Biol Cybern (2009) 100:5–10
DOI 10.1007/s00422-008-0286-1

INTERVIEW

Biological Cybernetics

Perceptions, reflections, and new directions in *Biological Cybernetics*: Horace Barlow in conversation with Leo van Hemmen and John Rinzel

Interview with Horace Barlow[1]

Barlow: I have always wondered why cybernetics went out of fashion and then came back in again, and I think it was because there was originally a great emphasis on feedback systems, but physiologists were already very well familiar with the idea of feedback and homoeostasis. So, after initially arousing interest, people said "Well, all the applicable parts of this we already know about". And it's true: neurophysiologists didn't often have an exact quantitative understanding of feedback, but they did have the qualitative ideas.
I don't think the exact treatment has added very much in the case of homoeostasis, but although Norbert Wiener gave them less emphasis, Cybernetics also introduced many of us to signal–noise ratios and other ideas connected with information theory, decision-making and statistics. I think it's these parts that have turned out to be novel and important in vision, and also for understanding many aspects of higher brain function. For example most of the things you can do with vision, you can also do with computer vision, but it's a different matter to rival the signal–noise ratios you achieve with natural vision. So the mystery remains to some extent. It's not that you can't make computers see, but there are many visual tasks we can still do better than computers can.

van Hemmen: Yes, but say, if I may go back to Wiener, with whom it apparently all started through, in particular, his book and his essay in Scientific American of the same year (1948). Can you try to make a summary? And what I found quite astounding is that he neatly explains cybernetics, the steersmen, and extensively discusses neurons, but it struck me Wiener never said anything on learning. Can you understand why?

Barlow: I think it was because machine learning was not well developed then. He didn't see cybernetics as having useful things to say to psychologists studying learning. Also at that time there were already several mathematical theories of learning going the rounds, so the fact that cybernetics was quantitative and precise did not seem to add anything new. Those old psychological theories have all turned out to be incomplete, to put it mildly, and I don't think anybody pays any attention to them now. So the answer to your question is that there were rival quantitative treatments of the subject, together with the fact that the existing state of the computational art of learning was not very well-developed. Machine learning now is in a quite different state through incorporating ideas from decision theory, statistical inference, and information theory, but it is interesting to note that Wiener did not foresee this development.

van Hemmen: That's certainly true, but you have to program a computer. Also in the old days that was a huge amount of work.

Barlow: I haven't really looked, with the computer-explosion in mind, at what Wiener wrote. But there's another aspect, namely the enormous importance of information theory in molecular biology, which was not connected with Wiener at all. But as people came to realise the basic importance

[1] This is the transcript of a telephone conversation between Horace Barlow, Leo van Hemmen, and John Rinzel on November 4, 2008. Professor Horace B. Barlow (born December 8, 1921) is at Trinity College, University of Cambridge. He is a visual neuroscientist, who started his career with precursor experiments to Hubel & Wiesel on frog vision. He stressed the importance of information-theoretic concepts, such as statistics of natural scenes, redundancy reduction, and factorial codes. Being a son of Sir Alan Barlow and Lady Nora Darwin he is the great-grandson of Charles Darwin. Professor Barlow is a Fellow of the Royal Society. He belonged to the Editorial Board of *Biological Cybernetics* from the very beginning, viz., Volume 1, as one can verify on p. 3 of the present issue.

of information in all biological control systems, they took another look at its importance in neuroscience, and I think this contributed to the delayed second take-off of Cybernetics.

van Hemmen: Yes, we'll come back to that later on. Could you say, Horace, because you were one of several people who joined the Editorial Board with volume one of *Biological Cybernetics*, there was already the distinction of cybernetics and biological cybernetics? Moreover, what in those days did you find attractive when joining a journal like *Biological Cybernetics*? How did it appeal to you?

Barlow: Two things: One was that it was a quantitative and theoretical approach, which were few and far between in biology in those days. And the other thing was its connection with information theory, which defined a different specific quantity to follow through the brain, in addition to energy, but of comparable importance. The first successful onslaught on vitalism was to show that the energy transactions in living things followed the ordinary physical rules. The retreating vitalists naturally said "Oh yes, but our ability to make rational judgements is something that no machine could do". So here again was a possibility of some non-physical things going on in the brain, but of course it has turned out that, as far as we can tell, there's no mysterious aspect to this, it's all quite straight-forward physics, but the important quantity to follow is information rather than energy.

van Hemmen: I think we could not but agree were it not for the fact that of course the brain takes quite a bit of energy, so it does cost energy.

