MODELLING AND SIMULATION METHODOLOGY
Knowledge Systems' Paradigms
M.S. Elzas, T.I. Ören, and B.P. Zeigler (Editors)
Elsevier Science Publishers B.V. (North-Holland), 1989

CHAPTER I.5

AN ADAPTIVE, SELFMODIFYING, NON GOAL DIRECTED MODELLING METHODOLOGY

P.Hogeweg & B.Hesper

BIOINFORMATICA
Padualaan 8
3584CH Utrecht
the Netherlands

Structure oriented models do not use (and do not need) apriori and externally specified goals in order to show adaptive selfmodifying or learning abilities. They employ the structure of the model behaviour to generate local structure enhancement. Such structure oriented, non goal directed methods, like non supervised pattern analysis methods (e.g. cluster analysis) seem often to be more powerful than their goal oriented counterparts even in those cases were apriori goals seem to be available. They certainly are, if no well formulated goals are available prior to the research under consideration, which is the more common case.

1. INTRODUCTION

Modelling and simulation methodology moves in the direction of incorporating more and more of the process of 'modelling in the large' in the formalised part of modelling. Notably methodologies exist for automated model generation, manipulation of knowledge bases used in and generated by models, and 'multifacetted modelling'. At first sight "goal-directed models" with variable structure, adaptive, selfmodifying or learning ability seem to be an important step in this direction. In our opinion such a methodology which presupposes preconceived goals can only contribute marginally to the 'goal' of formalising modelling in the large because the purpose of modelling is largely to discover, sift out and elucidate pursueable purposes. Moreover working with preconceived goals may harm the attempted modelling area.
In this paper we argue that variable structure models with adaptive selfmodifying or learning ability do indeed increase the power of our modelling formalisms considerably, but that it is a mistake to conceive such models as being primarily goal directed.
First we give our apriori reasons for considering non-goaloriented models (section 2), next we discuss variable structure models (section 3) and learning systems (section 4) and finally we combine these lines of our argument and show the feasibility and power of structure oriented models by elaborating on the theory of individual oriented variable structure modelling which we have developed over the past several years as MIRROR modelling methodology (Hogeweg & Hesper 1979,1981A,B, 1983,1985A,B,1986).

2. ON THE NEED OF NON-GOALDIRECTED MODELS

There are several types of reasons why we need non goaloriented models.

1. Whence do we get our goals?
The answer to this question is "by modelling". In particular modelling provides us with concepts and potential relationships between concepts. A classical example in case is the search for black holes. Likewise without a model no one would contemplate to develop a model which uses social structure as timing device (see Hogeweg & Hesper 1985).
Moreover in many cases it is much easier to formulate an interaction structure which could represent an object to be modelled than to formulate the behaviour of that object (e.g. Lotka-Volterra positive and negative interactions vs its dynamics of neutral oscillations (i.e. undamped oscillations with an amplitudo dependent on the initial state).

2. Presupposing well formulated goals is begging the question of modelling.
The simplistic image of modelling, i.e.: "Given a set of input/output relations find a model to describe them" has the form of strictly goal directed modelling and has of course a trivial solution: the list of input/output pairs itself. Such a model is 100% correct for the well formulated goal and 0% for the implied but not formulated goal: other input/output pairs.
Also in the much more sophisticated model methodology of Zeigler, Klir, Wymore and Elzas (see Elzas 1984 for a review) an actual simulation model is shaped relative to an "experimental frame", "observation channel" or "set of feasible test items", and such output (goal) oriented models tend to generate relatively little new knowledge because the model structure is expressed in terms of the same concepts as the model results. This is the easiest way to achieve a goal. In contrast a very important contribution of variable structure modelling methodology is the possibility of studying truly emergent properties of the model, i.e. properties which can only be expressed in terms of novel concepts, not used in the model formulation (Hogeweg & Hesper 1986, see also section 3).

3. Pattern recognition methodology demonstrates the role of non goal directed methods.
Methods which seek a structure in a dataset without defining classes a priori, are called non-supervised methods in pattern recognition. Even in cases that apriori defined groups are sought such (simple) non-supervised methods often prove to be more powerful than sophisticated supervised methods (e.g. if the predefined classes are subdivided in several noncontiguous 'clusters' in state space). Moreover, the newly defined concepts which reflect the structure of the data can often be related to additional (partial) knowledge about the data which is not included in the analysis.

4. Selfmodifying models do not need goals.
There is an existence proof that self-modifying, adaptive self-structuring systems do not need goals: the biosphere (at least as perceived by biologists!). Moreover the structure of science seems to fit better in this description than one in terms of predefined goals.
A selfstructuring modelling process does not imply randomness in

the resulting knowledge aquisition because the knowledge seeking path will be generated by the structure of the environment which supports the modelling: in science ultimately the structure of the universe. Only if the modelling environment is unstructured or is overhauled by the bulldozers of preconceived goals the knowledge seeking paths may lead to naught.

