Linguistic Computation with State Space Trajectories

Hermann Moisl

Centre for Research in Linguistics, University of Newcastle upon Tyne

Abstract. This paper addresses the key question of this book by applying the chaotic dynamics found in biological brains to design of a strictly sequential artificial neural network-based natural language understanding (NLU) system. The discussion is in three parts. The first part argues that, for NLU, two foundational principles of generative linguistics, mainstream cognitive science, and much of artificial intelligence –that natural language strings have complex syntactic structure processed by structure-sensitive algorithms, and that this syntactic structure determines string semantics– are unnecessary, and that it is sufficient to process strings purely as symbol sequences. The second part then describes neuroscientific work which identifies chaotic attractor trajectory in state space as the fundamental principle of brain function at a level above that of the individual neuron, and which indicates that sensory processing, and perhaps higher cognition more generally, are implemented by cooperating attractor sequence processes. Finally, the third part sketches a possible application of this neuroscientific work to design of an a sequential NLU system.

Introduction

The key question of this book is: 'What can we learn from cognitive neuroscience and the brain for building new computational neural architectures?'. This paper addresses that question in relation to natural language processing (NLP), and the answer involves application of chaotic dynamics found in biological brains to design of artificial neural network (ANN)-based string processing architectures. Specifically, it argues that, in designing and implementing NLP systems for semantic interpretation of natural language strings, henceforth referred to as natural language understanding (NLU) systems:

1. Two foundational principles of generative linguistics, mainstream cognitive science, and much of artificial intelligence (AI) and NLP –that NL strings have complex syntactic structure processed by structure sensitive algorithms, and that this syntactic structure determines string semantics– are unnecessary, and it is sufficient in principle for NLU purposes to process strings purely as symbol sequences, but there are substantial practical problems associated with sequential NLU.
2. Neuroscientific results support the principle of sequential NLU, and also provide potential solutions to the practical problems associated with it.

S. Wermter et al. (Eds.): Emergent Neural Computational Architectures, LNAI 2036, pp. 442–460, 2001.

3. It consequently makes sense to investigate the feasibility of sequential NLU by developing an NLU system based on the results in (2).

Two preliminary comments. The first is that NLU is here understood as language engineering –as the design and implementation of machines that process natural language for some purpose– and not as cognitive modeling. Secondly, given the current deluge of results on brain mechanisms generated by new imaging techniques, there are bound to be and in fact are controversies about their interpretation. Given the above construal of NLU, however, there is no need to engage in these controversies: it is legitimate to select and use brain mechanisms proposed by neuroscience purely on the grounds that they seem useful from an engineering point of view. The ideal, of course, is to engineer in accordance with the way the brain really does language, since it is the only known physical device that implements all human linguistic abilities and is thus the definitive model, but pending final understanding of the relevant mechanisms, there is no need to wait for the dust to settle.

1 The Need for Complex Syntactic Structure in NLU

This section is in two main parts. The first motivates the challenge to the need for complex syntactic structure in NLU, and the second presents the case for the principle of sequential mapping of NL strings to meanings.

1.1 Motivation

In the 1930s and 1940s, mathematical logicians formalized the intuitive notion of an effective procedure as a way of determining the class of functions that can be computed algorithmically. A variety of formalisms was proposed –recursive functions, lambda calculus, rewrite systems, artificial neural networks, automata– all of them equivalent in terms of the functions they can compute. Automata theory was to predominate in the sense that, on the one hand, it provided the theoretical basis for the architecture of most current computer technology, and, on the other, it is the standard computational formalism in numerous science and engineering disciplines. Automata theory was, moreover, soon applied to modeling of human intelligence, and has dominated thinking about human cognition ever since [43,40,20,39,80]. This approach to cognitive modeling attained apotheosis in the late 1970s, when Newell and Simon [61] proposed the Physical Symbol System Hypothesis (PSSH), where 'physical symbol system' is understood as a physical implementation of a mathematically stated effective procedure, a prime example of which is a programmed Turing Machine. The essence of the PSSH approach to cognitive modeling was set out by Fodor and Pylyshyn in 1988 [32]:

- There are representational primitives –symbols– called atomic representations.
- Being representational, symbols have semantic content, that is, each symbol denotes some aspect of the world.

