Elastic Systems: Case of Hebbian Neurons

Heikki Hyötyniemi*
*Helsinki University of Technology
Control Engineering Laboratory
P.O. Box 5500, FIN-02015 TKK, Finland

Abstract

As the neuronal (and cognitive) processes are, after all, so well known, they offer a nice testbench for complexity theories, and they are a nice prototype for understanding behaviors in cybernetic systems in general. It turns out that when applying the framework of *Hebbian neurons*, many observed brain functionalities can be attacked, including *sparse coding* on the low level, and *chunking* on the higher one. Even the mysteries of *causality* and *mind vs. matter* can perhaps be given new perspectives.

1 Introduction

Modeling of neurons is a nice application area — on the other hand, it is well-known (at least much-studied), but on the other, the neuronal system seems to be among the most complex ones to be understood. Brains and mental functions, or matter and mind, seem to be incompatible, and are a subject of age-old disputes. There are different levels, and emergence of higher-level structures from lower-level ones is evidently necessary — but, what is more, the different levels seem to be related to completely different *phenospheres*. First, there is the ecosystem of neurons; second, there is the infosystem of signals; and last, there is the "ideasystem" of concepts. All of these reside in the same physical medium, and powerful conceptual tools are needed to detect and distinguish between the phenomena.

The brain and cognition have long been seen as cybernetic systems. However, there are disputes: Heinz von Foerster coined the term *second-order cybernetics* and claimed that a cybernetic system (observer) cannot study another cybernetic system of the same complexity.

Here, it is claimed that *neocybernetics*, and, specifically, analysis of *elastic systems* offers a consistent framework that can capture brain and mind related phenomena in all of the phenospheres and in all of the levels. In the neocybernetic framework, all relevant constructs are *statistically determined dynamic equilibria*, being *attractors of the dynamic processes in the data space.* Hierarchies of higher-order cybernetics are reduced back to analysis of data properties.

From the point of view of AI, the key point in elastic systems is that when applying such a strong framework, it is possible to combine quantitative data processing and still reach qualitatively relevant models. This claim will be elaborated on in this paper, being continuation to the presentation in (Hyötyniemi, 2006a).

2 Modeling of neurons

It turns out that the neuronal system is *isomorphic* with the general elastic system. The main properties of such systems remain intact, but there are some issues that need to be studied from another point of view.

2.1 Hebbian perceptrons

Neurons are extremely complex electrochemical entities, and interactions among them are based on asynchronous, pulse-coded signals. In practice, such neuronal systems cannot be studied in detail, but simplifications are necessary. It turns out that when one abstracts away the time axis, studying average activity levels instead of the individual pulses, one can reach practical models — this is the approach in practically all artificial neural network structures (Haykin, 1999). In its basic form, the activities of the environment $\bar{u}_j$, where $1 \leq j \leq m$, and the activities of the neuron grid $\bar{x}_i$, where $1 \leq i \leq n$, can be coupled applying the *perceptron model*, so that $\bar{x}_i = f\left(\sum_{j=1}^{m} w_{ij}\bar{u}_j\right)$. Here, the activity of a neuron is simply a weighted sum (weighting factors w_{ij} being the synaptic strenghts) of the input activities, being further modified by the activation function f. When such expressions for the whole neuron grid are collected together, and if the nonlinearity is ignored, one has the simple matrix formula

$$\bar{x} = W\bar{u}, \tag{1}$$

where W is the matrix of synaptic weights mapping from the input vector $\bar{u}$ onto the neuronal state vector $\bar{x}$. This simplistic formulation is employed here, because it is all one needs. It directly complies with the elastic systems framework, with $\bar{\phi}^T = W$, and it fulfills the neocybernetic ideals[1]. The model (1) defines a static mapping between sets of variables, so that variables are coupled in terms of a set of constraints. Again, this pattern can be extended applying elasticity assumptions and dynamic tension considerations into another forms. The balance pursuit property of the neuronal system is crucial, giving rise to elastic master-slave behavior: No matter what is the input $\bar{u}$, neuronal state $\bar{x}$ finds its equilibrium. The properties of the converged system are reduced back to statistical properties of the environment.