Barlow: Yes, it does cost energy. But, nevertheless, to understand the unique work the brain does, you need to track information, rather than energy.

van Hemmen: What you are hinting at is very interesting. I would pose myself the thesis that information theory complements our understanding of the brain but is not the real access to understanding it, because information theory does not explain anything. It only describes how much information can go through the brain. Sorry, I'm rephrasing this as we check the arguments, but more or less that's what it is. Information theory was made for estimating the amount of information that something can handle. But how it handles the information, that's what information theory doesn't tell us.

Barlow: Showing how the brain obeys the law of conservation of energy does not tell you how it works either, but finding where information is conserved and where it is lost requires concepts like redundancy and its exploitation, error-correcting codes, sparse coding, adapting to properties of natural images, and so forth. I think these ideas are crucial for understanding the brain. Perhaps their importance is still under-appreciated.

Rinzel: Maybe what Leo was saying is something that I stress when I try to define computational neuroscience to students: it is that there are two pathways here and they overlap a lot. One is trying to figure out what computation a neural system is doing, and that involves much more information theory kinds of concepts and the other is, how is it implemented in neuralware, how exactly does a nervous system carry out such a computation? Does that resonate at all?

Barlow: Yes, I think what you are saying connects with David Marr when he wrote about the computational theory of neural computation. What is the brain trying to achieve by lateral inhibition, or colour opponency? You need the concepts of information theory to understand this. But it's also important to realise that the same computational theory can be implemented by many different methods. You need experiments to decide how goals are pursued and how a computation is implemented, but it can help the experimenter to know what the goals are! I don't think one can design an experiment properly without understanding that. Different methods are usually good at slightly different tasks, so when you find out what tasks the brain does well, you are already getting some guidance about the nature of the computation that's being done.

van Hemmen: Before rounding off this very neat piece of our discussion, when you were discussing error-correcting codes, the first name that came to my mind was Hamming. And why did Hamming think about the problem? As far as I can remember he was also at Bell Labs and information theory made clear to him that in order to produce an error-correcting code you have to provide more information than the message alone, and as you presumably remember the Hamming code was pretty smart. It was the first error-correcting code, but Hamming realized he must give the code a little bit more than just the message itself. And that's the key thing. But how to devise an error-correcting code, in particular a smart one, that's still a different story.

Barlow: Yes, you need some redundancy (in the information theory sense) for error-correcting codes to be possible, but aren't we all a bit surprised at how far a little redundancy goes? This is another example of the main point, that concepts of information theory are important in understanding what the brain does, in suggesting the goals of the computations it is doing, and for understanding the natural difficulties that impede successful computations for achieving them.

van Hemmen: Um, yes, certainly true. But then the hot question is, how does the brain implement, realize error-correcting codes?

Barlow: Yes, indeed. But it's a first step to know in informational terms what it's doing.

van Hemmen: Can I ask how you see the role of Reichardt in the early development of *Biological Cybernetics*? You met him years and years before I did.

Barlow: Well, I don't know what actual transactions took place in setting up the journal, but Werner played an enormous role in making the subject important and interesting in other people's eyes. And his work is still important: for example, that paper he gave at the 1959 MIT conference, on auto-correlation as an important function the brain must be able to perform, even the beetle's brain, was an absolute landmark. After all his model is the basis of almost all current models of motion perception. Actually I don't think the later models have been as good as his, for they seem to avoid the "auto" part of the correlation mechanism, and veer off towards the filtering aspect. They are based on following a filtered feature, rather than the auto-correlation aspect, which seems to me must form the logical basis of movement detection: as I think he may have appreciated, it's not just spatio-temporal filtering, but a form of symmetry detection. Having a Max Planck Institute to back his work was of course an enormous help, but the human hero was Werner Reichardt.

van Hemmen: That's certainly true. I visited him towards the end of his life, and I can only admit he had a really fantastic Max Planck Institute.