- A representational state consists of one or more symbols, each with an associated semantics, 'in which (i) there is a distinction between structurally atomic and structurally molecular representations, (ii) structurally molecular representations have syntactic constituents that are themselves either structurally molecular or structurally atomic, and (iii) the semantic content of a representation is a function of the semantic contents of its syntactic parts, together with the syntactic structure'.
- Input-output mappings and the transformation of mental states 'are defined over the structural properties of mental representations. Because these have combinatorial structure, mental processes apply to them by virtue of their form'.

Therefore, 'if in principle syntactic relations can be made to parallel semantic relations, and if in principle you have a mechanism whose operation on expressions are sensitive to syntax, then it is in principle possible to construct a syntactically driven machine whose state transitions satisfy semantic criteria of coherence. The idea that the brain is such a machine is the foundational hypothesis of classical cognitive science'.

Modeling of the human language faculty has been paradigmatic for the PSSH-based approach in cognitive science. The discipline concerned with this language modeling, generative linguistics, has developed considerably since Chomsky's pioneering work in the mid-1950s, but at least one foundational principle has remained unchanged: that NL strings have a structure beyond the strictly temporal or spatial sequentiality of speech utterances and text. More specifically, they are held to have a compositional phrase structure of the sort described above by Fodor and Pylyshyn, or, slightly more formally, given a phrase structure grammar (or equivalently, an automaton) that generates a language L, the tree diagrams that represent the structure of sentences in L must allow simultaneous left and right phrasal nonterminal branching from parent nodes; this is what is meant by 'complex syntactic structure' throughout the current discussion. The notion that such phrase structure determines sentence meaning is, moreover, central to currently dominant approaches to NL semantics [55,76, 25,11,26]. In model theoretic semantics [63,54], for example, the link between syntactic structure and meaning is made via the principle of compositionality, which says that the meaning of a syntactically complex expression is a function of the meanings of its constituent words and of the syntactic rules by which they are combined, and are realized by grammars in which each syntactic rule is associated with a semantic rule that specifies the meaning of a constituent in terms of the meanings of its own immediate constituents (the 'rule-to-rule' hypothesis). A derivation according to such a grammar on the one hand generates a sentence with an associated syntactic structure, and on the other the meaning of the sentence; this explains how syntactic structure systematically determines the meanings of the sentences of a language.

The PSSH has, furthermore, been influential not only in the cognitive sciences but also in AI, here understood as an engineering discipline which seeks to design and implement machines that emulate (aspects of) human behaviour. AI has historically been closely associated with PSSH-based cognitive science [43,73]: in essence, cognitive science has proposed models, and AI has attempted to implement them. NLU, more particularly, has been closely tied to developments in generative grammar.

Now, since the early 1980s, perceived shortcomings of the PSSH-based approach to cognitive science has generated a variety of challenges to its dominance –by artificial neural network [58,2,21,22] and dynamical system theories [83,48, 69,10,84,38,62,4] as paradigms for cognitive modeling, and by a radical shift of emphasis away from the study of the innate 'higher' functions like reasoning and language to concentration on developmental issues like the evolution of cognition and the interrelationship of mind, brain, body, and external environment [47,23, 46,39,62]. At the same time, the main thrust of AI research has retreated from the grand aim of general, human-level machine intelligence and has concentrated on developing systems that work well in restricted domains, such as expert systems, handwriting and speech recognition, and robotics and vision in specific industrial applications [31,57]. In response, at least some of those AI researchers who still believe in the achievability of general machine intelligence have increasingly turned to the same approaches as their cognitive science colleagues [6,3,82,29,57]. In NLU specifically [42,7,1,27,24], there now exist practical natural language interfaces in domain-specific applications, primarily databases and expert systems, but to my knowledge no reliable, broad-coverage natural language understanding system currently exists, nor do the prospects for one look promising. The problem is not with syntax –after several decades' effort, parsing algorithms capable of supporting reasonably large-scale general NLU systems are available [67,1,77,66,5] – but with implementation of semantics and pragmatics [42,1,68,56,71]. For natural language understanding, linguistic expressions must be related to the system's awareness of, interaction with, and expectations of the world, and the main lesson of the NLU work done in the 1970s and early 1980s is that any NLU system needs, at the very least, to represent, manipulate, and update real-world knowledge efficiently. Various knowledge and belief representation and update mechanisms have been developed [49,1,57], such as logic formalisms, semantic networks, schemas, frames, scripts, and rules, and while these have been shown to work well in small-scale, carefully managed applications, none has so far been scaled up successfully to usefully large general NLU systems.

