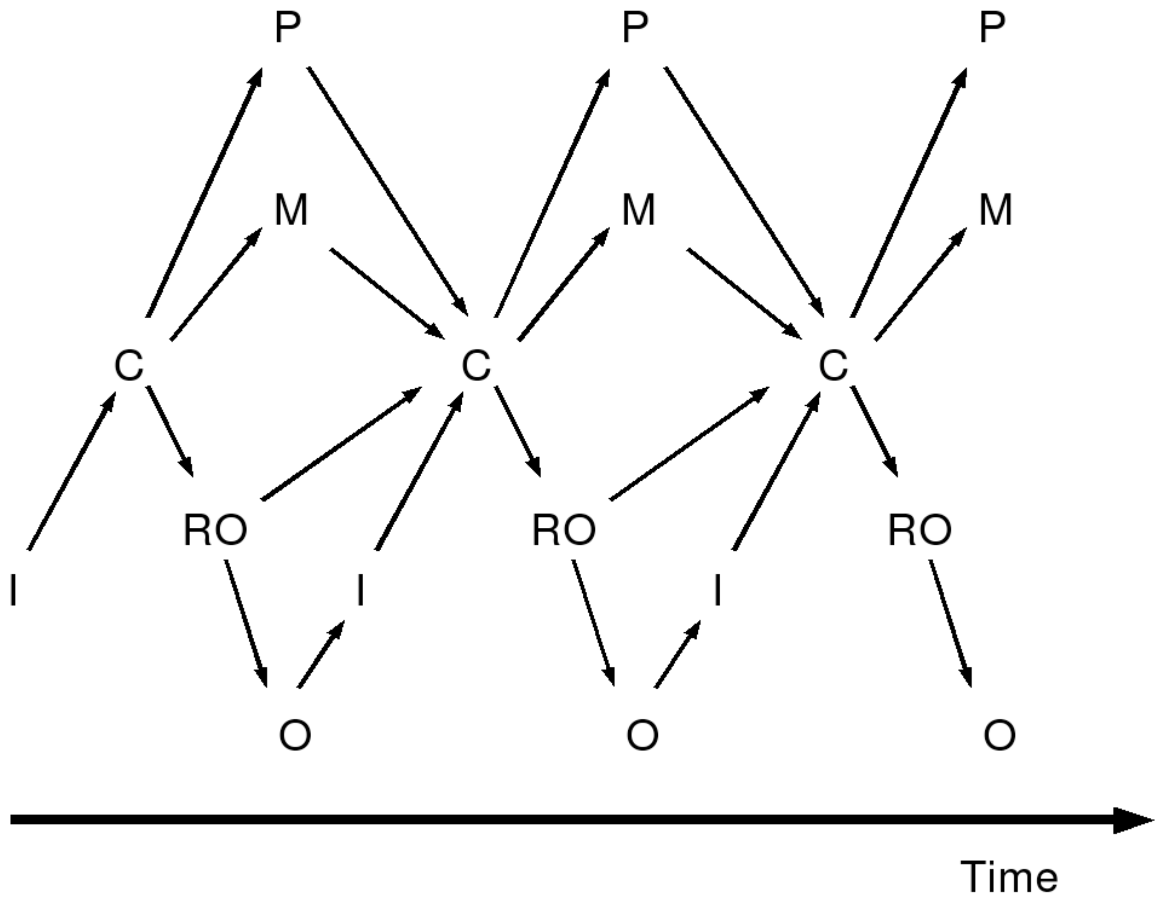

Figure 2: A conscious state (C) may be incited by sensory input (I). A conscious state can recrute memory (M) to use past experience for evaluation of the situation. C can also activate planning stages (P). Both M and P can lead to a new conscious state. A representation of motor activity (RO) leads to motor output (O). This output (O) influences environmental statistics and subsequent input (I). The subsequent conscious state (C) is influenced by I but also by the RO providing a context for I.

Furthermore, percepts are not independent from action. Sensations are interpreted according to previous action. That is, we use information about our body movements that constitute a context for incoming data. All of this branches are pictured together in Figure 2.