3. SYSTEM-ORIENTED VS INDIVIDUAL-ORIENTED VARIABLE STRUCTURE MODELS

Variable structure models are recognised as a class of models by Oren (1975,1979,1985). They are models in which the model specification at the level of coupling of systems (Zeigler 1976, Klir 1979) is not fixed but variable. Therefore such models must be specified at a higher level comparable to Klir's metasystem level (Klir 1979). The system behaviour consists of both changes in the variables (like any fixed structure model) and in changes of the components and their interconnections. This definition of variable structure systems is rather broad and includes e.g. neural networks in which the strength of the interconnections is changed during the simulation and cellular growth models (e.g. L-systems (Lindenmayer 1968) and their non-synchronous counterparts (Hogeweg 1980) in which the set of component systems (i.e. cells) changes (increases) during the simulation but in which the coupling remains as similar as is compatible with the varying components. Arguing that neural networks and cellular growth models do not need a novel methodology, Zeigler (1986) tries to narrow down the definition of variable structure models by distinguishing a class of "non-trivial" variable structure models which excludes the above mentioned cases. He requires for this set of non-trivial variable structure systems that the model specification distinguishes between"ordinary behaviour" (i.e. the change of variables) and "structural behaviour" (the change of components and couplings between components) and that phases of ordinary behaviour are punctuated by periods (of non-zero durations) in which the structure of the models is changed and the behaviour state of the former structure is linked to the behaviour state of the novel structure. Thus this class of variable structure systems correspond to a sequence of fixed structure systems and the transitions between them.

In contrast Hogeweg & Hesper (1979,1981a,b,1983,1985a,b,1986) have developed a class of variable structure models in which fixed structure is an (interesting) emergent property of the system which may occur locally (i.e. usually involves only part of the system) and may last for some time but can dissolve again. These systems are in several ways more truly a variable structure than the ones considered by Zeigler, in particular because:

1. Variable structure is the normal mode, fixed structure is an interesting "side effect"

2. The couplings between the (variable) set of components changes in a less restricted way: apart from the "ancestry based" inheritance of couplings which Zeigler uses (and also occurs in the above mentioned cellular growth models), and precoded couplings (as occur in many rule based "expert" systems) we also use a "neighbourhood" based coupling. Such neighborhood based coupling is much less predetermined than the "ancestry based" and "global

receptor based" coupling (compare even e.g. Beelers worm whose behaviour is shaped by meeting traces of its previous behaviour: this may lead after an arbitrary long time to death).

3. The behaviour of such systems is not studied in terms of predefined changes in variables and predefined structure changes but in terms of novel concepts which emerge during the simulation.

4. Such systems need novel methodology, not only for the specification of the system and consistency checking (as emphazised by Zeigler) but in particular for the observation of the system.

The differences between these two classes of variable structure models which go beyond the (too) simple variable structure models mentioned above, is rather profound. Zeiglers approach is essentially a global top-down approach. The specification focusses on the system as a whole. Component systems are defined as parts of the whole (compare the pruning of the entity structure to obtain a model in Zeigler's multifaccetted modelling methodology). In contrast our MIRROR models aim at a maximally local definition of entities to make a truly bottom up approach feasible. Thus the component systems are defined as autonomous entities (individuals) whose behaviour is dependent on their local environment. They actively extract information from their environment and are therefore viable in a large collection of environments (consisting of acollection of individuals), i.e. they are not defined as parts of a whole. This autonomy of the component systems is crucial in the context of bioinformatic models (i.e. models for studying informatic processes in biotic systems) because in that case we surely do not want to impose externally the "organisation of the system" but want to understand the organisation as an emerging property of the collection of individuals under consideration. This organisation can have the form of an emerging fixed structure but usually such fixed structures contain only part of the system (and last only part of the simulation period). Thus, viewed from the "system as a whole" viewpoint the systems consist of a sea of changing structures in which fixed structures take shape and dissolve again.

To distinguish the two approaches we will designate them global or system oriented variable structure systems and local or individual oriented variable structure systems: we do not think that non-triviality is a distinguishing property of the former approach.

4. SELFMODIFYING AND LEARNING SYSTEMS

In this section we review some of the concepts and distinctions which have emerged from the research on learning systems.

4.1 Degree of prestructuring of input

In classical learning systems (e.g. the classical neural network learning models mentioned above) the observational universe of the learning system is rigidly prestructured, i.e. a predefined set of objects is observed through a predetermined set of features (which are therefore known to be applicable to all objects). Thus the objects are known to the system as a vector of feature values, and no problems exist with respect to matching features (the homology of features). More recently, learning systems are studied with less prestructured input (e.g. Winston 1975,1982,1986, Michalski 1986a,b, Sankoff & Kruskal (1983), Hogeweg & Hesper 1984): the homology of features is not given but should be found by the system, and relevant features should be sought. This is done by incorporating the semantics of the features into the system and by considering the structuring context of the features. Note that features are only a weaker form of apriori structure definition; the point is, however that it is weaker. Learning systems operating in the "real world" (or in a model world) neccessarily are at least of the second type.