To keep the system balanced, non-idealities need to be employed. The normal approach is to introduce the non-ideality in the form of nonlinearity, but following the neocybernetic guidelines, another kind of non-ideality is used: Here it is assumed that the input u is exhausted as it is exploited, leaving only the residue $\bar{u}$ visible. The assumption of information flows always being carried on a material flow is a general principle, and it constitutes now the stabilizing (linear) feedback between the system and the environment.

In the neuronal system the assumed evolutionary dynamics is also nicely manifested. According to the Hebbian learning principle (Hebb, 1949) it has been observed that for real neurons there holds

> If the the input activity (now $\bar{u}_j$) and neuronal activity (now $\bar{x}_i$) correlate, the synaptic weigth becomes stronger.

This exactly equals the evolutionary goal of an elastic system of maximizing the coupling with the environment by adapting w_{ij} in the direction of $\mathrm{E}\{\bar{x}_i\bar{u}_j\}$, as studied in (Hyötyniemi, 2006a).

2.2 Sparse coding

What happens when Hebbian learning boosted with feedback takes place in the neuron grid?

[1]Note that nonlinearity is traditionally included in the perceptron model for pragmatic reasons: First, nonlinearity promises more functionality, and nonlinearity assures stability (compare to *Oja's rule*). In the neocybernetic spirit, linearity is pursued, because only then scalability of the models can be reached. It is the dynamic feedback structures that provide with nontrivial functionalities — like *self-regulation* and *self-organization*

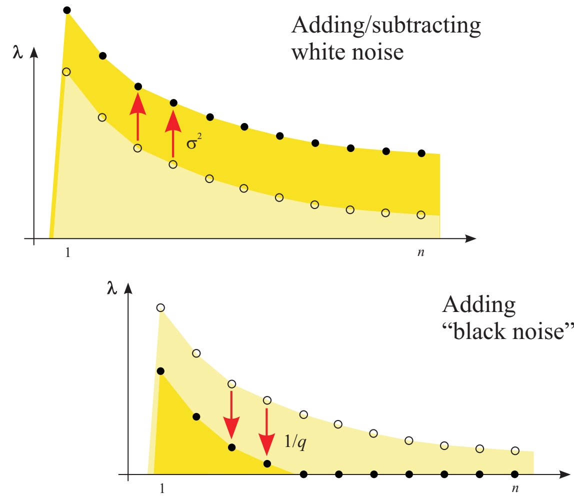

Figure 1: "Black noise" as compared to white noise: Variance structure of data in different directions

In (Hyötyniemi, 2006a) connections among $\bar{x}$ and $\bar{u}$, and among $\bar{x}$ and Δu were studied. When studying the theoretical mapping between $\bar{x}$ and the original undisturbed input u, it turns out (see Hyötyniemi (2006b)) that the eigenvalues of $\mathrm{E}\{\bar{x}\bar{x}^T\}$ can be expressed in terms of the n most significant eigenvalues λ_j of $\mathrm{E}\{uu^T\}$. Specially, if the coupling coefficients q_i and b_i are different for different neurons, the i'th eigenvalue (or latent variable variance) becomes

$$\frac{\sqrt{q_i\lambda_j} - 1}{b_i}, \tag{2}$$

indices i and j being ordered randomly. This reveals that there must hold $q_i\lambda_j > 1$ for that input variance direction to remain manifested in the system activity — if this does not hold, variable $\bar{x}_i$ fades away. On the other hand, for the modes fulfilling the constraint, interesting modification of the variance structure takes place; this can best be explained by studying a special case. Assume that one has selected $q_i = \lambda_j$ and $b_i = 1$ for all pairs of i and j. Then the corresponding variances become $\lambda_j - 1$ (see Fig. 1). In each direction in the data space, the effect of the system is to bring the variance level down if it is possible. Analogically, because white noise increases variation equally in all directions, one could in this opposite case speak of "black noise".

What are the effects of this addition of black noise in the signals? First, it is the principal subspace of u that is spanned by the vectors $\bar{\phi}_i$. But assuming that this subspace is n dimensional, there exist many ways how the basis vectors can be selected, and some of the selections can be physically better motivated.

For example, in *factor analysis* the PCA basis vectors are *rotated* to make them aligned with the underlying features, and the same idea takes place in *independent component analysis.* In factor analysis, it is assumed that the underlying features can be characterized in mathematical terms applying the idea of *sparseness:* When a data vector is decomposed, some of the latent variables have high scores while the others have low scores, increasing the differences among latent variable variances. This goal can be iteratively implemented in terms of criteria like *varimax* or *quartimax*, etc. In its extreme form, sparsity means that there are only a few of the candidates employed at a time, and the goal of modeling, rather than being minimization of the number of overall model size, it is the minimization of *simultaneously active constructs.* This means that the total dimension of the latent basis n can even become higher than the dimension m of the input data, the basis being *overcomplete.*

As shown in Figure 2, the *Hebbian feedback learning offers an efficient approach to achieving sparsity-oriented basis representation of the PCA subspace.* Whereas the overall captured variation (shown both in lighter and darker color in the figure) is not changed by orthogonal rotations, the variation over the bias level (shown in darker color) *can* be changed. As the nominal PCA approach typically distributes variation more or less evenly along each latent variable, it is most of the variation that remains below the threshold level; now, as it is the area above the threshold level that is maximized, non-trivial basis representations are reached.

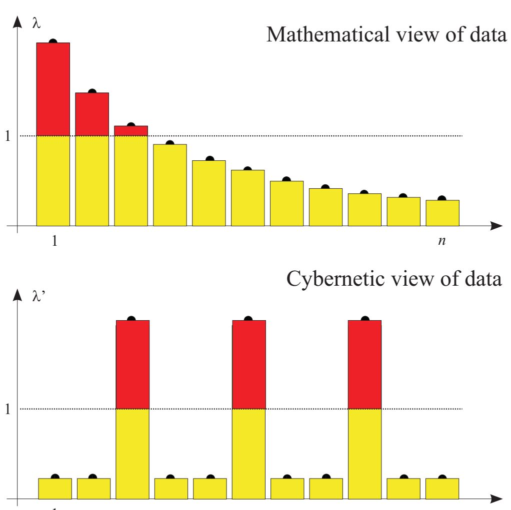

Figure 2: Schematic illustration of how black noise results in sparsity pursuit: Area above the threshold is maximized

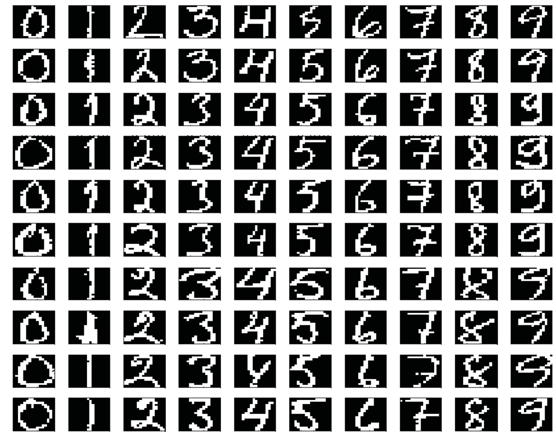

Figure 3: Examples of handwritten digits

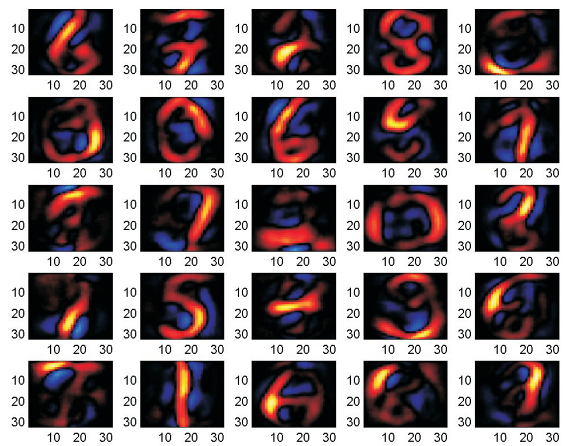

Figure 4: The 25 sparse components extracted from the handwritten digits. The 1024-dimensional vectors $\bar{\phi}_i$ are projected back onto the two-dimensional canvas, vector entries with high positive values being presented using lighter colors following the "hot/cold" color map. Clearly, different kinds of "strokes" emerge as separate memory representations. Ordering of the sparse components is random, but different runs result in essentially the same components being extracted

Figure 4 illustrates the sparse coding behavior of the feedback Hebbian neuron grid. As data material, there were over 8000 samples of handwritten digits (see Fig. 3) written in a grid of 32×32 intensity values (Laaksonen, 1997). The 1024-dimensional intensity vectors were used as data u, and each of the latent variables $\bar{x}_i$ was kept active by appropriately controlling the coupling factors q_i. To enhance convergence, the "cut" nonlinearity was additionally employed to emphasize sparsity of coding (see Hyötyniemi (2006b)). The behaviors differed very much from principal component coding: In the beginning, something like clustering emerged, each data region being represented by a separate variable, but as adaptation proceeded, the features started becoming more orthogonal, and patterns were decomposed further. What is interesting is that this kind of "stroke coding" has been observed also in the visual V1 cortex region.