The motivation for proposing to dispense with compositional structure and structure-sensitive processing in NLU, therefore, is simply that it has reached an impasse [30] (see however reviews of [30] in Artificial Intelligence 80 (1996)). This is not to claim that the PSSH-based approach to NLU cannot succeed. If human language, and cognition more generally, are computable functions [64,65, 9], then a PSSH virtual machine that implements them must exist. The problem is that no one has found it, nor does it look like it will be found soon, and as such it seems sensible to try other approaches.

1.2 Sequential Mapping of Strings to Meanings

It should be uncontroversial to observe that the research agendas of generative linguistics and NLU differ: to judge by their extensive research literatures, linguistics is science and NLU is engineering. Like any other science, linguistics aims to state empirically adequate and maximally expressive theories to explain its domain of interest, whereas the aim of NLU is to design and construct physical machines that respond to natural language utterances or text in an acceptably human-like way. Being physical, every NLU system is bounded in all aspects of its operation –in the lengths of its input and output strings, in the number of strings in can process in its operational lifetime, and in the memory and processing time at its disposal. This means that the unbounded input and output tapes and memories of automata theory need not be a design consideration, and in particular that the I/O sets which an NLU device will be asked to process must be finite. Now, a fundamental result in automata theory is that any finite mapping can be implemented by a finite state automaton (FSA) [44,8], and to an FSA every string has the same structure: strict sequence. It follows that any NLU function can be implemented by an FSA, and that sequential processing of input strings is sufficient for the purpose.

That FSA architecture, and thus sequential processing, is sufficient for NLU has been known since the early years of generative linguistics. Indeed, Chomsky himself pointed it out [12,14,15]. Despite that, finite state NLU has been all but ignored until fairly recently [74,51,18,50,19,52,70]. The various reasons for the lack of interest in FSAs are valid from a generative linguistics point of view, but irrelevant for NLU. These are briefly dealt with in what follows (see [16] and [17] for closely related discussions).

– Trivial finiteness
  One might want to argue that to insist on the finiteness of NLU functions is an excessively theoretical point –that string length, I/O set sizes, and the number of processing states / memory size in a real-world application might be so large as to be unbounded for all but the most abstemiously theoretical purposes– and that this disqualifies FSAs in practice. This has some initial plausibility, but consider these two arguments. Firstly, how long are NL strings in daily speech intercourse and text production? There are almost certainly studies that provide maxima, minima, means, and standard deviations for string length across different languages, but speaker / reader intuition serves to make the required point here (for a discussion of sentence length bounds see [53]). In everyday fact, most strings are very short – a dozen words, perhaps, or occasionally two or three times that number. Few extend beyond, say, 100 words, and those that do become increasingly incomprehensible. For NLU purposes, therefore, we are not dealing with what one might call trivially bounded strings which are so long –perhaps 100,000,000 words– that they are unbounded for any practical purpose. And, secondly, while it is true that, even with the above practical string length constraint, a native speaker of some language will typically produce a very

large number of strings in a lifetime, one has to keep in mind that NLU is engineering. Vocabulary size, string length, and the permitted range of syntactic patterns are under the designer's control. A fully human level of linguistic competence is very difficult to achieve, as the history of AI over the past several decades has shown, but if one is prepared to work to a less general specification the problem size can be scaled down by fiat.

