

The Blind's World I

GERD DÖBEN-HENISCH

A Philosophical Experiment on the Way to Digital Consciousness

I INTRODUCTION

When you enter the installation "The BLIND'S WORLD I" (= BW1) the first thing you see are the colorful funny pictographs representing the virtual world of BW1" The association with a computer game will inevitably suggest itself. However, the conclusion that BW1 is "just" a game would be a fallacy.

BW1 is primarily a "philosophical experiment". BW1 is designated to contribute decisively to issues of modern philosophy such as:

- What is consciousness?
- What functions can be discerned within consciousness?
- How do we gain "knowledge of the world"?
- How does language develop?
- What influence do emotions exert?

The notion of BW1 as a "philosophical experiment" represents a marked contrast to the idea of philosophy considering "philosophical cognition" as "self-contained", "based on itself", "independent", "needing no additional aid", with the aim to grasp "the nature of things", uncovering "eternal truths" or "general principles of all recognition",

Furthermore, philosophers usually attribute the idea of the "experiment" as such to the empirical paradigm of cognition; in the past, it was after all the "canonization of the experiment" as a criterion to validate scientific statements which led to the schism between the modern natural sciences and classical philosophy. For many philosophers — and the majority of scientists, for that matter — this "schism" is still an unbridgeable gap.

BW1 holds another provocative idea in store: mention is made of "consciousness". In the modern experimental sciences there is no room for a "consciousness sure of itself", not even in psychology. Even in modern philosophy, the term has been strongly discredited in certain fields — e.g. analytical philosophy.

Understanding BW1 in its philosophical and scientific significance will require a few additional explanations that go beyond the simple description of the BW1 program.

The history of BW1 starts with Alan Matthew TURING.

2 TURING THE VISIONARY

In 1936 Alan Matthew TURING's famous work "On computable numbers with an Application to the Entscheidungsproblem" was published. Not only did he answer — in the negative — the question for decideability in mathematics, which had been explicitly posed by HILBERT in 1928, he also proposed a definition of computable processes which during his lifetime became a household word among mathematicians and logicians under the name TURING MACHINE [TM]. For GÖDEL, the TM was the most satisfactory of all definitions ever given of a "mechanical process" (DAVIS 1965: 72, footnote).

The concept of the Turing Machine and especially of its generalized form, the "Universal Turing Machine" [UTM], is based on the hypothesis that it can represent all computable processes.

With his ideal machine TURING was as much ahead of the technical possibilities of his time as were his practical and philosophical visions inspired by this new concept.

The topics of his considerations, seen by many as provocative, fall into two categories: (1) the construction of an electronic brain and (2) learning processes in computers.

In spite of his strong interest in concrete physiology and particularly in neurophysiology — after all, he was a trailblazer in that he wrote one of the first mathematical works on the chemistry of morphogenesis in 1952 -he refused to imitate the physiological brain structure in hardware. He was more interested in analyzing "the logical structure of the brain", proceeding on the assumption that every continuous system can be approximated with any degree of accuracy by a "discrete system". These assumptions opened a new perspective, i.e. that of simulating a brain approximated by discrete states in a UTM.

Given this process of creating a structural relation between the human brain and the UTM, the possibility of imitating human intelligence in a UTM seemed to be within reach. As early as 1941 TURING had dealt with the issue of chess-playing machines, extending the question as to whether a machine, i.e. a UTM, would be able to "learn" in a general sense.

TURING seemed a bit ambivalent in pursuing this issue (cf. TURING 1948 and TURING 1950).

On the one hand he is apparently aware of the fact that a general ability to learn like a human being requires appropriately "elaborate interactions with the world". The machine would have to be equipped with TV cameras, microphones, loudspeakers etc. so that A was able to interact with the outside world as much as possible. Moreover, it would have to be able to "roam the land", to have "personal initiative" to be "trained" and "educated", in short, it would have to have everything a human child has when learning.

On the other hand there are passages where he is against an over-imitation of human beings and their natural qualities. Machines with a "reduced body" were just as interesting for him as candidates.