Barlow: Oh, yes. There were many other people there: Tommy Poggio was there for a time, wasn't he? And also Valentino Braitenberg as another director, so I think altogether it was very effective. Of course some people, usually ones without Max Planck Institutes of their own, used to tease Werner for building a whole institute around a fly, but he made good use of the support they provided; it was great.

van Hemmen: I fully agree. I also gratefully remember my many discussions with Valentino Braitenberg in Tübingen. That was indeed not only rewarding but also very stimulating. How do you think Reichardt viewed the role of mathematics in computational neurosciences or, as I prefer to call it, theoretical neurobiology? Did he think about mathematics explicitly, or just take it for granted?

Barlow: I think mathematics was the basis for a great deal of his thinking. He took the mathematics of an actual mechanism, such as his navigational device for estimating the ground speed of airplanes, and applied this to biological mechanisms. At heart I think he was more a mathematician than an engineer or biologist.

van Hemmen: Okay. That's really nice. Making a jump, because I think we both agree that cybernetics more or less kept silent for twenty years, and it reappeared at the end of the eighties.

Barlow: Of course another factor in it reappearing was that people, by the eighties, had their own desktop computers, so that they could do for themselves some of the things which had previously required hordes of slaves typing up IBM cards and the like, as was mentioned earlier.

van Hemmen: That's neat. Yeah, that's what I didn't think of.

Barlow: And also by that time people had gotten used to the fact that computers weren't just as good as human brains for doing some things, they were a hell of a lot better than humans at many of them.

van Hemmen: Yes, certainly true. But as I noted also elsewhere, I also think people rediscovered learning. Wiener had forgotten learning, and for quite a while there was quite a bit of silence regarding learning, and then by the end of the eighties, the people saw again, hey, we have a learning problem. How do we learn for instance spatiotemporal patterns? That is, neuronal patterns evolving not only in space but also in time. How do we do this? At least that was the way in which my former graduate students Andreas Herz and Wulf Gerstner and I myself entered the field. It all culminated in spike-timing-dependent plasticity, which, by the way, we predicted a year before it was confirmed experimentally, by studying the barn owl and the way in which it localizes its prey with such a high precision, microseconds instead of milliseconds, and asking: How can we explain this? Learning spatiotemporal patterns was greatly neglected. And you know there were several people, starting with Terry Sejnowski and many others in the US, who also got fascinated by learning. Do you think that, say, learning joining the neurobiological scene as a theoretical problem was an important input?

Barlow: Yes, specifically with regard to learning, as soon as people started thinking of little devices that did brain-like operations, they got into learning. The earliest robots, even in Grey Walter's day, tended to have some true kinds of learning mechanisms, but what was not appreciated to begin with was that what's important in memory and learning is not just the hardware that stores the information; it's largely the way the information is organised that determines what you can learn and what you cannot. Learning is essentially a statistical decision-making problem, but that wasn't appreciated until the second phase of cybernetics, because people thought the big problem of memory was to know where it occurred: was it the synapse, was it perhaps the accumulation of sodium inside the cell body, were reverberating circuits involved, and so on?

van Hemmen: As a final question regarding the past, if we now look back, you can imagine that, as Editor and Coeditors-in-Chief we have had extensive discussion as to whether we would stick with *Biological Cybernetics*. Well it's such a nice historical name for research having the Wiener focus on feedback together with a bit of what we would now call biomimetics—which attempts to realize and mimic neuronal functions in hardware. Wouldn't you think that, say, sticking to *Biological Cybernetics* might well be justified? Or do you have a good proposal for another name, if we had to?

Barlow: I have not thought at all about changing the name, but I'll try to think on line about it. I think the title *Biological Cybernetics* tends to make one think "Oh, this is going to be about homeostasis and feedback." And perhaps you don't really want to do that these days. Important though feedback is, it's not the only thing that we think is important. So maybe sticking to cybernetics is not so smart. If I were going to change the title I would try to make it more inclusive.

van Hemmen: Well, I provided a subtitle for *Biological Cybernetics:* "Advances in Computational Neuroscience." That was just making clear what we are aiming at in Biological Cybernetics. So I think at least as a subtitle it's fair. You could take it as the title or stick to the original name, which will soon have existed for half a century.