4.2 Supervised vs nonsupervised learning

This distinction, also labeled with the terms learning s.s. vs discovery or learning from observation reoccurs in all subfields of learning and pattern analysis. In supervised methods there is a prior definition of what is to be learned by the system (there is a teacher who knows better), whereas no such priori definition exists in the latter case. Supervised learning very often takes the form of learning by examples (and counterexamples) from these examples (i.e. from this "paradigmatic definition of the concept 'class') the system is to generate a description of the class and/or a decision rule for the membership of the class (i.e. generate an intensive definition of the concept/class).
In contrast non supervised learning methods generate concepts/classes which are "interesting" for the information under consideration. Interesting because the available information can be conveniently expressed in terms of the generated classes/concepts or because the classes/concepts are useful as building blocks for higher level concepts.

The feasibility of both approaches depends on the structure of primary information. Non supervised methods degenerate if

1. The primary information is unstructured: any class/concept then is as good (i.e. bad) as any other. Pleas for adding goals to non supervised methods usually use such unstructured examples (e.g. Stepp & Michalski 1986). Such examples remain however useless whatever is done to them.

2. Several conflicting structures are present in the data. In that case most methods do not expose the conflicting structures but come up with some compromise which renders the generated concepts uninterpretable. This is harmful if the concepts are to be used as building blocks for higher level concepts. Better interpretability can be achieved by:
(a) pattern enhancing techniques which filter out one of the structures (e.g. by iterative character weighting, Hogeweg 1976)
adding constraints to the concepts to be formed (e.g. monothetic clustering techniques (e.g. Williams & Lambert 1960), Constrained clustering (austin 19..), Conceptual clustering (Michalski & Stepp 1984) or adjusting the concept definition so as to allow for easy recognition (Oligothetic characterisation of clusters (Hogeweg & Hesper 1981c))

Like the non supervised approach, the supervised approach may degenerate into a useless exercise in several ways:

1. To learn well defined concepts is useless: why not use the definition.

2. Correct recognition of the previously given examples is always possible by matching to the (completely stored) examples

3. Generalisation beyond the given examples is only possible if the concepts to be learned are 'nicely' represented in the data and if no conflicting structure prevents its recognition. For example if a concept is represented in the data as two clusters separated by the cluster representing the concepts to be distinguished from it, most supervised methods fail, (whereas nonsupervised methods may come up with 3 concepts (corresponding to the clusters) which can easily be combined into the two desired concepts). Also, a priori defined concepts which are just a little askew with the data available are hard to recognise while the skewness may arise from the not well definedness of the concepts (compare 1). Also in this case non-supervised methods may help by coming up with a more useful definition of the concept.

Thus we conclude that both types of learning are feasible only if the data are well structured. In addition, supervised methods are only useful if the supervised concepts match the structure of the data, which is usually only the case if the concepts are derived (learned) from observation of the data (i.e. by (implicit) non-supervised learning). Although the concepts some people want to teach their machines may be implicitly derived in this manner (e.g. 'arches' or 'dogs') they cannot be if the data are never observed before, as is the case in novel situations, e.g. those created by model universes.

4.3 Engineering vs modelling viewpoint of learning

From the early days (in the fifties) of machine learning onwards, many of the learning systems have had a dual interpretation: that of a model of a (biotic) learning system and that of a tool to perform a learning task. Examples of such model/tool interpretation of identical (or very similar) systems are e.g. Perceptrons (Rosenblatt 1958, vs machines using discriminant functions (Nillson 1965); Neural networks vs Distributed parallel systems (see e.g. Rumelhardt 1986, Martin 1986 but also much older literature); Evolutionary models vs Genetic algorithms (cf Holland 1975, 1986); Human learning vs adaptive expert systems (cf Simon 1984, Michalski 1986, Rosenbloom & Newell 1986, Anderson 1986, Davis & Lenat 1982, Lenat 1984).

All learning models fall in the class of variable structure models: without a changing structure they would not learn. Some fall in the above mentioned class of "trivial" variable structure models (e.g. neural networks) in which only parameters (strength of couplings) are changed. Others fall in the class of individual oriented or object oriented variable structure models (object oriented models are similar, but less local, than the individual oriented variable structure models discussed above) (e.g. genetic algorithms, Actor (message passing) systems (Hewitt 1977, Lenat 1975), Transfer-frame based and Censor-based learning (Winston 1978,1986), Chunking and Knowledge compilation (Rosenbloom and Newell 1986, Anderson 1986), Society of mind models (Minsky 1987) etc.).

Not only are learning models variable structure models, the converse is true as well: variable structure models can virtually always be interpreted as learning systems (see below).