The existence of "artificial intelligence" was to be recognized by the criterion of "imitation". Whenever a human being reached the "conclusion" that the partner communicating with him/her exclusively via a terminal could just as well be human, the descriptions which would normally be applied to human beings only could also be applied to this "artificial intelligence".

3 MEANING, KNOWLEDGE OF THE WORLD AND CONSCIOUSNESS

The discussion triggered off by TURING as to whether machines could develop intelligence comparable to that of human beings and whether they could even develop something like "consciousness" has continued until today and the problem can definitely not be regarded as solved. In spite of the fact that we have grown accustomed to computers, the issue remains philosophically explosive (DÖBEN-HENISCH 1993). Language seems to be a key to further clarifications in this context. TURING already saw this quite clearly. When enumerating

potential applications for an intelligent machine, he said something along the following lines: "Among the above-mentioned applications, learning languages would be the most impressive one because it is the most human of these activities." (TURING 1948. German version 1987: 98)

TURING also predicted, albeit half-heartedly, that a UTM in a supposed position to learn any natural language similar to a human being, and to use it as required in the respective situation, would also need the appropriate interactions with the world so as to acquire the knowledge of the world necessary to use the language properly (e.g. TURING 1948, German version 1987: 98).

However, the statement that one needs knowledge of the world does not necessarily say anything about (1) "what knowledge" this would exactly be, and (2) "how" such knowledge could be "acquired" or not, as well as (3) "how" this knowledge would have to be "represented internally" so that it can interact with the linguistic system.

Without a theoretical framework of reference within which the above-mentioned questions are to be answered, the answers to (1) through (3) will be random because they have no locus. This also holds true for questions such as (4) "how" a "language system evolves" in a speaker-listener and (5) "how" the "language system" can begin to -interact with the knowledge of the world".

Contrary to the empirical sciences, the questions (2), (4) and (5) have only played a marginal role in philosophy until today — if they played any role at all. As a consequence, questions (1) and (3) as well as special partial aspects of (5), such as the question about the meaning of verbal expressions, were treated as isolated static aspects of a process of learning that is per se dynamic. It is against this background that one can at least partially understand why the view that the meaning of verbal expressions can also be reconstructed without reference to facts embedded in consciousness specifically developed in philosophy (cf. FREGE 1892. WITTGENSTEIN 1921, DAVIDSON/HARMAN 1972, BARWISEIPERRY 1983, to name just a few).

However, if one adopts the position of learning or of the learner -referred to as the "agent" here — and accepts questions (2), (4) and (5) as the guiding questions, one is forced not only to assume an agent-world system as a minimal framework, but also to think about the "interface with the world" which the agent has to have to translate the processes required in (2), (4) and (5) into reality.

TURING himself did not answer these questions in their entirety.

To be able to construct the internal function f of an agent we have to make a fundamental decision: should the assumptions concerning f (a) be "arbitrary" or will we (b) orient ourselves on certain "givens"?

A decision in favor of (a) would place the project of constructing f in the realm of a general structural science such as mathematics and would possibly limit R to the computable functions.

A decision in favor of (b) leads to two further decisions: should the construction of f (b.1) depend on certain behavioral data from human test subjects or (b.2) will we introduce the knowledge every human being has from the angle of his/her self-consciousness?

A decision in favor of (b.1) leads to the paradigm of the cognitive sciences trying to develop computer models on the basis of behavioral data (phonetics, psychology of language, physiology...), with the functions of these models corresponding as far as possible to the behavioral data. However, proceeding according to (b.1) has a serious drawback: the formation of models on the basis of behavioral data is "highly underdetermined" as regards the formal possibilities of the structures internal processes may have.

The question for variation (b.2) remains open. The first response to this variation will be the prejudice (rather widespread nowadays) that data acquired from our consciousness are unsuitable for serious reconstruction.