Barlow: I like your subtitle. Perhaps you could expand it: "Advances in Computational Neuroscience, Biomimetics, and something else. The "something else" might be "control of movement". There is a very flourishing group in engineering here, including Daniel Wolpert and Zoubin Gharamani, working on this, doing elegant experiments that have real implications about how the brain does it. Zoubin is actually more on the machine learning aspects. So the complete subtitle might be "Advances in Computational Neuroscience, Biomimetics, Control of Movement, and Machine Learning". This would of course need changes in the Editorial board, which, by the way, it is time for me to take part in!

van Hemmen: Fair enough. Say I'm allowed to make a small detour. Which aspect of computational neuroscience or theoretical neurobiology would you consider most promising? With this of course I'm asking you a bit for your personal taste. But what would you consider to be the more promising or most promising directions in theoretical neurobiology?

Barlow: Well, alongside machine learning and movement control, which I've already mentioned, I think real progress is being made is in decision theory and inference theory, which are perhaps parts of machine learning. It does seem to me that our whole approach to statistics is changing and has perhaps already changed radically. Nowadays the equivalent of Laplace saying that he had no need of God to explain the orbits of the planets would be to say that we no longer have need of God to explain the capacity of humans for rational judgments.

van Hemmen: No, I fully agree with you that decision-making, the way in which the brain does it so as to reach what we then call a conclusion, that's a fascinating topic. I cannot but agree.

Barlow: I was just going to make the comment that—perhaps I should speak for myself, but I don't think I'm capable of doing a very good job of foreseeing the major areas of advance around biological cybernetics and information theory and inference and so on. What we have to ask ourselves is how to attract the up-and-coming authors who have ideas in this direction. It is not just a matter of accepting their papers. We must show that *Biological Cybernetics*, and the readership of *Biological Cybernetics*, provide a good forum and audience for what they have to say, that here's where they will find people who understand their ideas. I don't quite know how one gets that across, but the main thing is to get good, original, submissions, rather than for us to say "this looks like a promising field", for we would very likely get it wrong if we tried to do that.

Rinzel: How can we still serve a role? I mean, early on, as you said, Horace, there were few opportunities to present theoretical notions about neuroscience. And now we have several options; there are several journals in this direction. Are we still serving the same sort of useful purpose, or are we providing an avenue for people that don't want to try to get into experimental literature?

Barlow: Gerald Westheimer says that, in the old days, when you found a bright young thing in your class, the offer of doing some experiments in a lab was very attractive. Now all they want to do is to sit in front of a computer and write a simulation or something. Offering that as an alternative definitely has its negative aspect. In any proper science, getting the facts right is always more important than any theory, even one that turns out to be right. As Darwin pointed out, false speculations are much less harmful than the promulgation of false facts, for the latter hang around polluting the literature, sometimes for generations. On the other hand all one's colleagues join eagerly in attacking what they regard as false speculations. There is simply nothing in science as important as getting the experimental facts right.

van Hemmen: Yeah, but that's okay. After all, a good theoretician has to first master the mathematics. Mathematics doesn't come by itself spontaneously. You have to work on it. And then you can reap the rewards. Say, for making a good theory, you needn't do experiments. What you have to be acquainted with is the experimental literature, and even more importantly, I think, is the discussion with experimentalists to see how they understand their own results. That, I would say, is a key aspect of any theoretical work: you talk with the experimentalist, you ask him or her "How do you see it?", and you discuss in depth what they think they have measured.

Barlow: Where you say "…for making a good theory, you needn't do experiments", I would add "although it's better if you do". Otherwise I think you are absolutely right.

van Hemmen: And in this context, I'm strongly urging that you should also learn your mathematics. Otherwise, at least as theoreticians, you cannot express your new ideas in a suitable language.

Barlow: Yes, but I don't think many people learn much mathematics after they've got really embedded into the biological aspects, or the experimental aspects.

van Hemmen: Horace, you're right: I learned it beforehand! (laughs)

Barlow: And I wish I had too! (laughs)

Rinzel: So now I guess, Horace, would you agree that we've seen an increasing amount of modelling papers, theoretical papers, in the traditionally experimental literature?

Barlow: I think there's been an increase, yes, but I'm also slightly disconcerted. Can you name any model from, say, the last fifty years that is as complete and revealing as the Hodgkin–Huxley modelling?

Rinzel: Well, I would point to Wilfrid Rall's dendritic cable theory.

van Hemmen: Because cable theory, the way that Rall did it, allowed you to understand mathematically what's going on there.