4.4. Learning as adaptive behaviour vs learning as building and manipulating representations of experience.

Any system which changes its behaviour dependent on the environment in which it finds itself can be considered as learning or "adaptive" system. If no apriori definition is given of what is to be learned (i.e. in a non-supervised system, e.g. the biosphere, mathematics (see Davis & Lenat 1982)) any behaviour change is almost bound to be equivalent to learning (or optimising) something. Although this fact is sometimes cited as a criticism of evolution theory, it can better be interpreted as a very powerful fact about variable structure systems. Nevertheless it is sometimes useful to employ the more restricted definition of learning proposed by McCarthy (1968): "Learning is building and manipulating representations of experience". Learning in the latter sense increases considerably the power of variable structure models (which automatically are learning systems in the former sense).

5. SELFMODIFYING VARIABLE STRUCTURE MODELS

From the previous lines of argument we conclude that a valuable class of models has the following properties:

1. No global goal is defined apriori. Instead an apriori structure definition generates locally pursuable purposes

2. the structure definition is of the type of an individual oriented, local, truly variable structure one.

3. Learning is an integral part of the model and includes selfmodification, adaptation and representation generation and manipulation. The prestructuring of the data needed for interesting learning is generated by the model, the emphasis is on non-supervised learning methods although supervision for learning can be generated by the model as well.

4. The variable structure and learning capabilities of the models serve both the model itself (s.s.) and the modelling methodology.

MIRROR modelling is our attempt to develop a such a methodology. Within this framework we have previously shown that:

1. We need individual oriented models if the models are to reflect the informatic structure of the system studied (Hogeweg & Hesper 1979,1983,1985); individual oriented models generally have a variable structure.
2. Ww need variable structure models to model variable structure systems (Hogeweg & Hesper 1979,1983,1986)
3. Surprisingly simple individual oriented variable structure models can generate surprisingly complex, selfregulating and 'multifacetted' behaviour (Hogeweg & Hesper all references, see especially 1983,1985,1986)
4. We need sophisticated knowledge seeking processes in such models (Hogeweg & Hesper 1981,1986)
5. Selfstructuring through building and manipulating representations of experience is feasible and desirable (MICMAC modelling) (Hogeweg & Hesper 1979,1981)

In this paper we go beyond this previous work in showing that structure recognising and structure enhancing processes within such systems provide a means to include a still larger part of modelling in the large in the formalised part of modelling, and forms a (better) alternative to "goal oriented modelling". In particular we discuss the recognition of locally fixed structure by DWARFs (section 6) and its use in gradually restructuring models (section 7). Finally we discuss how different modelling paradigms (i.e. "quantitative" models vs "qualitative" models (i.e. models in the qualitative reasoning paradigm of AI) and classification vs modelling s.s. are unified through these concepts.

6. HOARDING STRUCTURE IN MIRROR MODELS

6.1 DWARFs, JEWELs and GEMs

The basic structure recognising entities in MIRROR worlds are DWARFs. DWARFs simply search for invariant relations by keeping track of activities and revivals of other MIRROR entities (possibly other DWARFs). REVIVALS represent the "who reacts on whom" structure of the system. Activities include procedure calls and changes in variables

If a DWARF finds an invariance (i.e. the tracked entity is repeatedly revived by the same other entity, calls repeatedly a procedure with a same parameter or changes a variable to the same value) a JEWEL is placed in its HOARD.

A HOARD is shared by DWARFs with similar concern, e.g. because they track the same entity. Placing a JEWEL in a HOARD causes other DWARFs to reexamine their memory tracks using the JEWEL as a tool. The DWARFs revived by the addition of a JEWEL to a HOARD are all those which:

1. are concerned with the same individual (looking at other properties than the DWARF who made the JEWEL).
2. are concerned with the entities (individuals, procedures, variables) which occur in the JEWEL as values of the JEWEL. For instance a JEWEL representing the fact that individual A was repeatedly revived by the activity P of individual B will revive the DWARF tracking who revives B and the DWARF tracking the activity P of B.

The DWARFs revived by a JEWEL use it to focus on:

1. the time period indicated by the new JEWEL. Being able to focus on a time period makes it feasible to detect a wider class of patterns than just invariance as the DWARF does without any focus. For example they look wether a value changes monotonically in the period, or wether it takes only a few values.

2. The 'interesting events' (Hogeweg & Hesper 1981) punctuating the invariance. The emergence of an invariance and especially the end of an invariance are interesting and the events happening at those times are important in a higher level description of the system. Therefore the revived DWARFs check wether their tracks took a turn at the beginning or end of the invariance represented by the JEWEL, i.e. whether the entities tracked by them participated in the interesting event. Such an analysis of the interesting events is added as GEMs to the JEWEL.

3. the entities involved: did they occur in their tracks, e.g. as the value of a parameter or variable?
Any structure found in this way is added to the HOARD (and linked to the JEWEL which was used as tool). The addition of these new JEWELs again revive DWARFs.