We would like to emphasize two points here. First, this interaction between perception and action, relying on higher-order processes, always happens in time and cannot be thought without the notion of (supposed) temporal correlations. Particularly, consciousness is only needed on time-scales exceeding direct input-output relations. Only in situations where we have time for reflecting things, consciousness can have value. Indeed, being aware may be effective for action only indirectly, e.g., by storing new, potentially valuable, memory. With longer effective time-scales both in past and in future, relatively complex information can be integrated. Ressources are allocated to flexibly handle novel situations that cannot be (or have not yet been) automatized by subconscious processes.

Second, only a small part of all possible perceptions appear. A specific part of our environment is represented in a percept and this percept is enriched with past experiences giving a context to the situation. In other words, the cortex is saved from many other possible situations that are, by judgement of filtering processes such as attention, irrelevant. Hence, the percept can be seen as an information bottleneck between the sensory past and future planning, being relevant in complex situations. Stable perceptions are a final product of filtering processes getting rid of irrelevant information and keeping only those that may be important for further action and cannot be processed by established non-conscious circuits. These perceptions are then presented, for example, to the frontal lobe which is responsible for planning.

This approach points to the limited information processing capacities of neural systems and the need to efficiently handle coding and energy resources. Furthermore, we suggest to see conciousness as direct elaboration of simpler neural sensory-motor circuits, enabling the computation, learning and handling of relevant temporal regularities. The unconscious processing makes use of learned temporal regularities, whereas the percepts are needed to handle non-standard situations. The behavioural output is not evident but has to be inferred recruiting memory and higher cortical planning states. But we don't know in advance which memory states are of importance. Hence, it is a sensible strategy to broadcast information to many cortical regions to see which memories and planning operations give feedback to successive consciousness stages, modifications of previous states. The process of broadcasting corresponds to having a conscious state, the identity of the source, i.e., the originating neural bundle, corresponds to the specificity of state.

We can put these aspects into one phrase: the conciousness constitutes a past-future information bottleneck of input-output relations. The function of consciousness can then be seen as interpreting the world on long time scales such that action can be *effective* in the evolutionary sense. As a zombie, i.e., without consciousness, we could not develop complex plans regarding, for example, our career. Finally, we should clearly state, that major parts of the previous paragraphs don't exceed speculations.

Rather than a scientific theory, this essay offers an additional viewpoint that may help to understand this scientific and philosophical problem.

References

Attneave, F., Some informational aspects of visual perception. Psychological Review, pp.183--193, 1954.

Barlow, H.B. Possible principles underlying the transformation of sensory messages, Sensory Communication, pp. 217--234, 1961.

Barlow, H.B. Redundancy reduction revisited. Network: Computation in Neural Systems. pp. 241—253, 2001.

Berkes, P. and Wiskott, L. Slow feature analysis yields a rich repertoire of complex cells. Journal of Vision, pp. 579–-602, 2005.

Britten, K.H. and Shadlen, M.N. and Newsome, W.T. and Movshon, J.A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. Journal of Neuroscience, pp. 4745--4765, 1992.

Creutzig, F. and Sprekeler, H. Predictive coding and the slowness principle: an information-theoretic approach. Neural Computation (20) 1026-1041, 2008.

Franzius, M and Sprekeler, H. and Wiskott, L. Slowness leads to place cells. In: Proceedings CNS 2006, 2006.

Fuster, J.M. Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. Journal of Neurophysiology, pp. 61-78, 1973.

He, S., Cavanagh, P., and Intrilligator, J. Attentional Resolution and the locus of visual awareness, Nature pp. 334-337, 1996.

Jackendoff, R. Consciousness and the Computational Mind. Cambridge, MA: MIT Press, 1987.

Koch, C., The Quest for Consciousness: A Neurobiological Approach. Englewood, Colorado, 2004.

Logothetis, N.K. Single units and conscious vision. Phil. Trans. R. Soc. Lond. B, pp. 1003—1037, 1998

Milner, A.D. and Goodale, M.A. The visual brain in action. Oxford, 1995.

Profitt, D.R. and Bhalla, M. and Gossweiler, R. and Midgett,J. Perceiving geographical slant. Psychonomic Bulletin & Review, pp. 409—428.

Rao, R.P.N. and Ballard, D.H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nature Neuroscience, pp. 79--97, 1999.

Srinivasan, M.V. and Laughlin, S.B. and Dubs, A. Predictive coding: a fresh view of inhibition in the retina. Proc. R. Soc. Lond. B Biol. Sci., pp. 427–459, 1982.