The above view of neuronal functioning can be extended to more complex data, just by interpreting the data-based constructs in new ways.

3 Towards modeling of cognition

The neocybernetic framework not only allows modeling of the coding of individual patterns, it can perhaps give tools to attack the functioning of the complete brain. There is the intuition backing up us here: The cognitive system simply *has to be* cybernetic — even in various ways (Heylighen and Joslyn, 2001).

3.1 Population of neurons

Employing the idea of looking at the neurons as a population of competing individuals, one can see the neuronal "ecosystem" as a self-regulative entity. The computer paradigm with "memory registers", etc., collapses; there is no transfer of information, no separate long-term memory or short-term memory elements, but it is all an integrated whole. No central control is necessary, nor some "operating system", it is all about distributed pursuit for resources in terms of excitation, or variation in signals. The winning neurons start representing the corresponding association structure, defining a (subconscious) "concept atom". As atomary concepts are connected to previously activated ones, sequences of concepts emerge. In the long run, when the time structure is ripped off, some kind of a *semantic net* emerges.

This all is more or less familiar — the added value, the main contribution of the neocybernetic perspective, comes from the ability of explaining how the above declarative representations change to associative ones, or how the *shift from novice to expert* can be explained. The key functionality is the self-regulation and self-organization property of the Hebbian feedback system: As the Hebbian adaptation takes place, locally and independently in each synapse, the declarative structures become swallowed in an "associative medium". As correlating concepts are appropriately connected together, following the neocybernetic adaptation principles, links of the semantic net become denser and numerically optimized, and bidirectional. All concepts are defined in terms of examples and associations to other concepts — because of the numeric, dynamic platform, the hermeneutic cycles converge to a balance of connotations.

The above process of automatisation is the key process in the mental system. But the inverse process is equally important: The high-dimensional associative representation has to be decoded into a one-dimensional representation to facilitate any information transfer, or processes like reasoning, and *thinking* in general. Such issues are not studied here — but the claim is that if the link between declarative sequential representations and associative parallel ones someday is found, ultimate homogeneity of mental functions can be reached: There is no need for special structures when mental faculties are implemented.

3.2 Role of semantics

It is assumed here that intelligence is an illusion that emerges when a large number of simple structures cumulate. The principles assumedly remain the same also on the new emergent level, so that the processes can be reduced back to processing of data. When proceeding from the level of signal processing to processing of information and knowledge, one is facing new challenges, because one needs to address issues that are the most relevant to the human mind: A cognitive model is void, its essence escapes, giving rise to *Chinese room* type arguments, if it does not somehow capture the *semantics* of the constructs. One needs to extend from the infosphere, where it was simply data (co)variation that needed to be captured, to "ideasphere".

For concrete modeling purposes, one needs to be capable of reductionistically decomposing the cascaded system hierarchy into self-contained entities. Now, assuming that these "idea atoms" emerge from lower levels, being individual cybernetic models for subdomains, how to avoid the infinite recess, concentrating on a single level, truncating the succession of models? In other words: How to assure that the data delivered to a cybernetic system constitutes a "cybernetic closure"? How to fix the grounding of semantics, or make the concrete data contain the necessary "atoms of semantics"?

The concept of semantics needs to be formalized at some level. When processing signals, information being interpreted as (co)variation, one concentrates on *contextual semantics*, where the meaning of the structures is determined in terms of their interconnections, finally reducing back to the system inputs (*natural-*

istic semantics). For a cyberenetic system, however, this kind of static definition is not enough, one once again needs to extend the studies to dynamic domain. It was dynamic balances that were the key issue in neocybernetics, and the cybernetic models are models over such equilibria. These balances need to be buried in data, or, the data needs to be made balanced.

In each state there is a tendency to move in some direction. This "flow" is proportional to the unbalanced tensions in that state, and can be utilized to quantify the counteracting forces. Such tensions are also visible in the observation data: State changes, or differences between successive states are proportional to the flow. When such derivatives are included in the data, they stand for the additional compensating forces that are needed in that state to reach balance:

$$u'(k) = \left(\frac{u(k)}{\frac{du}{dt}(k)} \right) \approx \left(\frac{u(k)}{u(k+1) - u(k)} \right). \tag{3}$$

Such "preprocessing" of observations, emphasis on changes or differences between successive ones, can also be motivated in terms of psychological and neurophysiological studies.

As an example of the relevance of the above discussion study a case where chess configurations are modeled. Chess is the "banana fly" of cognitive science, being a simple domain, but still being far from trivial. There were some 5000 configurations from real games used for modeling[2]. The coding of the configurations was carried out so that for each location on the board (altogether $8 \times 8 = 64$) it was assumed that there are 12 different pieces that can in principle be located there, and for each of them there was a separate entry in the data vectors. This means that there are altogether $64 \times 12 = 768$ binary entries in the highly redundant data vectors — and when the derivatives were included and u' was defined as in (3) the data was $2 \times 768 = 1536$ dimensional. In Fig. 5 the results are presented when 100 memory representations or feature prototypes, or chunks, were allocated for this chess coding task. These chunks $\bar{\phi}_i$, where $1 \leq i \leq 100$, were extracted from the training data, and after convergence a typical configuration was reconstructed as a weighted sum of the chunks. In the figure, the modeling results are approximatively illustrated by projecting the numeric representations back onto the discrete-valued realm.

It is interesting to note that it has been claimed that some 50000 chunks are needed to satisfactorily represent the chess board (Chase and Simon, 1973). Now the numeric nature of the chunks and inherent optimization of the representations makes it possible to reach a much more compact model for a domain.

The results are interesting, remotely reminding the mental operationing of a real chess expert: It is known that chess experts only concentrate on the "hot spots" on the board, but this kind of attention control has not been satisfactorily explained. Of course, the current experiment only studied very elementary patterns on the board, and to capture phenomena like *functional chunks,* to reach towards really "understanding" the game, one could introduce more complex (cybernetic) preprocessing of the observations:

$$u''(k) = \left(\frac{\bar{u}'(k)}{\bar{x}'(k)} \right). \tag{4}$$

3.3 Epistemology of constructs

In today's AI paradigms (like in *semantic webs* and earlier in expert systems), it seems that one is interested in *ontologies.* However, the essence of knowledge is not in the objects but it is in the ways of conceptualizing and representing them. What kind of *epistemologies* are dictated by the underlying "wetware"?

As it is assumed that it is essentially the same Hebbian perceptrons that implement all the functionalities, there is the common neural basis of the cognitive constructs, dictating their structure. The linear, sparsely combined features are best characterized as *degrees of freedom* in data space. The "conceptual spaces" are now not based on clusters in the data space but on additive axes of degrees of freedom. Because of this uniformity, it must be so that for example *categories* and their *attributes* have essentially the same kind of structure, each determining the other. The "is-a" hierarchies and "has-property" structures become unified. Also subclasses, and, specially, instances of classes, are similarly represented as interconnected degrees of freedom. This means that the framework of *fuzzy subsets* offers an appropriate epistemology for mental constructs — subclasses belong to superclasses, but also *superclasses belong to subclasses* (see Fig. 6). Normally, category properties (couplings to other categories) are stored in the prototype itself if they are not shared, but if an attribute is common to many categories, it becomes manifested as a separate structure of its own (compare to the "strokes" as atoms of relatedness in Fig. 4).

Perceptions are lower-level observations that are filtered through the mental model. This structure is not, however, purely hierarchic — higher-level perceptions also affect lower-level ones. In concrete terms, $\bar{x}_i$ determines the relevance of the concept

[2] I am grateful to Professor Pertti Saariluoma for the data material and for encouraging discussions

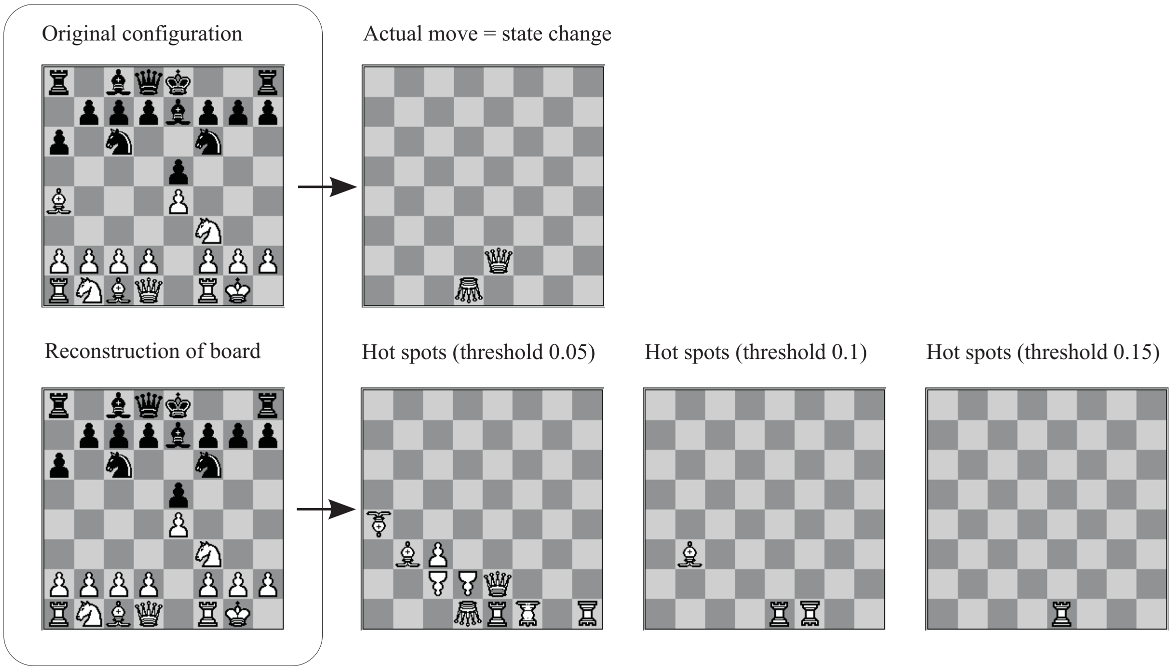

Figure 5: An example of how high dimensionality makes it possible to mimic cognitively relevant functionalities. First, when using only static configuration data, it turns out that *chunking* can be studied; second, when the derivatives are also included, it becomes possible to attack the challenges of *attention*. On the leftmost images, the observed chess piece configurations $u(k)$ are presented: On top, there is the outlook of the original board, and on the bottom, there is the reconstruction when using a storage of only 100 numeric chunks that are appropriately stacked on top of each other. In such a typical case, almost all pieces can be correctly recalled (the vector $\hat{u}(k)$ is thresholded so that only pieces with relevance $\hat{u}_j > 0.5$ are shown). The remaining images illustrate the "flow" of the game, or derivative $\frac{du}{dt}(k)$ in the current state k: Again, on top, there is the observed change in the configuration, and on the bottom, there is the estimate, visualized applying three different threshold levels. The pieces upside down denote vanishing pieces. It turns out that the "hot spots" are located in relatively appropriately (even the missing bishop is now there), and, as it turns out, it is indeed the expert-selected move that has a strong representation — even though it is not the winner. Note that the reconstruction is purely asociative, and no check for validity is here carried out, so that some "ghost spots" also exist. On top of the associations, higher-level reasoning is needed to screen the reasonable moves

(category/attribute) number i when perceiving the input. As seen in another perspective, the sparse coded momentary weights $\bar{x}_i$ stand for the cognitivistic notion of *short-term memory*, containing "indices" to *long-term memory* constructs. These LTM constructs are the profiles $\bar{\phi}_i$ expressing the elementary patterns of exhaustion of available activation. This scheme is completely distributed and locally controlled; the computer paradigm with its centralized registers, memory units, and data transfer among them, can be abandoned in this framework.

In the system theoretical perspective, relevant concepts are *attractors* in the data space: In appropriate conditions, when the incoming signals match the model, dynamics ends in the corresponding basin of attraction. Because of the distributed nature of the model, many of the available attractors can be simultaneously active. The nonlinear, multimodal distributions of natural data are thus represented by separate mental substructures, each competing for activation.

The contents of a concept are determined essentially in terms of relations to other constructs, in the simplest case these constructs being *examples*. The uniformity of mental phenomena can be extended outside the nominal concepts: Even the content of *feelings* is determined by prototypical experiences.

3.4 On expertise

One can claim that *expertise* in a specific domain is based on appropriate features or chunks existing in the conceptual space. An expert matches the observations against his mental view, thus compressing the data into domain-oriented representations. Applying this low-dimensional representation, missing variables are "filled in" as the known variables are

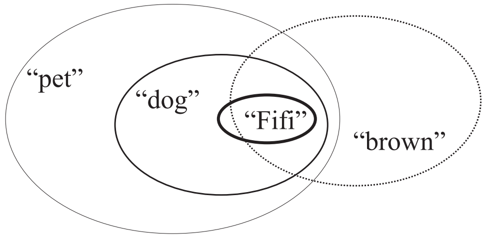

Figure 6: When interpreting the neurons-defined structures in epistemic terms, it can be said that the concept hierarchies become fuzzy: A dog is a subclass of a pet, and Fifi is a subclass of a dog — but, simultaneously, a dog is a part of the contents of a pet, and Fifi is part of dog. Inheritance is not hierarchic but becomes a network: For example, examples of a dog determine what brown color is like, and the concept of brown partly define what dogs are like. Speaking of dogs activates associations to pets, and *vice versa.* Such discussions become less trivial in higher dimensions

matched against the model, and this way, “associative inference” is implemented (see Fig. 7).

The distribution-oriented view of expertise allows subtle, non-binary reasoning, and also merciful degradation of mental capacity as a function of scarcity of resources is manifested. This is visualized, for example, by the chess experiments, where the errors are “expert-like” (again, see Fig. 5).

There is a close connection to case-based reasoning (CBR) here, but now there is some kind of “functional” matching of the mental model against data taking place: All input entities do not have the same significances, but the knowledge structure is taken into account. A solution to the *frame problem* is also reached: The high-dimensional knowledge representations are never isolated from their surroundings, but different kinds of “defaults” are automatically carried along.

4 Contribution of neocybernetics

There are many additional intuitions that are offered by the neocybernetic approach.

Causality. The mapping from u to $\bar{x}$ is valid only if the closed-loop system is stable, that is, the system can affect the environmental variables and balance them. This is not always the case, and analysis of covariance properties of u do not necessarily reveal the system properties. A more fundamental relationship is that between $\bar{x}$ and Δu, the mapping as seen by the system itself. This mapping is interesting because it essentially represents the *effect of the system*

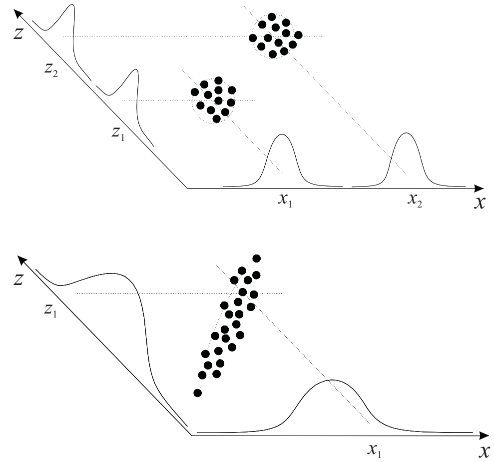

Figure 7: Traditional view of expertise (on top) makes it possible to implement rules of the form `IF` $x = x_i$ `THEN` $z = z_i$, etc., whereas when distributions are employed, inference becomes subtler, being an associative (maximum likelihood) pattern matching process against the existing high-dimensional knowledge model

back onto the environment.

As observed originally by Hume, one cannot ever see *causalities* in data, only *correlations*, that is, one cannot detect cause/effect relationships. But as Kant has said, it is causal structures that are constructed by the human mind. How is this possible then without employing some *a priori* understanding?

It can be claimed that the neocybernetic model offers a solution to this dilemma. Because it is only ones own actions Δu, as induced by the environment, that are being coded in $\bar{x}$, one implicitly knows the structure among causes and effects — there is no paradox any more here. True causality structures are indeed built deep in the Hebbian feedback models.

Consciousness. There are many contradicting intuitions of how consciousness should be defined — the heated controversies being, of course, caused by the fact that consciousness is the essence of our speciality among animals. The views vary from the highest (consciousness is the culmination of intelligence) to the lowest level (consciousness is ability to feel something like pain), or even below that (consciousness can only be explained in terms of quantum effects).

Awareness of *self*, or “knowing that one knows that one knows”, is assumedly a holistic, emergent phenomenon that cannot be reduced. However, in the adopted framework this structure of infinite recess can again be collapsed, and a concrete definition for consciousness can be proposed. In the neocybernetic spirit, it can be assumed that the mental machinery constructs a more or less sophisticated model of the environment; when this model becomes complex enough, the "self" becomes a relevant entity in the model that succesfully helps in structuring the observations and behaviors in the environment. When such self emerges as an independent entity in the mental model, there is consciousness. This would mean that animals have consciousness in varying degrees — but also non-biological cybernetic systems would be conscious to some extent. On the other hand, a small child not distinguishing itself from its mother is *not yet* conscious — but the "brain prosthesis" can truly capture the mind.

Intersubjectivity and interobjectivity. In complexity theory that is based on chaos theoretical starting points, it is often claimed that the value of the models is questionable as small deviations in initial conditions result in completely different outcomes. Now, on the other hand, stochastic variations are not significant: It is statistically relevant constructs or attractors of dynamic processes within the phenosphere that are being captured, and role of the transients fades away.

This observation has its effects also on AI, or, indeed, to the *theory of mind* itself: If the modeling principles are the same, reaching towards optimized representations, and if the environments are the same, two different modeling processes end up in essentially the same "world view". This applies not only to humans that have their subjective experiences — this *intersubjectivity* can be extended also between natural and artificial brains. Constructing "smart machines" in the deep AI sense *is* possible.

One can even extend these considerations beyond the mind, to the principles of modeling in general: The goal of AI is to make the the computer somehow understand — or model — the world. In (Hyötyniemi, 2006b) it is observed that nature itself constructs models — a higher trophic layer that is evolutionarily reasonable tries to capture the behaviors of the lower layer to maximally exploit it. If a human (or computer) then models the same ecosystem (applying appropriate modeling framework), is the artificially-made and nature-made models somehow essentially different?

In modeling theory one is traditionally humble: It is assumed that a model can only capture a shadow of the real-life complexity. But if the real life is also a model or a model-based simulation, *interobjectivity* can be reached, where the man-made model is essentially the *same thing* as the reality itself. Essence of systems is not in the building blocks but in the information structures. Epistemology becomes ontology. It may be that *metaphysics* finally becomes a science.

5 Conclusions

To goal of complex systems thinking is to find simple underlying principles beyond different kinds of complex phenomena. In the neocybernetic perspective, it seems that there truly exist common basis for neural/cognitive systems and other cybernetic ones.

It is not only so that cybernetic understanding can be exploited in AI and analysis of cognitive systems, as was done above — there is also contribution in the inverse way. Perhaps the intuitions from AI can be employed when studying what kind of ontologies there can exist in other cybernetic systems. For example, can the structures emerging in *genetic systems* be studied in terms of similar concepts? What is more — can the genetic code be translated into a human-understandable language?

And the social systems among individual people — it may turn out that extending the cognitive system beyond the limits of a single brain can be understood and analyzed in the neocybernetic framework. Constructing a common "supermodel" among the individuals is all dependent of the information transfer between the minds. Perhaps the idea of EQ or "emotional intelligence quotient" is the IQ of intelligent organizations?

References

W.G. Chase and H.A. Simon. *The mind's eye in chess*. In *Visual Information Processing* (ed. Chase, W.), Academic Press, New York, 1973.

S. Haykin. *Neural Networks — A Comprehensive Foundation*. Prentice–Hall, Upper Saddle River, New Jersey, 1999.

D.O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. John Wiley & Sons, New York, 1949.

F. Heylighen and C. Joslyn. *Cybernetics and Second-order Cybernetics*. In R.A. Meyers (ed.). Encyclopedia of Physical Science & Technology (3rd ed.), Academic Press, New York, 2001.

H. Hyötyniemi. *Elastic Systems: Another View at Complexity*. SCAI'06, Helsinki, Finland, 2006a.

H. Hyötyniemi. *Neocybernetics in Biological Systems*. Helsinki University of Technology, Control Engineering Laboratory, 2006b.

J. Laaksonen. *Subspace Classifiers in Recognition of Handwritten Digits*. Acta Polytechnica Mathematica; Mathematics, Computing and Management in Engineering series No. 84, Espoo, 1997.