However, it can be said in favor of the use of data from our consciousness that there are a number of philosophical works showing that many aspects of cognition only become accessible by reference to one's own consciousness (e.g. MACH 1922. HUSSERL 1913. MERLEAU-PONTY 1945). Publications from the field of experimental psychology on perception and memory (e.g. MURCH[WOODWORTH 1978. KLIX 1980, HOFFMANN 1982. SHIFF 1980) have indirectly corroborated the philosophical arguments.

This approach to cognition, which can roughly be described as phenomenological, was in the past characterized by a lack of sufficient instruments of linguistic critique, which has had quite markedly adverse effects on A's acceptance in the field of philosophy until today. The introduction of modern linguistic critique in phenomenology as well as developments in the modern theory of science after 1970 (LUDWIG 1970. BALZER/MOULINES/SNEED 1987) allow for the formulation of a paradigm of cognition enabling us to elaborate formal theories of consciousness structures on the basis of phenomenological analysis. The formal theory formation we thus have at our disposal makes it possible to "represent" "almost all" (DÖBEN-HENISCH 1994c) aspects of consciousness and renders this representation "controllable".

The approach to a solution described under (b.2) is designed in such a way as to enable "the application of the theory to itself" (an answer to NAGEL 1986). Due to A's choice of means, such a philosophical theory is also "compatible" with any empirical theory because on a formal level, A contains every conceivable empirical theory as a partial theory. From this point of view, the suggestion that the gap between the sciences and the humanities cannot be bridged turns out to be a mere artefact of misplaced boundaries separating paradigms.

The decisive point is that "simulation models" can be defined on the basis of such a formal philosophical theory of self-consciousness. Simulation models limited to the use of UTMs are a special partial category. However, simulation models are by no means mere spin-offs of a self-contained formal theory of consciousness. In view of the mind-boggling complexity of processes of consciousness we should rather assume that the simulations are a necessary aid so that theories of self-consciousness with a certain degree of complexity can be developed at all. In this sense, simulation models are "instruments for philosophical experiments".

The procedure briefly described here may well form the basis of a new "discipline" which — depending on your point of view — could be "called Computer-Aided Philosophy" [CAP], or better "Computational Philosophy" [CP] if you feel more akin to a philosophical interest in cognition, or "Artificial Consciousness" [AC] if you want to place it in the general framework of computer science and pursue it as a new branch parallel to the existing Artificial Intelligence [AI] branch.

4 KNOWBOTIC INTERFACE

To translate the concept of computational philosophy into reality it is necessary for a formal theory of selfconsciousness as well as a suitable simulation model to be developed simultaneously.

BW1 is designed to be "the first prototype" of such a theory-guided "simulation model of self-consciousness" implemented in the framework of a Knowbotic Interface "[KInt]".

First formulated in the spring of 1994 (DÖBEN-HENISCH 1994a) and further developed theoretically in cooperation with Prof. HOCHÉ in a research group in Bochum (DÖBEN-HENISCH 1994b), the KInt addresses the consciousness issue and falls back on an idea proposed by TURING, whose vision of a learning machine was a child-machine. R would, of course, be a UTM, the nature of which would allow for it to be subjected to processes of education and training like a child (TURING 1950: 1987. pp. 177f).

Along with the KInt, a set of software "building blocks" would be made available so that the UTM Child Machines can be defined as TURING imagined them, complete with random UTM Worlds which "normal people", too, can enter, albeit in the guise of UTM Child Machines. The author calls these UTM Child Machines Knowbots (derived from to know + robot = knowbot) whereas the UTM Child Machines "disguised as" human beings are called "Pseudo-Knowbots". (The name Knowbot was inspired by the author's intense discussions with Christian Hübler of the group knowbotic research. The group kr+cf also uses the term 'knowbot', albeit in a different sense. The word "knowbot" is occasionally also used on the Internet to denote intelligent agents collecting information on the net. These "ordinary" knowbots and the knowbots of KInt have only their name in common.)

Thus, the KInt offers human beings the opportunity to train and educate knowbots without the knowbots necessarily having to recognize the pseudo-knowbots as something that differs from them. In this way, the KInt opens up a new variation of TURING's imitation test.

In contrast to TURING's approach, the KInt greatly emphasizes the fact that the internal structure of the knowbot, viz. its consciousness, and the structure of the KInt's world, correspond to the world known to us human beings, in so far that all facts relevant to meaning in the context of the natural languages known to man can, in principle, be simulated.

Due to the flexibility of the KInt other world and consciousness structures than those known to the human world can be simulated, too. In principle, topics to be dealt with in the framework of the KInt could well be derived from the realm of science fiction novels: artificial consciousness structures substantially different from ours, speaking languages of their own, possibly understanding human languages whereas we do not understand theirs. In the KInt framework such experiments can be conducted directly.

5 THE BLIND'S WORLD 1




Illustration 1 shows the basic structure of BW1.

First, one has to differentiate between "a server program" administering a certain world and a "client program". The latter can be a knowbot or pseudo-knowbot Client and server programs can either be started on the same or different computers, i.e. the "client-server architecture" is fully "network-compatible". Several server programs may exist side by side on the same computer or on different computers.


The "world" implemented by the server program is initially loaded from a text file and internally built up by the server as a world data structure. The text file may be edited like a normal text document viz. the user may define the world at will the interior of a house, a town, an authority, a certain country or an entire planet. In the first version, the possible world representations are limited to two dimensions with five different strata. The BW1 worlds all resemble colorful maps, with multi-colored pictographs moving around on them.

All objects in BW1 are "multi-sensorial objects". Each object is depicted as a figure with an internal structure and colored patterns in the visual dimension. It may also come with certain sounds in the acoustic dimension, a certain smell in the olfactory dimension, a specific taste in the gustatory one and various haptic qualities in the tactile dimension. These qualities may also change continuously according to internal conditions or exterior influences. The multi-sensorial character of all objects is a necessary consequence of the demand that the (pseudo-) knowbots communicate with the BW1 world exclusively through sense organs.

All things considered, the prototype BW1 world is very simple: it is a world seen as a huge surface area, the boundaries of which have not been defined so far. Inspired by the biblical account of the creation of the world, these worlds are basically "expanses of water" into which isolated or linked "landmasses" can be inserted at will. Each country can be designed from a stock of "landscape types". These include, among other things, "deserts", "grassland", "woods", "rocks", "rivers", "lakes" and "paths". Any number of individual objects can then be placed on the land characterized this way. The "types of objects" available are "plants" and "living beings". Plants fall into the categories "small plants", "bushes" and "trees". Living beings can be "small animals", "predacious animals" and "big game"; (pseudo-) knowbots may become special objects here if they register as clients in the world process.

Apart from the sensory qualities, plants and living beings have additional characteristics. Plants have e.g. certain "nutritional values" which can have various effects on living beings. Living beings are endowed with at least the basic needs such as "hunger", "thirst", "fatigue" and the "need to procreate". They may also feel "pain" and "fear".

Once the clients have logged in, registering as "citizens of the world" on a server, "the world process" passes as a sequence of world cycles. A "world cycle" is a processing cycle of the server, which requires a certain physical period of time. (ill. 2)



Depending on the quantity of data to be processed and on hardware capacity the physical time required for a single world cycle may vary greatly. One cycle within the world process corresponds exactly to a minimal unit of world time, i.e. the world process functions on the basis of "logical time".

Within a world cycle there are basically two activities: (1) the "interaction manager [IM]" reads all action messages which have accumulated in the actionmessage mailbox since the last evaluation. It checks each message for possible collisions. Whenever R detects a collision, it calculates a corrected final position for the movement and passes it on. Once the repercussions of the action messages have been recorded, (2) the "sensory manager [SM]" takes over. It calculates all possible sensory stimuli that may occur "according to the world laws" in the places where living beings are located at that point of time, and it does so for each living being, especially for the (pseudo-) knowbots. The totality of the values the SM calculates for a certain living being is then "rolled into" a sense message and sent to the being. When the SM has completed its task, a world cycle is over.

Action messages and sense messages are strings of signs (ASCII strings) the syntactic structure of which is determined by context-free grammars.

6 KNOWBOTS

The KInt described above is required to make a minimum of "ambient conditions" available for the simulations of a structural equivalent to consciousness.

The following passages refer to the framework conditions envisaged for the first BW1 experiments involving the phenomenon of self-consciousness.

6.1 Minimum Reactive System

In BW1 it is assumed that the knowbots have at least sense and efferent organs.

The sensory capacities are limited to the senses of "hearing", "smelling", "touching" and "tasting" (vision was left out in BW1). The developments of values communicated to the "interior" of the knowbot by the sensors are processed there: they are transformed into a three-

dimensional sensory map maintaining their topology. From here they are available for further use.

The efferent capacities are based on a finite set of elementary actions (walking, turning around, moving one's hands, grasping something, putting something away, eating, drinking, sleeping, mating, playing) which can be expressed in terms of parameters. Several such elementary actions can be summarized in complex actions with the conditions for the expression in parameters remaining intact. In BW1 it is assumed that each knowbot has a number of behavior patterns existing "ab ovo" which ensure basic behaviors such as searching for food, eating, procreation and flight. "A behavior pattern is a set of behavior rules, with each behavior rule in turn being bi-conditional as follows: OBJECT 0 can be attained if CONDITION C1 is fulfilled & ... & CONDITION Cn is fulfilled. The left-hand side of the rule is its "head" and the right-hand side of the rule is its "condition". In a general case more than one object can be formulated in one's mind and disjointed statements are possible in the condition. The "pattern" character of behavior rules evolves because variables can also appear as arguments for the objects and conditions. Such a behavior scheme thus stands for a whole set of possible specific behavior rules, depending on the concrete values that may be substituted for the variables.

Response mechanisms also include simple motor planning activating the respective behavior patterns and coordinating the actions to be taken. This encompasses the utilization of sensory feedback as well as the implementation of "tendencies" given by the emotional control unit.

"Emotions" and various "physical states" form the third basic component. Physical states include the "energy balance", "balance of fluids" and "waking/sleeping states". Instincts and drives are counted among emotions (hunger, thirst, fatigue, sexual desire, play behavior), as are pain, fear and an "overall emotional state" that is not identified any further. These components have a bearing on one another, with the overall emotional state "dominating"; however, depending on their respective extent, instinct and pain may "override" it.


There are "activators" and "deactivators" for each emotion. Hunger is e.g. activated by a negative energy balance and deactivated by food intake, depending on the properties of the food.

At any given time there is a "dominating" emotion. It has a bearing on the design of the action plans. Changes in the emotions may lead to the interruption of ongoing actions or the reversion to a plan previously abandoned.

The motor planning unit thus completely depends on the current emotional state unless a behavior plan activated at the time has a "higher" priority.

Physical states are influenced by actions as well as sensory inputs. In turn, they have a bearing on the emotions.

The entirety of these basic components forms what in BW1 is called a "reactive system" (ill. 3). viz. these components are per se a functional unit in a position to move around in the world and to feed itself, albeit poorly. Further components may now supplement the reactive system, thus modifying it.



6.2 The Necessity of Learning

One of the major drawbacks of the reactive system is its inflexibility; except for simple variations within given boundaries, it is not capable of learning anything. It converts impacts from its surroundings into certain modes of reactions "in a certain way", but the mode of reaction f does not undergo any appreciable change while the system is in existence. The mode of reaction f can be regarded as a function mapping elements from the set S of stimuli on elements of the set R of reactions.

To turn a reactive system incapable of learning into a system "capable of learning" substantial changes in the mode of reaction f must be possible, i.e., one must be in a position to replace the current version of f by a different version f' from the set of "possible modes of reactions" F .

In such a system there is an additional learning function L , which substitutes a new version f' for the current version of a mode of reaction f by reverting to the set F . However, "blind" learning, which replaces a current mode of reaction f by "any mode of reaction f' that is different", is not very effective. Improvements attained this way require "long" periods of time as well as high population density and multiplication rates. By contrast, behavior related learning is limited to the individual and happens in a comparatively "brief" period of time. This is only possible if one assumes that behavior-related learning is put into practice as "informed learning". The least one could then demand from a system capable of learning would be the existence of "rating criteria" which it can use to co-determine the design of the mode of reaction f . Such an evaluation could e.g. base ratings on environmental impacts.

Apart from the pure behavioral function f , there would also exist a rating function ev indexing environmental impacts with positive or negative ratings, and a modified learning function L , which takes ratings into consideration when selecting new modes of behavior f' from F .

- (1) $F: S \rightarrow R$
- (2) $ev: S \rightarrow E$
- (3) $L: F \times E \rightarrow F$

6.3 Learning in BW1

In BW1 the abstract postulate that knowbots are conceived of as systems with an informed learning function is translated into reality in several places simultaneously.

In a knowbot the "rating function" ev materializes first and foremost through the impacts which ingestions and perceptions have on physical and emotional states. If e.g. eating a certain foodstuff appeases hunger, such food, including its taste and smell, correlates with a positive rating. This figure can be used to evaluate a certain object as well as a certain action in the context of a certain need gratification.

Making available, at a later point of time t_{j+c} , correlations of actions, perceptions, ingestions and their effects taking place in a certain interval of time Δt_{ij} requires the ability to "remember" them appropriately, viz. they have to be "stored in such a way that they can be remembered". This requires a memory structure capable of "filing objects away" and "finding them when required" while taking into consideration their different sensory qualities as well as their positive or negative impacts on emotional states and recording the actions involved and the relevant spatial structures.

Finally, there must also be a possibility of utilizing correlations that have been rated and made retrievable as described above when planning concrete behavior. That is to say, if the action-planning unit has decided that a certain sequence of actions is required to reach a certain goal, a system capable of learning must allow for the "given behavior patterns" to be modified on the basis of "experiences". Such modifications may take place on various levels. All in all, the necessity of an option to include "correlations that can be remembered" in current action plans is an additional challenge to the method of "storing" such correlations.

As BW1 has not been in existence very long, these memory structures have so far been implemented in a rather rudimentary way.

6.4 Language

Apart from its expressive material (sounds, signs) a natural language also consists of what the verbal expressions "signify", what they "say" and the effects they "produce" when they are used. These aspects of verbal expressions that go beyond the expressive material will, for our purpose, be called "meaning" by way of simplification.

For us human beings, meaning in language is inseparable from the relationship we have with ourselves and with the world.

For the knowbots in BW1 this relationship with themselves and the world is only available to the extent to which it is provided by the reactive system capable of informed learning as described above. As will be shown below, this structure is still far too simple for the realization of discerning structures of meaning as we know them from human every-day language. No more than the most primitive naming and simple sentences consisting of one, two or three words are possible.

The language module is composed of several submodules.

The dictionary consisting of various sub-functions is an important basic unit. A word generator tries to form phoneme sequences to be stored in the dictionary from a set of elementary sounds and intonation patterns. The phoneme sequences will, however, only be kept in the dictionary if they obtain external "confirmation" within a certain period of time. A word-object generator forms word-object hypotheses. i.e. It tries to link words with states or actions. The respective current situation has priority in the formation of such hypotheses.

Word-object hypotheses, too, will only remain in the dictionary if they obtain confirmation from outside.

To be able to understand the grammar, one has to be aware of the tasks that have to be fulfilled in communication by means of language. Basically, BW1 differentiates between speaking and "listening". Ill. 4 demonstrates the procedure involved in speaking.

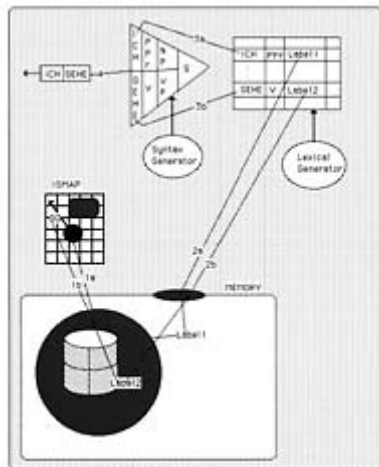


Bild 4 Bild 4

On the basis of an "interpreted sense map [ISMAP]" a pragmatic module called "language map [LMAP]" develops the differences between speaker and potential listener as well as various relations existing at that time between the speaker, its environment and the listener. The results are suggestions for various speech acts. Once a speech act has been selected, the objects to be included in the speech act are activated from the dictionary. The word forms in the dictionary, including various lexical categories, then refer to the grammar. The grammar has grammatical rules determining how words of a certain lexical category can be combined with one another while taking into consideration their object index and the speech act involved. The result is a concrete sequence of words that can be "heard" as a specific utterance.




Bild 5

Ill. 5 shows the reverse case, listening to an utterance, which is considerably more demanding than speaking. A speaker, not the same person as the listener, says, "I go". The schematic grammatical rules first take the sequence of words apart: the individual elements are able to refer to the appropriate objects via the dictionary. With the help of the LMAP, the listener

tries at the same time to comprehend the potential speaker and the resulting relations to it so as to determine possible forms of speech acts. In this context, the listener has to perform a rather complicated transfer, as R has to "regard" the "I-object" attributed to the word "I" as the model for a different object. In keeping with this, the accessible aspects of the utterance "I go" have to be transferred to the "other- I -object" and its relations to the environment.

It must be noted that the "notion of comprehension in every-day language" is much more complex than the process of linguistic decoding during the grammatical analysis of an utterance as described above. In everyday life the result of verbal recognition usually leads to further processes of cognition such as logical deductions, conclusions by analogies and graphic associations. Purely verbal comprehension has to be delineated as an -albeit necessary — sub-process of this general notion of comprehension. Due to the short time available, only the first tentative stages of verbal comprehension have been implemented in the BW1 knowbots.

As all participants in BW1 are blind, learning the connections between linguistic signs and random perceptual events can only be influenced "from outside" to the extent to which it is possible to create a potential link among portions of a verbal utterance on the one hand, and perception or the current situation model on the other by virtue of the momentum of "simultaneity". The non-availability of vision requires that any ambiguity in the process has to be set off by supplementary haptic values, which is not always easy.

7 HARDWARE, SOFTWARE, MANPOWER

When BW1 was designed, the enormous demands on inter-process and inter-object communication in the installation were met by one operating system only, i.e. NEXTSTEP. Hewlett Packard were the only ones to provide the high-performance hardware on which NEXTSTEP could be installed during the development period. For this reason, the development platform we selected consisted of HP workstations with NEXTSTEP 3.2 as the operating system: one HP 712180 and one HP 712160 with 64 MB RAM as well as a 1 GB harddisk each. In addition we used an Aquarius PC 486/66 with 32 MB RAM and a 1 GB harddisk to run NEXSTEP 3.3. For the development of the XWindow Motif version for the normal-user client, an ESCOM PC 486166 with 16 MB RAM and a 1.3 GB harddisk were available. The operating system was Unifix 1.5. and included Motif 2.0. Furthermore, a version of the XClient was compiled under IRIX on Silicon Graphics computers. All computers were interlinked by means of EtherNet.

Leo POS and Thore SWINDALL programmed the server software and the NEXTSTEP version for the super-user client. Thore SWINDALL and Raoul SCHOLZ made the perception and memory module in the knowbot. The action module and parts of the language module were programmed by Leo POS. Michael KLOCKNER was in charge of the language module.

Joachim RASCH and Sonja SCHELLENBERG programmed The XWindow-Motif version of the normal-user client in the course of their theses.

The working paper "Agents with Consciousness" inspired Software programming under NEXTSTEP. The Theoretical framework for the Knowbotic Interface Project, Phase 1: The BLIND'S WORLD I" by the author. In late March, the work comprised 130 pages plus numerous color illustrations.

All work was done at the Institut für Neue Medien, where director Michael KLEIN generously made available a project space for us. Another special note of

[-EOF-]