An interesting special case is if DWARFs in 'jeweling mode' are revived by JEWELs, i.e. DWARFs who detected an invariance which remains valid during the interesting event terminating the invariance detected by the reviving JEWEL. The two JEWELs are likely to represent conflicting structures of which only one survives the interesting event. This is noted and a new DWARF tracking potential conflict is generated: it tracks the DWARFs.

6.2 Example of a functioning HOARD

As a simple example consider a HOARD in the SKINNY universe (cf Hogeweg & Hesper 1985,1986; SKINNIES are social individuals who interact by a dominance interaction (DODOM) and who know each other personally, i.e. they build a representation of the other individuals in their SKINSPACE (mental SPACE)). DWARFs track e.g. DODOM interactions and look for invariance of the DODOM parameter (the one which receives the DODOM behaviour). Such invariances are readily detected: pairs of SKINNIES tend to split off the group. Surprisingly such invariances also exist in the larger group but in that case they are mostly not symmetric: SKINNY A only interacts with B but B interacts also with other SKINNIES. This discrepancy is noted because the JEWELs of A revive the DWARFs tracking DODOM of B. The latter DWARF notes that B interacts with several (but a small subset of all) SKINNIES and links a JEWEL to the first with this less stringent invariance.

After several such augmented JEWELs are formed an interesting question (adressed by OBSERVERs, see Hogeweg & Hesper 1986) becomes what are the differences between SKINNY A and B? In this case the DWARFs supply the set of objects and the supervision for a pattern analysis (learning) task of the OBSERVERs. It appears that SKINNY B is always the more dominant one (Both in the mental space of A and in that of B).

6.3 Multilevel modelling and HOARDs

The resulting HOARDs can be seen as representing the behaviour of the system in a way similar to the apriori definition of the system in system oriented variable structure modelling. Augmented JEWELs represent (parts of) the system in 'ordinary' (= (semi) fixed structure) mode and their behaviour pattern. The GEMs, representing the interesting events punctuating JEWELs, represent the systems structural behaviour pattern. However the important differences are:

1. This description is a result of the modelling, not its apriori definition;

2. It is a local description which involves (changing) parts of the entire system;

3. It is a multilevel representation: structural behaviour at one level (represented by one set of JEWELs) is 'ordinary behaviour' at an other level which has recognised the pattern in interesting events.

7. JEWELS AS TOOLS

DWARFs hoard JEWELs which are used by DWARFs as tools to make and hoard more JEWELs: Jeweling is both means and end for DWARFs. JEWELs can also be used by other entities of the MIRROR world, for example by the OBSERVERs mentioned above (obviously JEWELs are nice to observe!). Moreover the JEWELs can also be used to change the universe. Changing the universe can mean:
1) the (macro) behaviour remains the same but the mechanisms which generate this behaviour change: i.e. the informatic structure generating the behaviour is changed (optimised): MICMAC modelling.
2) the informatic structure of the model remains the same (similar) but the macro behaviour is changed (e.g. by a change of the parameters or a local change in the interaction pattern): structure enhancement by WIZARDs.

7.1. MICMAC modelling

The MICMAC modelling principle has been developed by us several years ago (Hogeweg & Hesper 1979,1981). The basic ideas behind this methodology are: (1) only 'interesting' events should be represented in the model; (2) Interesting events are those events which redirect a (sub) process from its "expected" course of behaviour; (3) The purpose of modelling is to form and refine expectations about (sub) processes; (4) such expectations can be formed by the model itself; (5) these expectations can be used by the model to select interesting events; (6) in this way, piecemeal, the model can restructure itself in a multilevel, heterarchical manner; (7) The original definition of the system should be retained so that the system remains "open" to unexpected events, which it handles in the basic lowest level definition.

Thus MICMAC modelling aims at generating new higer level definitions of the model without changing its (original) behaviour. JEWELs and GEMs provide this higher level definition. The (partial) invariance of the JEWELs can be exploited: a much simpler model often suffices to generate the variable changes during this period (much remains invariant and the changes which do occur can often be 'predicted' by mimicking previously hoarded JEWELs (Hogeweg and Hesper 1981).

However, in order to use this shortened behaviour definition the occurrence of a state which will generate a JEWEL similar to the already hoarded JEWELs should be recognised. Thus, the cost of simple large-scale behaviour rules is paid in terms of classification ability. This classification ability can be obtained using simple learning schemes because the learning problem is prestructured by the DWARFs: they define the set of instances to be classified as well as the features which can be used for the classification. Athough the JEWELs can also provide a useable supervision for this learning we find it useful to use nonsupervised methods on the set of systems states selected by the JEWELs: although JEWELs may be superficially similar they can arise from different circumstances: recognising these differences even when the results appear to be similar enhances the reliability of the classification.

7.2. Structure enhancement by WIZARDs

Like in MICMAC modelling the system is restructured piecemeal. However, unlike MICMAC modelling the behaviour of the system is not supposed to remain identical to the original system, but should enhance the patterns occurring in the original system. Moreover, the restructuring occurs at the level of (or below) the level of the the original definition of the system, using, however the higher level representations of the system generated by the DWARFs.

Structure enhancement is the work of WIZARDs; WIZARDs try to change the original definition of the system so that the structure represented by the JEWELs will occur more frequently in the system. However, in doing so they are not supposed to enhance the information processing capabilities of the original entities in a goal directed manner. For example it is not allowed to tell a SKINNY to stick to a SKINNY it has interacted with during some time: such a redefinition would annihilate the interestingness of the studied phenomenon and structure enhancement would degenerate to MICMAC modelling without error checking. Moreover WIZARDs are often not supposed to differentiate at the definitional level between entities which were originally defined identically: the differentiation of entities is to remain a purely 'environmental' phenomenon. Therefore WIZARDs, unlike DWARFs and MICMAC modelling, should look beyond the local structure of the model: the changes they make affect the entire universe.

In a very preliminary implementation, a WIZARD, having noticed that pairswapping between SKINNIEs (who usually stick to their original partner) occurred mainly when (a) each pair had one very low ranked individual and (b) the original partners lost sight of each other, lowered the stepsize of lowranked SKINNIES which resulted in an increased pair formation and an increased fidelity of the pairs.

A world change effectuated by a WIZARD generates a rival WIZARD: one which wants to undo the change and even wants to go beyond that by affectuating the opposite change. The rival WIZARD of the above metioned WIZARD increased the walking speed of low ranked individuals, which results in quite a different structure: meetings of pairs almost always result in a partner exchange such that the SKINNYs which are closest in rank remain together.

Thus WIZARDs explore possible universes. How their conflicts should be resolved is not yet known: it looks as if they can only fight it out (as has been recounted in many tales).

8. DISCUSSION

Classical mathematical and simulation models can be represented as a fixed network of interacting entities, and the next state function of the system is maximally generalised, i.e. it is defined as much as possible without reference to the particular state in which the system occurs. This is especially true in mathematical models and simulation models closely linked to mathematical formalisms (e.g. continous systems models). In event and process (object) oriented models next state functions can be specified relative to the state (change) of the system, but, no elaborate state recognition capabilities are used because the relevant states/events are fixed.

In contrast, much "naive" (but very powerful!) modelling depends largely on the correct recognition and classification of the state of the system. For each class of states a different next state function is used which is largely derived from experience. Clancy (1988) has correctly recognised that expert systems fall into this class of classification dependent models. Moreover, Hardt (1988) demonstrated the power and limitations of a classification oriented approach (called "qualitative reasoning" in AI) by comparing classical diffusion equations (which are general but hard to solve) and their qualitative counterparts. Very simple qualitative rules (whose applicability depends however on the correct recognition of subproblems) suffice to derive estimates of the time it takes of a substance to diffuse from A to B. However, trying to convince the audience of the validity of her qualitative rules she took resort to the diffusion equations and to "it turns out that".

In their representation of the aquisisition of knowledge about physical systems Forbus & Gentner (1986) recognise four stages, i.e. "naive physics, causal relations, process models, and mathematical theory. The earlier stages are knowledge rich, the later stages progressively incorporate less knowledge. This sequence corresponds to moving from more classification oriented models to general nextstate function models, and this seems indeed the way science has progressed. Indeed simplicity (Occams razor, cf e.g. Russel 1946) has always been an important principle in science and has indeed been interpreted as minimising a.o.the state-space knowledge in the model formulation. Simple models which incorporate little explicit knowledge are so powerful because their simplicity ensure the relevance of the results, whereas knowledge oriented models are only relevant relative to previouly observed behaviour or relative to a prior goal setting.

The classical conception of simplicity often implies homogeneity, a small number of variables and a fixed structure. This concept of simplicity leads to a global desciption.
However, simplicity is not a fixed concept. In last analysis simple is what we conceive of as simple, and this depends on the tools we have.
The tools we have today invite a new concept simplicity: one which involves local, simple infomation processing entities with varying interactions, forming varying patterns. Models which are very simple in this sense can generate very complex macro structures. In order to perceive these macro structures we need knowledge-rich classification type qualitative models of them. Thus, the course of science seems about to turn around 180 degrees: Now that we can formulate truly simple informatic models we need to model them in a manner we originally modelled the universe, i.e. in terms of "naive", knowledge rich, classification oriented, large scale models. MICMAC modelling does exactly this: given a small scale, local simple definition of a universe, it generates models, at multiple levels, which are progressively more knowledge (classification) oriented. Again, these models can be validated in terms of "it turns out that": However, now there are many universes which can be compared.

Goal oriented modelling considers the structure of the universe as a constraint for the attainment of the preconceived goal. In contrast, in structure oriented modelling the structure of the universe is used to recognise what is interesting to know and learn. In fact without a structured universe knowing and learning is not well defined. Only by creating many alternative possible universes (e.g. MIRROR worlds) and traveling the paths which are opened up by them, can we enhance what can be known and pursued.

REFERENCES

Anderson, J.R. (1986). Knowledge Compilation: The general Learning Mechanism. In: Machine Learning II, An Artificial Intelligence Approach. Michalski, R.S., Carbonell, J.G. and Mitchell. T.M. (Eds.), Morgan Kauffman Publ in pp. 289-310.

Clancy, W.J. (1988) Viewing Knowledge bases as Qualitative models. Chapter I.4 this volume.

Davis, R. and D.B. Lenat (1982). Knowledge based Systems in Artificial Intelligence, Mc Graw Hill Publ., 490 p.

Elzas, M.S. (1984).System Paradigms as reality mappings. In: T.I. Oren et al (Eds) Simulation and model based methodologies, An integrative view, Springer Verlag.

Forbus, K.D. and D. Gentner (1986).Learning Physical Domains: Toward a Theoretical Framework. In: Machine Learning II, Michalski, R.S. (ed) op cit.).

Hardt, S.L. (1988).Qualitative Reasoning and Simulation for Knowledge intensive Explanations and Predictions. Chapter I.7 this volume.

Hewitt, C.E. (1977). Viewing control structures as patterns of passing messages. Journal of Artificial Intelligence 8 pp. 323-364.

Hogeweg, P. (1976). Iterative Character Weighing in Numerical Taxonomy. Comput Biol Med 6, pp. 199-211.

Hogeweg, P. and B. Hesper (1979). Heterarchical Selfstructuring Simulation Systems: Concepts and Application in Biology. In: Methodology of Systems modelling and Simulation, B.P. Zeigler, M.S. Elzas, G.J. Klir, T.I. Oren (Eds),North Holland, pp 221-232.

Hogeweg, P. (1980). Locally synchronised developmental systems, conceptual advantages of the discrete event formalism. Int. J. General Systems 6, pp. 57-73.

Hogeweg, P. and B. Hesper (1981a). Two predators and a prey in a patchy environment: An application of MICMAC modelling. J. Theor Biol 93, pp. 411-432.

Hogeweg, P. and B. Hesper (1981b). On the role of OBSERVERs in large scale systems. In: UKSC Conference on computer simulation, Westbury House, Harrogate, pp. 420-425.

Hogeweg, P. and B. Hesper (1981c). Oligothetic Recognition of Clusters. Pattern Recognition 14, pp. 131-136.

Hogeweg, P. and B. Hesper (1983). The ontogeny of the interaction structure in Bumble Bee colonies: a MIRROR model. Behav. Ecol Sociobiol 12 pp. 271-283.

Hogeweg, P. and B. Hesper (1984). The alignment of Sequences and the construction of phyletic trees: An integrated method. J. Mol Evol 20 pp. 175-186.

Hogeweg P. and B. Hesper (1985a). Socioinformatic process, a MIRROR modelling methodology. J. Theor Biol 113, pp. 311-330.

Hogeweg P. and B. Hesper (1985b). Interesting Events and distributed systems. SCS Multiconference 1985 p9 81-87.

Hogeweg, P. and B. Hesper (1986). Knowledge seeking in variable structure models. In: Simulation in the Artificial Intelligence Era M.S. Elzas, T.I. Oren B.P. Zeigler (Eds), North Holland pp. 227-243

Holland, J.H. (1975). Adaptation in natural and artificial systems, Univerity of Michigan Press, Ann Arbor.

Holland, J.H. (1986). Escaping Brittleness: The possibilities of general purpose learning algorithms applied to parallel rule based systems. In: Michalski, R.S., Carbonell, J.G. & Mitchell, T.M. (Eds) Machine Learning, An artificial Intelligence approach Vol 2, Morgan Kauffman, Publ.

Klir, G.J. (1979). General Systems Problem solving methodology. In: Methodology of systems modelling and simulation. Zeigler et al (Eds) op.cit.

Lenat, D.B. (1984). The role of heuristics in learning by discovery: three case studies. In: Machine Learning an Artificial Intelligence Approach, Michalski et al (Eds) op. cit.

Lindenmayer, A. (1968). Mathematical models for cellular interactions in development I and II. J. Theor. Biol. 18, pp. 280-299 & 300-521.

Martin, L. (1986). Use of explicit constraints to evaluate a potentially unique, neural model for single interval concept manipulation. In: Modelling and Simulation methodology in the Artificial Intelligence Era M.S. Elzas et. al (Eds) op. cit.

McCarthy J. (1968). Programs with Common Sense. In: Semantic Information Processing M. Minsky (ed) MIT Press.

Michalski, R.S. and R,E. Stepp (1984). Learning from observation: Conceptual clustering In: Michalski et al. (Eds) op cit.

Michalski, R.S. (1986). Understanding the nature of learning, Issues and research directions. In: Michalski et al. (Eds) op. cit.

Minsky, M. (1986). The Society of Mind. (draft, forthcoming)

Nilson, N.J. (1965). Learning Machines. McGraw Hill

Oren, T.I. (1975). Simulation of time varying systems. Proc Int Congress of Cybernetics and Systems Gordon Breach Oxford, pp. 1229-1238.

Oren, T.I. (1979). Concepts for advanced computer assisted modelling. In: Methodology in Systems modelling and Simulation Zeigler et al (Eds) op cit.

Oren, T.I. (1985). Simulation: Taxonomy. In Encyclopeadia of Systems and Control. Pergamon Press Oxford (in press).

Rosenbloom, P.S. and A. Newell (1986). The Chunking of Goal Hierarchies A Generalised model of practice. In: Machine Learning II, An Artificial Intelligence approach. R.S.Michalski et al. (Eds) op. cit.

Rumelhardt, D.E., J.L. McClelland and the PDP Research Group (1986). Parallel Distributed Processing, Exploration in the Microstructure of Cognition Vol 1, 2, MIT Press, Cambridge.

Rosenblatt, F. (1958). The Percepron: a probabilistic model for information storage and organisation in the Brain. Psychological Review 65 pp. 386-407.

Russel, B. (1946). History of Western Phylosophy. Allen & Unwin.

Sankoff, D. and J.B. Kruskal (1983). Time Warps, String edits and Macromolecules: the theory and practice of sequence comparison. Addison Wesly Publ Comp 382 pp.

Simon, H.A. (1984). Why should machines learn?. In: Michalski et al (eds) op cit.

Stepp III, R.E. and R.S. Michalski (1986). Conceptual clustering: Inventing goal-oriented classifications of structured objects. In: Michalski et al (Eds) op. cit.

Williams, W.T. and J.M. Lambert (1960). Multivariate methods in Plant Ecology. J. Ecol. 48, p. 689.

Winston P.H. (1975). Learning Structural Descriptions from Examples in Psychology of Computer Vision. P.H. Winston (ed) McGRaw Hill New York.

Winston, P.H. (1982). Learning and Reasoning by Analogy. Communications of the ACM Vol 23, no 12, pp. 689-703.

Winston, P.H. (1986). Learning by augmenting rules and accumulating censors. In: Machine Learning II, Michalski et al (Eds) op. cit.

Zeigler, B.P. (1976). Methodology of Systems Modelling and Simulation. Wiley.

Zeigler, B.P. (1986). Toward a Simulation Methodology for variable structure modelling. In: Elzas, M.S. et al (Eds) op cit.

GLOSSARY

DEMON: *Activates its TIE when its TARGET event happens; is generally generated by its TIE (typically a DWELLER, an OBSERVER or a PATCH).*

DODOM: *dominance interaction for socioinformatic models*

DWARFs: *The basic structure recognising entities in MIRROR universes; They hoard JEWELs and GEMs and can make tools (possibly of tools...).*

DWELLER: *Spatially embedded locally defined autonomous entity in MIRROR worlds.*

GEM: *hoarded and shaped by DWARFs; represents period of changing structure which bridges a gap between JEWELs; can be used as 'interesting event'.*

HOARD: *ordered collection of JEWELs and GEMs.*

Interesting Event: *event which changes the structure of the universe, i.e. redirects the universe from its 'expected' course of behaviour (expected on het basis of observed behaviour in the previous fixed structure universe).*

JEWEL: *hoarded and shaped by DWARFs; represents period of fixed structure, punctuated by GEMs.*

MICMAC modelling principle: *model transformation by gradual replacement of MICro entities by MACro entities on the basis of the observed behaviour. (Thus reducing the set of events to 'interesting events').*

MIRROR modelling: *Modelling methodology for creating artificial universes (MIRROR worlds) consisting of autonomous entities with variable interactions.*

OBSERVER: *most versatile of the output generating entities; tries to represent only 'interesting' features of the MIRROR world.*

PATCH: *homogeneous part of a SPACE; can be 'active', i.e. change its state autonomously.*

RECORDER: *simplest type of output generating entity; records variables and events explicitly defined in the model.*

REPORTER: *generates output on some predefined global property of the MIRROR world which is however not explicitly represented in the MIRROR world.*

SKINNY: *a specialisation of a DWELLER used to study spatial and social interactions; has a SKINSPACE to represent its estimate of other SKINNIES.*

SKINSPACE: *SPACE associated with a DWELLER (but not the space in which it dwells) in which a partially independent world of interacting DWELLERs exist; used e.g. as mental space or as parasite space.*

SPACE: *space in which DWELLERs dwell. (Can be 1,2,3,..D Euclidean or graph structure).*

SPACE-SPACE: *SPACE in which the primary model entities dwell, representing 'real' space.*

WIZARD: *entity which may initiate a 'reality change'; it tries to enhance the structure discovered by DWARFs and represented in JEWELs.*