- Weak generative capacity
  A fundamental result of formal language and automata theory is that FSAs cannot generate the language a^nb^n, that is, sets of strings in which a sequence of some symbol a is followed by exactly the same length sequence of another symbol b, where n is any positive integer. It has further been claimed (misleadingly, but that is another matter) that NLs contain strings with the a^nb^n pattern, and that FSAs are consequently unable to generate the class of natural languages [13,41,63]. This argument assumes unbounded n, which is legitimate given the aims of linguistic theory. But, in terms of NLU as language engineering, n must be bounded; where n is bounded it is prespecifiable, and in such a case the a^nb^n pattern can be generated by an FSA. The argument against finite state NLU from weak generative capacity is thus irrelevant. An analogous argument applies to the cross-serial dependency pattern that was used to disqualify the class of context free grammars / pushdown automata, and by transitivity regular grammars / FSAs, as a basis for NL modeling [41,79,63].
- Strong generative capacity
  Complex syntactic structure is fundamental to the explanatory capacity of linguistic syntactic theory because it allows intuitions and empirical findings about natural language word order to be expressed in satisfying generalizations. Because they cannot deal with complex syntactic structures, FSAs are effectively useless for linguistic modeling. Chomsky pointed out their impoverished explanatory capacity in the 1950s, and they have for that reason rightly been ignored by generative linguists ever since. On the present view, however, NLU is not linguistic modeling, but is interested in constructing physical devices with some particular I/O behaviour. And, in view of the foregoing discussion, there is no good theoretical reason to prefer any one automaton class over the others in NLU design: a finite string set does not imply the computational class of the automaton that generates it. In other words, explanatory adequacy is irrelevant to the choice of computational architecture for NLU purposes.
- Compositional semantics
  As noted, compositionality explains how syntactic structure systematically determines the meanings of the sentences of a language. An FSA can support a compositional semantics, but that semantics is explanatorily trivial. Because an FSA treats the syntax of every string as strictly sequential, an FSA-based compositional semantic theory can assert only one thing –that the meanings of all sentences in all languages are sequentially concatenative. This amounts to proposing a semantic theory which pairs each syntactically legal string with a meaning, and simply lists the pairs. Because a list cap-

tures no generalizations, FSAs are of no interest for NL semantic theory. Once again, however, NLU is not primarily concerned with explanation, and this conclusion is consequently irrelevant. A compositional semantic theory that uses complex syntactic structure gives a particular explanation of how sentences map to meanings, but from a language engineering point of view a list would be equally valid in principle. NLU is, in short, not required to adopt compositionality.

There is, then, no theoretical obstacle to a finite state, and thus sequential, approach to design and implementation of NLP and more specifically NLU devices. In fact, the past decade has seen a marked increase of interest in a finite state approach to NLP in such areas as speech recognition, phonological and morphological analysis, syntactic parsing, and information extraction from text [42,16,59,75,17]. There is, however, a significant silence on semantics, for reasons we are about to come to.

In principle, design of a device that maps strings to meanings is easily formulated. Because NLs are finite sets for NLU purposes, it is possible to define a string-to-meaning function f of (string, meaning) pairs. Once the set exists, all that is required for implementation is an algorithm for table lookup: given a string as input, the device returns the associated meaning as output. An appropriately configured FSA is one possible algorithm. Each lexically distinct string drives the machine through a characteristic state sequence, and the final state, which uniquely identifies the string, is mapped to the associated meaning. As a generative linguistic account of the human language faculty, this is of course completely inadequate since it amounts to no more than a list and hence explains nothing, as noted. For NLU, however, the only issue is whether or not such an approach is feasible in a specific NLU domain. And there are in fact difficult problems with it; the two main ones are:

1. Meaning and its representation
   To be able to construct a (string, meaning) pair list, it is necessary to have a clear idea what 'strings' and 'meanings' are, and to have a way of representing them so that they can be processed. The ontology and representation of symbol strings is straightforward. For meanings they are anything but. There is a long history of philosophical debate about how the concept of linguistic meaning might be understood, and, at the moment, there is a range of possibilities with little agreement among them [Craig 1998]. Model-theoretic NL semantics, pioneered by Montague, is the currently-dominant approach within generative linguistics [63,54], but others, such as the 'use' theory of linguistic meaning propounded by Wittgenstein, are still being developed [45]. Because it relies crucially on the notion of complex syntactic structure in the generation of sentence meaning, model-theoretic semantics is unavailable to a sequential NLU system for practical purposes, and choice among remaining alternatives is not clear. Nor, assuming a choice has been made, is it clear how abstractly-characterized meanings should be represented.
2. List construction
   Assuming a suitable meaning representation has been adopted, it is possible

to compile (string, meaning) pair lists explicitly for small, domain-restricted applications. Extension to general, real world NLU would, however, clearly be extremely onerous and almost certainly impractical.

The neuroscientific results described in the next section offer a good starting point for solutions to these problems.

2 Meaning and Sequential Processing in the Brain

On the basis of his work on biological sensory systems, Walter Freeman proposes fundamental principles of brain function and how these implement essential aspects of human cognitive behaviour. These proposals relate directly to sequential NLU, in that they include both a coherent view of the nature of meaning, and an account of the biological implementation of meaning that features purely sequential brain dynamics, with no reference to complex syntactic structure. We look first of all at Freeman's view of meaning and then at its implementation (what follows is based on [81,33,34,35,36,37,38]). A caveat, however. What follows is a very brief summary of a comprehensive account of human behaviour and its biological basis developed over a lifetime's research. As such, it inevitably and grossly oversimplifies, and may also misconstrue or misrepresent Freeman's work; apologies, where appropriate, are offered in advance.

2.1 Meaning

Intentionality is fundamental to Freeman's view of meaning. This is not the intentionality of twentieth-century analytic philosophy and of PSSH-based cognitive science, where the term denotes the 'aboutness' of mental representations, and is used to designate the relation between mental states and objects or events in the world, whether real or imaginary. Rather, his understanding of the term is derived from that of the thirteenth-century philosopher Thomas Aquinas, for whom intentionality had to do with goal-directed action in the world, and modification of the self in response to that consequences of that action as a way of coming to understand the world and the place of the self in it. Analyzed in terms of intentionality in this sense, an organism's existence in the world over time is a sequence of intents, where an intent has three stages: (1) formulation of a goal to whose realization an action will be directed, (2) perception of the environment, execution of the intended action, and perception of consequences for the self, and (3) learning from these consequences relative to the goal. Intents in the sequence are not independent, but are linked to one another in that goal formulation, execution, and adaptation at any given point in the sequence occur in the context of, and are informed by, the organism's lifetime intentional history up to that point. Thus, a hungry cat formulates a goal of catching food. On the basis of previous intentional action, it knows that it can achieve that goal by hiding in tall grass, priming itself for a spring, and then pouncing on an animal smaller than itself. It has a visual perception of suitable prey and pounces, but because the smaller animal on this occasion is an urban rat that fights back, the cat experiences pain. From the perceptions of rat and of pain, the cat learns to pick on different small animals in future. This sort of intentionality is characteristic of all vertebrates, in that they observably act in furtherance of their survival in specific environments and modify their behaviour in response to the consequences of their actions. There are degrees of intentionality, however: a salamander has much simpler, shorter-term goals and capacities for action and learning than does a human.

Meaning arises from intentionality. It is the organism's learned awareness of the interrelationship of perceived states of the world, its goals, and its actions in pursuance of those goals, at any given point in its intentional evolution. In the above example, the visual perception of a rat comes to mean pain to the cat as a result of learning. Meaning is, moreover, organism-specific: to a bigger, tougher cat perception of a rat might mean something very different. And, again, there are degrees of meaning commensurate with the richness of the intentionality available to a species. A given odour can presumably mean only a limited range of things to a salamander –food, danger, safety– whereas, to a human, it can not only mean such things, but also, for example, a subtle blend of emotional experience.

2.2 The Implementation of Meaning

Freeman models the vertebrate brain as a hierarchy of interacting nonlinear dynamical systems whose global evolution implements intentionality. In humans, these dynamical systems correspond to the modules of the limbic system, the sensory cortices, and the areas of the neocortex associated with higher cognitive functions. All three are required to implement full human intentionality, but to keep the discussion tractable the focus will be on the first two only. This will be sufficient for present purposes: the interaction of limbic system and sensory cortices is the necessary basis for human intentionality, and the principles that Freeman proposes for its operation extend straightforwardly to interaction with the neocortex (on which see [38] chapters 5-7).

The rest of this subsection develops the implementation of meaning in three parts. The first looks at the methodology and results of Freeman's work on the olfactory system, which is the empirical basis for his proposals on general brain function. The second then outlines his interpretation of these results in terms of dynamical systems theory. And the third describes his proposals for extending the dynamical systems analysis of the olfactory system to other perceptual modalities such as vision, and to the integration of these modalities in the enthorinal cortex and the hippocampus so as to implement perceptually grounded meaning.

The Olfactory System. Freeman has concentrated on olfaction because it is the primary perceptual system in evolutionary terms, and, in his view, is fundamental to understanding of the other perceptual modalities.

The operation of the sensory modalities has two distinct aspects: sensation and perception (see also [80]). Sensation is the transduction of external signals to spatial patterns of neural activation through stimulation of neural receptors on the sensory surface, here the nose, whereas perception involves transformation of this spatial pattern within the brain. Freeman has been concerned with perception. He has studied the behaviour of the olfactory system, comprising the nasal receptors, olfactory bulb, and olfactory cortex, in rabbits and cats. This was done by surgically attaching an 8 x 8 array of electrodes to the olfactory bulb. The animals were then trained to respond to conditioned and unconditioned stimuli and, when training was complete, an EEG was used to record the responses of the bulb both to training odorants and to odorants not previously experienced. The main observations relevant to the current discussion were:

- At rest, that is, when no learned odorant was present for the animal to smell, the EEG showed continual background activity with an aperiodic waveform.
- With each inhalation of a learned odorant there was a burst of activity which ended at the onset of exhalation. The burst was an oscillation with a common waveform throughout the bulb, but whose amplitude varied from place to place in it. This amplitude modulation (AM) constitutes a bulb-wide spatial pattern of activation. There was a characteristic AM pattern associated with each learned stimulus in the sense that each presentation of a specific odorant produced an AM pattern that differed from those associated with other learned odorants.
- When the animal was presented with a stimulus that was unconditioned during training, or a stimulus not previously encountered, there was no burst.
- When an animal trained on some set of stimuli was further trained to add a conditioned stimulus to its existing repertoire, two things happened: a new AM pattern was produced when the new odorant was presented after training, and all existing AM patterns were modified to a greater or lesser degree.

The Olfactory System as a Dynamical System. Neural populations within the brain are instances of macroscopic ensembles in an extensive range of naturally-occurring complex systems that evolve over time; examples are ecologies, social groups, and weather systems. They share the following characteristics with such ensembles: (i) there are many semi-autonomous elements, (ii) each element interacts with many others, (iii) the input-output relations of the elements are nonlinear, and (iv) there is an energy source and sink. The behaviour of such natural systems over time is standardly modeled using dynamical systems theory. A brain consists of a very large number of elements with nonlinear input-output behaviour that interact by means of rich and recurrent interconnections, and has a metabolic energy source and sink. As such it is a natural step to extend the dynamical systems approach to the study of the brain as well, and that is what Freeman has done.

The olfactory bulb and cortex constitute a system of coupled oscillators whose nonlinearity and recurrence permits the full range of behaviours associated with nonlinear dynamical systems, including not only point and limit-cycle but also chaotic attractors:

- The aperiodic wave which the EEG detects in the absence of sensory stimulus is the olfactory bulb in a chaotic attractor. Each module on its own has only a point attractor and a limit cycle attractor, that is, an oscillation with a characteristic frequency. The chaotic activity results from the coupling of the modules by both forward and feedback connections. Their characteristic frequencies differ, and they cannot agree on a common frequency, and so oscillate aperiodically. That this activity results from the modules' interaction is clear: on the one hand, it does not arise from stimulation originating elsewhere in the brain, since it persists even when the olfactory system is surgically isolated from the brain, and on the other, if the olfactory modules are surgically separated, they immediately settle into their characteristic oscillation. This chaotic basal state keeps the olfactory system in a constant state of readiness, so that even a small perturbation of the right sort can move it quickly to another attractor.
- The AM patterns are chaotic attractors in the bulb's state space. In a trained system, the stimulus from a specific odorant increases the gain in the bulb, and causes a state transition from the basal chaotic attractor to an attractor associated with that stimulus. On exhalation, the stimulus from the transduction pattern is removed, the gain is reduced, and the bulb returns to its basal chaotic attractor. In other words, the bulb is destabilized by gain from external stimulation.
- The odorant-specific attractors, that is, the AM patterns, are formed during training by the development of assemblies in the bulb via Hebbian synaptic modification. The result, in a trained animal, is that the olfactory bulb has a landscape of chaotic attractors including a basal attractor to which the system reverts in the absence of stimulation, and one for each of the learned odorants. The change in all existing AM patterns when a new odorant is added to the animal's training results from crowding of the attractors in the landscape as a new one is added.

Perceptually-grounded Meaning. On the above understanding of meaning, perception of an odour on its own means nothing; indeed, training on unconditioned odorants does not result in formation of an AM pattern / attractor for the odorant. To mean something, an odorant has to correlate with perceptions from one or more other sensory modalities. How is this implemented in the brain? By perceptual integration in the enthorinal cortex. The enthorinal cortex is one of the modules of the limbic system, and is remarkable first and foremost for the large number of other brain areas with which it interacts. Of particular relevance here is the fact that all the sensory cortices send their output to, and receive feedback from, the enthorinal cortex (see Fig. 1).

This indicates that the enthorinal cortex is the place where perceptual integration is implemented. In support, Freeman notes that, by observing the sensory cortices of rabbits trained in various sorts of conditioned input –light,

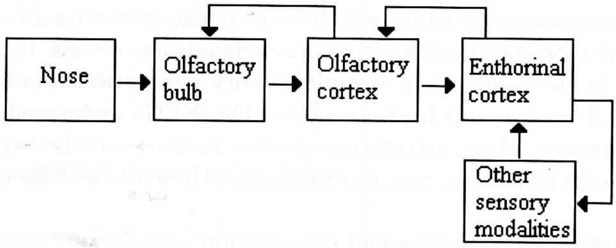

Fig. 1. Multisensory integration in the enthorinal cortex

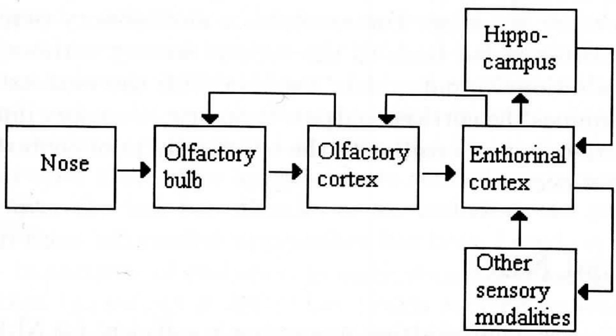

Fig. 2. Sequencing of multisensory perceptions via the interaction of enthorinal cortex and hippocampus

sound, touch– he was able to detect essentially the same dynamic behaviour as he observed in the olfactory system. The implication is that all the sensory systems use the same dynamics, and that, given this uniformity, the enthorinal cortex could in principle integrate its various perceptual inputs into multisensory perceptions or gestalts. Such a gestalt carries the meaning of a particular stimulus for the organism.

It needs, furthermore, to be kept in mind that intentionality is sequential. Individual sensory perceptions and integrated multisensory perceptions are embedded in a dynamic process which implements the organism's interaction with the world in real time. There are two issues to be addressed here (see Fig. 2):

– Implementation of sequencing

The enthorinal cortex, which as we have just seen is proposed as the locus of multiperceptual integration, sends most of its output to the hippocampus, and via feedback connections also receives most of the hippocampal output, so that the two modules are in constant interaction. Now, the hippocampus has been shown experimentally to be deeply involved in the orientation of behaviour in space and time, which is what the intentional loop does. It is, therefore, reasonable to infer that intentional space-time behaviour is implemented in the sequencing of multisensory perceptions via the interaction of enthorinal cortex and hippocampus. The EEGs generated by this interaction, moreover, show waveforms similar to those of the sensory cortices, which suggests that here, too, dynamics like those of the olfactory cortex are at work.

– Implementation of attention and expectation
Perception is an active process. Organisms do not simply respond to environmental stimuli, but fix their attention on aspects of the environment and have expectations of the environment for which they prepare. Attention and expectation are implemented by feedback signals that modulate response to the next sensory stimulus. For example, a multisensory perception in the enthorinal cortex is fed back to the various sensory cortices, thereby providing an additional input which, together with the next external sensory input, determines the cortices' output. Response to sensory input is thus not merely reactive, but is sensitive to the temporally prior context in which the current input occurs.

3 Sequential NLU

We have seen that sequential string processing is sufficient for NLU in principle, but that there are serious problems with it in practice. Freeman's proposals on how the brain implements intentionality and thus meaning both support the principle of sequential processing and offer an attractive approach to dealing with these problems. True, the proposals are based to a large extent on research into the role of the limbic system in sensory processing the results of which, when extrapolated to cognition more generally, are hypothetical and need to be empirically tested to gain acceptance in the brain and cognitive sciences. As noted in the Introduction, however, it is sufficient for NLU that they be plausible enough to be worth investigating relative to the engineering problem in question irrespective of their validity for brain function, and in my view they are. This section therefore sketches the relevance of Freeman's proposals to design of a sequential NLU architecture.

3.1 Sequential Processing

Freeman identifies attractor trajectory in state space as the fundamental principle of brain function at a level above that of the individual neuron. Each of the modules in sensory processing subsequent to transduction is driven by input through a sequence of attractors; there are no complex syntactic structures, only cooperating sequential processes. If that view is correct, then, since the brain in fact implements language, sequentiality in this sense must be sufficient. