SYSTEMS THEORY VS. THEORY OF MIND: TOWARDS A SYNTHESIS?[1]

Heikki Hyötyniemi[2]

Helsinki University of Technology, Control Engineering Laboratory
Konemiehentie 2, FIN-02150 Espoo, Finland

Systems theory offers tools for analysis and construction of models for different kinds of systems. It turns out that these tools may open new horizons also when studying mental phenomena. This paper illustrates the new possibilities, and proposes a new view of "mental imagery" that attacks some of the old paradoxes in the field.

1. INTRODUCTION

Systems theory analyses systems. What, then, is a system? Perhaps the most revealing definition for this is that "a *system* is an entity that can be seen as a system"! This definition is based on intuitive notions; it embraces everything because all more specific definitions would be too restricting. For example, see Fig. 1 – there are two "black boxes", the ideal mixer in 1a and the "idea mixer" in 1b; both of them can be seen as systems, having some internal dynamics and well-defined connections to outside world. Of course, the mystery lies in the very different basic nature of these two systems: Whereas the ideal mixer maximizes the entropy, the mental machinery minimizes it, creating order and structure.

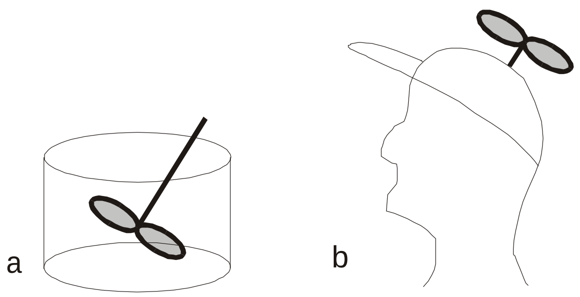

Figure 1. Two vessels

No matter what is the application field, systems theory takes the observed facts and tries to construct a model that would explain the observations, optimising some more or less concrete criterion. In our field there is a plenty of conflicting evidence available. If the Kuhnian view of scientific progress is adopted, it can be said that now we have the antitheses, and some kind of synthesis is waited for. It can be claimed that systems theory is the "science of syntheses" … Let us see what can perhaps be reached.

---

[1] This presentation was originally given during the Symposium on "Representation", organized in Dec. 17, 1999, by the Finnish Artificial Intelligence Society and Finnish Philosophical Society, at the University of Helsinki. The paper is, however, original
[2] Currently working for Nokia Research Center

2. MAPPING THE ARENA

Systems theory is more like philosophy than a compact toolbox of ready-to-use methodologies. Regardless of the application field, though, the objective is always to find models that would reach a good compromise between conflicting needs and boundary conditions. The model is just an approximation of the reality, absolutely correct models do not usually exist; the validity of a model can be evaluated using different criteria. In addition to matching the observations, there are various other aspects that should also be taken into account (see Fig. 2): For example, in what respects does the new approach outperforms the old ones, would it really be motivated to change the established practices? How about the practical issues – how can the model be analysed, and how easily it can be applied in different cases? These questions are concentrated on separately in what follows.

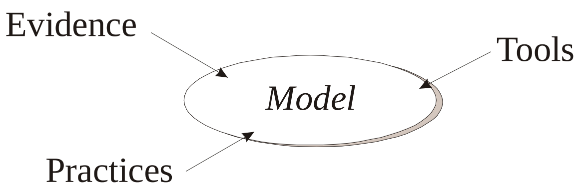

Figure 2. Boundary conditions

2.1 Available Evidence

Usually when constructing models, the observations are rather compact and well-defined; it is not a surprise that now the observations are more difficult to quantify. Now it is all kinds of intellectual activities that tell us something about the principles of mental functions, and some of these branches are discussed below. Very much evidence has been ignored; only those facts that are explicitly attacked are mentioned: These are now our "starting hypotheses", from the "micro-world" level towards the "macro-world" level.

Neurophysiology. The brain functions are implemented using neurons; even though the gap between the lowest physical level and the level of cognitive processes is huge and conclusions should be drawn with extreme care, some experiences are still perhaps useful. The analysis of sensory signals, specially in the visual cortex, shows that a special kind of pattern reconstruction takes place: the visual image is presented in terms of elementary features. The cortical cells constitute a reservoir of feature prototypes – at any instant only a subset of all available features is utilized [12].

As motivated in [1], there cannot exist separate faculties for different mental capabilities in the brain – the same principles are responsible for all different kinds of high-level activities. This uniformity principle can be extended further: It can be assumed that this kind of *sparse coding of features* is characteristic to all levels of mental behavior.

Cognitive psychology. In cognitive science a wealth of concrete facts have been found that are in conflict with mainstream knowledge engineering methodologies; for example:

- The operational differences between the long-term memory (LTM) and the short-term (working) memory (STM) is not addressed.
- The observed fuzziness of categories [14] has not been embedded in the knowledge representations.
- The pecularities in the shift from novice to expert [3] have not been explained.

In expertise research it has been noticed that the expert reasoning becomes faster and more automated as compared to the novice reasoning. This is known as the *speedup principle* [2]. The traditional way to enhance knowledge bases is through growing the number or sophistication of the rules – but how could the expert make the conclusions faster than the novice if the rules were more complex?

Different cultures. It has been recognized that the cognitive categories are not uniquely determined: Even though the physiological machinery is the same and the observations of the world may be the same, in different cultures different mental classes may be constructed; for example, the colors differ between cultures. This refers to the fact that even the simplest categories are not hard-wired or Platonian ideals, but the interaction with the environment may affect the resulting mental constructs.

Another source of insight is the Eastern view of Knowledge and Understanding. The verbal level knowledge is only the first level; only after the words have been eliminated, real enlightment can take place. In Western cultures, this kind of non-thinking is allowed only in arts and poetry – but this is what makes as special as humans (see the Preface of this Volume).

2.2 Current Practices

"Good Old-Fashioned AI" was based on symbols and constructs often inherited directly from linguistics. The shortcomings of the linguistic representation (and, more notoriously, the "biases" caused by this starting point) are discussed in [4]. However, the language-based approaches offer a very sophisticated framework for representing structure among entities. This underlines the main shortcoming what comes to connectionistic approaches: It is their lack of structure – traditional neural networks training algorithms that concentrate on the input-output mapping do not facilitate *emergence*. Neural networks and fuzzy systems all share the problem of *unscalability*: Even though they may work in "toy worlds", their performance may collapse in more complicated environments of higher dimensionality.

The self-organizing maps are a step towards emergence; however, the structures that can emerge are rather simple. How to combine scalability, emergence, and flexible structure in the same framework?

2.3 Modeling Tools

There are various frameworks that are proposed for modeling of mental phenomena, most notably perhaps ACT-R [1] and SOAR [11]. The basic problem is that these models become increasingly complex, thus making them intuitively less appealing; they can be used to explain behavioral patterns, but they are not suited for predicting or gaining intuition. In a way, they are "too powerful" frameworks; the same criticism applies also to many neural networks formalisms: It has been proven that recurrent perceptron networks can implement *any computable function.* However, having no prior restrictions to the class of functions to be searched for, immense amounts of data are needed for training.

"Make it as simple as possible (but not simpler)" – this is the famous Occam's (or Einstein's) razor. The simplest models are reached if the supported model structures comply with the problem to be modeled. What is then the appropriate model structure when mental phenomena are studied? What are the classes of emergent structure that one should look for? In [4], it is assumed that *sparsely coded linear structures* would do (see the next section).

Usually some fancy nonlinear functions are introduced to implement complicated tasks. However, it has not been rigorously proven that nonlinearities (sigmoid functions or the like) would be essential in the construction of mental functions. Perhaps the sigmoid nonlinearity that is found in a neuron is only a nasty limitation of the biological components that are used for computation … compare this to transistors: The global characteristic curve is, of course, nonlinear (almost like sigmoid, by the way), but the electronic constructions containing transistors usually only utilize the local, linear part of the curve (in analog devices), or if the transistors are used as switches in digital devices, only the saturated extreme parts of the curve are utilized. What is interesting is that if these two operating modes of the transistor are used in the same device, its operation can again be expressed using a sparse linear model!

The (piecewise) linearity makes it possible to utilize the "divide and conquer" idea for analysis, so that a simple substructure can be analyzed separately and later be included in the overall system. This results in reductionism; it is assumed that the illusion of intelligence emerges when large numbers of simple operations are combined. Sparsity makes the overall system nonlinear, though, so that the expressional power need not be compromised.

Now we need a formalism for representing parallel, linear phenomena; natural language is not suitable for this purpose, but *matrix calculus*[3] is.

3. MENTAL MODEL

The idea of *mental images* is a useful concept. Originally, mental imagery was studied exclusively in the context of concrete visual scenes (see [10]). However, the nature of the mental imagery is not agreed upon [13], and parallel "mental views" seem like a good approach to discuss expertise as well – the expert has internalized a sophisticated set of mental images governing the problem area. As presented below, the specialized imagery consisting of the domain-specific prior "observations" (original or modified; see later) constitutes a "filter" that preprocesses the observation data, creating a compact internal representation of the situation at hand.

3.1 Ontology and Epistemology

The fundamental role of the mental representations is to convey *semantics*, or the meaning of constructs: What is the link between the internal structures and the outside world? To evaluate the representational power of a mental model, and to compare different models reasonably, a concrete starting point is needed.

In philosophy, empirism offers a fruitful framework for AI, as contrasted to the more metaphysical paradigms. Rather than assuming that there were predestinated *a priori* mental structures, one assumes that knowledge emerges from observations. This view results in the *naturalistic semantics*, where semantics of constructs is defined by their context: how a concept is related to other ones dictates its meaning. In connectionism, these questions are called computational or procedural semantics [9]; perhaps a better name would be *contextual* or *associative semantics*, to emphasize the need of parallel processing.

When the contents of the semantic universe are learned empirically, the observed data directly dictates the contents of the mental representations. This leads to the epistemic problem setting: what is *knowledge*, and what is *truth* in the first place? Things that have been observed together many times, become coupled together – it is relevance that is of primary importance rather than "truth", determining what are the "beliefs" of the system. This way, the difficult philosophical problems can be avoided.

[3] Incidentally, the essential role of "Matrix manipulations" in the future AI environments is recognized also by Keanu Reeves *et al*. during their adventures in the Matrix (Warner Bros., 1999)!

3.2 Representation of Knowledge

The observations are now assumed to span a high-dimensional space; each scalar piece of information (call it "feature") has its own entry in the observation vector φ. The long-term memory (LTM) consists of this kind of observation vectors (or other compatible vectors; how these are constructed is explained later).

A fixed point in the observation space represents one "state of affairs": It reveals how the different features are related in that special case. Depending of the interpretation of the input features, this point can represent one single fact, or it can just as well be a rule: If some of the entries can be seen as conditions, the other features can be seen as results corresponding to the condition values. In what follows, the simple "static" facts and rules need not be discussed separately; it is how a observation vector is utilized that determines how it is being interpreted. It turns out that we need to define different "levels of knowledge":

1. "Incomplete" point in the observation space where only a subset of the entries in the high-dimensional observation vector are known; this could also be called "linguistic projection", because verbal facts (or rules) seldom determine all of the feature values.
2. "Complete" point in the observation space; the incomplete points can be (hopefully) somehow completed, resulting in "(extended) fuzzy" facts or rules.
3. "Axis" spanning a line in the observation space; using other vocabulary, these axes can be called *latent variables* or *independent components*. The structure and outlook of the vectors is the same as above; they are just used in a different way (see below).
4. "Subspace" in the observation space spanned by various latent variables corresponds to the highest level of domain-area expertise (see below).

3.3 About Perception

In principle, the modeling task is to find a set of vectors in LTM so that the new observations (hopefully having the same statistical properties as the training samples) could be represented as accurately as possible as a weighted sum of these. Simultaneously, there is the sparsity objective; of course, the two conflicting criteria (minimization of the reconstruction error $\hat{\varphi}-\varphi$, and minimization of the number of active elements) make the optimization task difficult, and no explicit formulae probably exist; iteration is needed. The algorithm below continuously tries to make the contents of LTM more appropriate (see Fig. 3):

1. **Input.** If available, input new feature-vector-form observations (or pieces of knowledge, or hypotheses, or queries) as candidates in the long-term memory.
2. **Selection.** Select next φ from the long-term memory according to some kind of novelty or "interestingness" criterion (for example, the newest one, just imported or modified).
3. **Attention control.** Select the focus, or the features that are concentrated on (for example, the elements in φ that are assumed to be known). Accordingly, construct the diagonal binary matrix W so that there is 1 on the diagonal corresponding to these focus features, all other entries being zeros.
4. **Associative matching.** Construct $\hat{\varphi}=\sum_i \Phi_i \varphi_i$ so that $(\varphi-\hat{\varphi})^T W(\varphi-\hat{\varphi})$ becomes small. Here the value of i varies among the long term memory vectors, all of them having a unique index. The total number of selected vectors cannot exceed the short-term memory capacity, though; this means that most of the weighting values Φ_i are zeros. The vector Φ stands for the sparse coded, "internal representation" of φ; this "perception" can be used further as an observation on a higher levels of mental imagery. The decomposition of φ into its components its discussed more, for example, in [6].

5. **Inference.** Define $\varphi' = \varphi + (I - W) \cdot \hat{\varphi}$ and store it in the long-term memory (here I is the identity matrix compatible with W). This means that in the new vector the focus features remain unchanged, whereas the other entries may be completed.
6. **Update.** If the capacity of the long-term memory is reached, apply some criterion for elimination of "non-interesting" (or perhaps conflicting) constructs … yes, but *how* …?
7. **Return.** Go back to Step 1 (indefinitely).

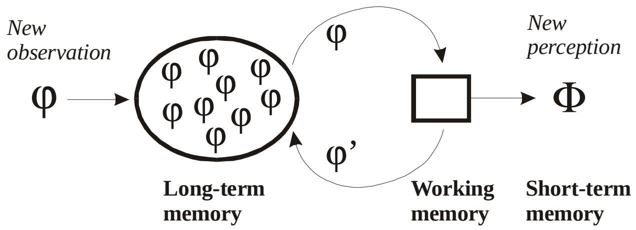

Figure 3. "Mental imagery"

3.4 Discussion

The above algorithm has its shortcomings; for example, how to control the "control of attention", and how to maintain the contents of the long-term memory? More fundamentally, how could *sequences* of incoming observations be captured? This would be the key to modeling causality and more complex somehow hierarchical observations. For example, in [5], the problem of modeling successive observations was neglected; the problem space was represented in a static, high-dimensional form. Only one level of chess expertice could be studied (see Fig. 4).

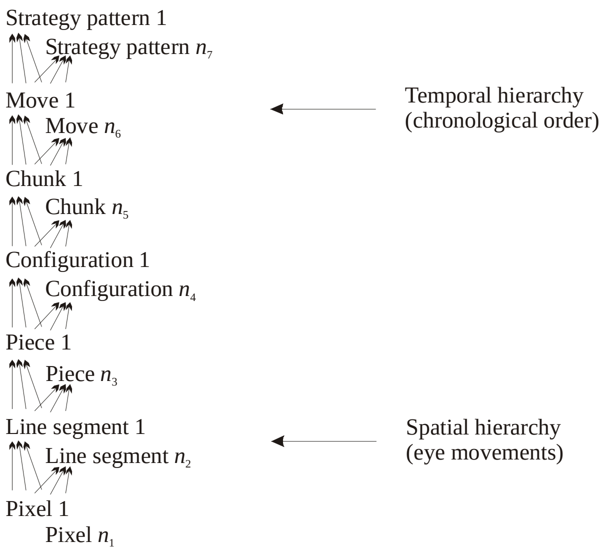

Figure 4. Levels of chess expertise

It is difficult to compare the expressive power of this kind of "data-based knowledge" model to traditional logic formalisms. For example, there are no AND or OR connectives explicitly available for constructing rules; NOT does not either exist. One has to think in terms of *relevance* – it is the correlation between vectors that counts, so that all non-zero elements in the vectors contribute in the final outcome. The same signs in the elements results in a positive effect (combined, fuzzy AND and OR together), whereas negative effect can be interpreted as fuzzy NOT. Zero elements implement a kind of "don't care" option.

From the point of view of expressive power, the linearity of the representations, being determined by sums of features, seems very restrictive. However, the sparse nature (only a subset of features is used simultaneously) makes the approach considerably more powerful; for example, the XOR problem is solved when the alternative options may be included in mutually exclusive features. This sparseness, and the changing sets of features in different environments, also makes nonmonotonic reasoning possible.

It is not necessarily only observation variables that can be included in the "observation" vector; output (motoric) signals can also be included. This means that when the observation is reconstructed the appropriate responses can be associatively constructed by the model. Thus, it is not only knowledge that can be modeled, but also *skill*.

4. EXAMPLE APPLICATION

Assume that we are told the following: "Tim, Tom, and John are children; one of them is three years old, one is six years, and one is nine. Additionally, it is known that Tim is three and Tom is *not* nine. How old is John?" Note that this is an (extremely) simplified version of the very common type of logic problems, where the pieces of information are presented in a very fragmented, incomplete form – the goal is to find a consistent set of points in the observation space.

4.1 Representing Information

Concentrate on the current level of abstraction; assume that categorizations like "Tim" or "three" are already available. There are six independent capsules of information now – the concepts of "Tim", "Tom", "John", "three", "six", and "nine", so that there needs to exist an entry for all of them in the observation vector. So, the utterance "Tim is three" can be coded as ("++" symbols denoting a (large) positive value; actual numbers are now not important):

$$\varphi_{\text{TimThree}} = \begin{pmatrix} ++ \\ 0 \\ 0 \\ ++ \\ 0 \\ 0 \end{pmatrix} \quad \begin{array}{l} \leftarrow \text{"Tim"} \\ \leftarrow \text{"Tom"} \\ \leftarrow \text{"John"} \\ \leftarrow \text{"three"} \\ \leftarrow \text{"six"} \\ \leftarrow \text{"nine"} \end{array}$$

Using the defined knowledge hierarchy, the above vactor is on the level 1 only. There is some additional information ("frame knowledge") that is not explicitly expressed in the above world model: for example, if somebody is "Tim", he cannot be "Tom" or "John"; also, if somebody is three, he cannot be six or nine at the same time. This knowledge can be reptresented using "rules" like

$$\varphi_{\text{TimOnly}} = \begin{pmatrix} ++ \\ - \\ - \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad \text{and} \quad \varphi_{\text{ThreeOnly}} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ ++ \\ - \\ - \end{pmatrix}.$$

Using the above vectors we can already construct a "complete view" of Tim. This takes two inference steps: First, selecting focus on the entry "Tim" (in other words, interpreting this entry as being the condition part of a rule), the vectors φ_{Tim} and φ_{TimOnly} can be combined, and after that, the result can be combined with $\varphi_{\text{ThreeOnly}}$ using "three" as focus; the result is then

$$\varphi_{\text{TimThree}} = \begin{pmatrix} \boxed{++} \\ 0 \\ 0 \\ ++ \\ 0 \\ 0 \end{pmatrix} \xrightarrow{\varphi_{\text{TimOnly}}} \begin{pmatrix} ++ \\ - \\ - \\ \boxed{++} \\ 0 \\ 0 \end{pmatrix} \xrightarrow{\varphi_{\text{ThreeOnly}}} \begin{pmatrix} ++ \\ - \\ - \\ ++ \\ - \\ - \end{pmatrix} = \phi_{\text{Tim}}.$$

The focus entries used during the inference steps have been drawn in boxes, and the model vectors that are matched are shown above the arrows. This is now a completely determined point in the feature space, expressing a "frame-integrated" piece of knowledge; to emphasize its nature of "second level" information, the vector symbol here has been changed slightly. Further, starting from "Tom is *not* nine" vector $\varphi_{\text{TomNotNine}}$ the forward-chaining process results in

$$\begin{pmatrix} 0 \\ \boxed{++} \\ 0 \\ 0 \\ 0 \\ -- \end{pmatrix} \xrightarrow{\varphi_{\text{TomOnly}}} \begin{pmatrix} - \\ ++ \\ - \\ 0 \\ 0 \\ \boxed{--} \end{pmatrix} \xrightarrow{-\varphi_{\text{NineOnly}}} \begin{pmatrix} \boxed{-} \\ ++ \\ - \\ + \\ + \\ -- \end{pmatrix} \xrightarrow{-\frac{1}{2}\cdot\varphi_{\text{TimThree}}} \begin{pmatrix} - \\ ++ \\ - \\ - \\ + \\ -- \end{pmatrix}.$$

The final result that can be denoted ϕ_{Tom} is the complete (unpolished) version of knowledge about Tom; it seems that he probably is six years old. The inference steps seem somewhat heuristic – but who would say that declarative knowledge is easy to use!

Finally, starting from a query "What is known about John?", the following sequence might be found:

$$\begin{pmatrix} 0 \\ 0 \\ \boxed{++} \\ 0 \\ 0 \\ 0 \end{pmatrix} \xrightarrow{\varphi_{\text{JohnOnly}}} \begin{pmatrix} - \\ \boxed{-} \\ ++ \\ 0 \\ 0 \\ 0 \end{pmatrix} \xrightarrow{-\frac{1}{2}\cdot\varphi_{\text{TomNotNine}}} \begin{pmatrix} - \\ - \\ ++ \\ 0 \\ 0 \\ + \end{pmatrix}.$$

Further steps could be taken to fill in the missing entries in the above vector, but the goal has already been reached: The question can be answered – John probably is nine years old.

Applying the given pieces of knowledge, three completed points in feature space can (realistically) be determined in this case – one for Tim, one for Tom, and one for John. Updating the set of long-term memory contents is a challenging task: Invalid contents easily result in “computerized schizophrenia”!

4.2 Knowledge Emerging

Assuming that, for example, the GGHA algorithm [6] is used to search for the center (average) of the cluster of the feature data; the extracted new observation vector could look something like

$$\phi_{\text{Boy}} = \begin{pmatrix} + \\ + \\ + \\ + \\ + \\ + \end{pmatrix}.$$

This can be interpreted as the “boy” category prototype (the “latent vectors” can be defined explicitly – see [7] – or they can be found statistically – see [8]). Further, there is some relationship between the ages three, six, and nine; consider the following (fuzzy) knowledge structure (something like it can emerge when automatic feature extraction [6] is applied; or it can be given directly as input from outside):

$$\varphi_{\text{Young}} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ ++ \\ 0 \\ -- \end{pmatrix} \quad \text{or, better,} \quad \phi_{\text{Young}} = \begin{pmatrix} + \\ 0 \\ - \\ ++ \\ 0 \\ -- \end{pmatrix}.$$

Note that here the age of six is assumed to be “neutral”, not young but not “not-young” either. The entries “Tim”, etc., also become involved in the adaptation process, because of their correlations with the age feature entries. This vector can implement a *third-level* knowledge structure, or an “age-axis” around feature cluster centers. Assume that the same utterance as above (“Tim is three”) is given now – in a constructivist way, the internal representation may become very different this time:

$$\varphi_{\text{TimThree}} = \begin{pmatrix} \boxed{++} \\ 0 \\ 0 \\ \boxed{++} \\ 0 \\ 0 \end{pmatrix} \cong \begin{pmatrix} \boxed{++} \\ + \\ 0 \\ \boxed{++} \\ + \\ 0 \end{pmatrix} \approx 1 \cdot \begin{pmatrix} + \\ + \\ + \\ + \\ + \\ + \end{pmatrix} + \frac{1}{2} \cdot \begin{pmatrix} + \\ 0 \\ - \\ ++ \\ 0 \\ -- \end{pmatrix} = 1 \cdot \phi_{\text{Boy}} + \tfrac{1}{2} \cdot \phi_{\text{Young}}.$$

The above should be interpreted so that the focus entries can be matched by adding the prototypes of “boy” and “young” together … a “rather young boy” is also perceived. This reconstruction is carried out in the Step 4 of the Algorithm. Note that in the previous phases the associative matching was trivial, when only one vector was matched at a time, but the number of inference steps became larger; now, on the other hand, there is just one step, but matching process is more complicated, more component vectors being involved.

5. CONCLUSION

Above, a compact framework was constructed that seems to address some painstaking problems of current AI research and cognitive science from a new, seemingly fruitful standpoint. There are various shortcomings; but if the reader wants to absolutely disagree with the above discussion, here are some comments already available:

1. “We cannot study the mental phenomena using the same mental machinery”. This complaint is not valid; mathematics is a neutral language free of Wittgensteinian biases: If an algorithm works, it works regardless of the interpretation of the data structures.

2. “It is only a matter of computing power; the current approaches will work fine when we have faster hardware”. This is not true either; it has been approximated that the computing power of the current computers is at the level of a simple fly – but still the fly is qualitatively on a much higher level of intelligence.

3. “We just cannot attack these questions using any form of analysis”. This may, of course, finally turn out to be true. However, hopefully the only motivation for this opinion is not the all too common mental laziness … you know, in the late 1800’s they thought that “there just is so much coal in the Sun”!

REFERENCES

[1] Anderson, J.R.: *The Architecture of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1983.

[2] Chase, W.G. and Ericsson, K.A.: Skill and working memory. In *The Psychology of Learning and Motivation* (ed. G.H. Bower), Academic Press, New York, 1982.

[3] Elio, R. and Scharf, P.B.: Modeling novice-to-expert shifts in problem-solving strategy and knowledge organization. *Cognitive Science*, **14**, 1990, pp. 579-639.

[4] Hyötyniemi, H.: On Mental Images and “Computational Semantics”. *Proceedings of the 8th Finnish Artificial Intelligence Conference STeP'98* (eds. Koikkalainen, P. and Puuronen, S.), Publications of the Finnish Artificial Intelligence Society (FAIS), Helsinki, 1998, pp. 199-208.

[5] Hyötyniemi, H. and Saariluoma, P.: Chess – Beyond the Rules. In *Games, Computers and People* (ed. Honkela, T.), Finnish Artificial Intelligence Society, 1999, pp. 100-112.

[6] Hyötyniemi, H.: *GGHA Toolbox for Matlab*. Helsinki University of Technology, Control Engineering Laboratory, Report 115, 1999.

[7] Hyötyniemi, H.: Explorations in “Naturalistic Formalisms”. In the *Proceedings of STeP 2000* (this volume).

[8] Hyötyniemi, H. and Hasu, V.: Independent Components of Expertise. In the *Proceedings of STeP 2000* (this volume).

[9] Jackson, S.A.: *Connectionism and Meaning: From Truth Conditions to Weight Representations*. Ablex Publishing, New Jersey, 1996.

[10] Kosslyn, S.M.: *Image and Mind*. Harvard University Press, Cambridge, Massachusetts, 1980.

[11] Laird, J.E., Rosenbloom, P.S., and Newell, A.: Chunking in SOAR: The anatomy of a general learning mechanism. *Machine Learning*, **1**, pp. 11-46.

[12] Olshausen, B.A. and Field, D.J.: Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Vision Research*, **37**, 1997, pp. 3311-3325.

[13] Pylyshyn, Z.: The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, **88**, 1981, pp. 16-45.

[14] Rosch, E.: Principles of categorization. In *Cognition and Categorization* (eds. Rosch, E. and Lloyd, B.B.), Erlbaum, Hillsdale, New Jersey, 1978.