Barlow: Well of course cable theory actually predates Hodgkin and Huxley.

Rinzel: Sure, but the use of it in understanding dendritic function I think belongs to Rall.

Barlow: Yes, I don't want to run down his contribution, but I think there may be more things going on in the neuron than the cable theory of the dendrites tells you. No doubt Wilfrid Rall's work is the basis for understanding the electrotonic structure of a neuron, but I'm very uncertain about the extent that this is crucial for understanding what the neuron does. Is it really more important than the biology of intracellular signalling and control mechanisms, the control of gene expression, calcium release and uptake, molecular diffusion, and all sorts of internal biochemical processes whose names I forget? I fear these may all be more important than anything that the electrotonic structure of the dendritic tree has to tell us.

Rinzel: Right.

van Hemmen: Well, normally these different factors complement each other, so a full understanding means that we have to understand each of the factors and therefore all of them…which of course we can't manage, but that's a different story.

Barlow: I was also thinking of neural network models. They have certainly introduced some valuable new concepts, such as parallel distributed processing, back-propagation, and Terry Sejnowski's idea of the projective zone of a neuron, which are all important, and I think the latter has been rather neglected. The fact that a neuron whose cell body sits, say, in the primary visual cortex can control the destination of its axons, maybe not down to the nearest neuron, but perhaps down to the nearest column, right at the opposite extreme of the brain; that seems to me to be an astonishing thing which we really hardly begin to understand, although the brain depends absolutely crucially on the neurons having appropriate projective zones. And more recently neural network thinking has introduced many statistical concepts into brain modelling, but here it is the statistical concepts that are important, not the network modelling. So altogether I'm disappointed.

Rinzel: Some people feel very strongly that a useful model is one that can be falsified. And I'm not so sure that we're teaching our students that aspect of modelling. You know, many people come away with the idea that we have to make a model that describes the situation, but don't go so far as to say: "How can we test it, how can we break the model?" Do you have any comments about that?

Barlow: Yes, you're spot on, you know. And perhaps it's a gap in how we educate the modellers. One should be constantly asking the question: "How could I disprove this, how could this be shown to be wrong?" If you can show that a whole class of models isn't going to do what you want it to do, you've made a really giant step, whereas if you just show that one specific realization of a model works, you've made a tiny step. The big steps are those that exclude a whole lot of possibilities. I would like to add a point here.
I think measurements of absolute efficiencies are especially important because they can exclude whole classes of inefficient possible mechanisms. I am not thinking so much of metabolic efficiencies—for example, that many animals are so inefficient at generating mechanical energy that they cannot possibly fly, but rather of statistical efficiencies. For example, the statistical efficiency of low luminance vision is reasonably close to the fraction of light actually absorbed by rhodopsin, and this single fact rules out all models of the transduction process that are not close to 100% efficient in making use of every quantum absorbed.

van Hemmen: Looking back, what do you think the key step in developing a good model is?

Barlow: Well, I think one of the benefits of modelling perceptual processes was that it revealed the natural difficulties of doing perceptual tasks. I mean, it's obvious when you start trying to make something that flies, that the low density of air is a prime obstacle. If you've got to keep something aloft with it, you've got to be moving rather fast, and you've got to shape things carefully: all sorts of things that, once you've started into the subject, seem to be absolutely obvious, but they're not obvious at all until you try to do it. And it seems to me that a useful model is often one that identifies what I call a "natural difficulty" in performing that task, and how that specific problem can be solved. Does that make sense?

Rinzel: Yes, I think so. Leo, I think a key step is to identify among the different experimental observations: which are the ones that you can address with some class of models, and which are those that you have to put on the back burner—although not forget about. But the important thing is to identify which set of questions you think you can answer with a model.

van Hemmen: Yes.

Barlow: … and make sure they're questions that need to be answered, too! (all laugh)

van Hemmen: So, Horace, it's you who—you needn't!—but who may formulate a piece of advice for future generations of computational neuroscientists, for future generations of theoretical neurobiologists. What should they keep in mind? What would be your advice?

Barlow: Well, learn how experimentalists think. Because it's only when you understand how experimentalists think, that you can say things to them which they will find interesting and important.

van Hemmen: Ah, excellent! That's more or less why I advocated this discussion with experimentalists. Thank you very much; this was great.

Barlow: I enjoyed it, too.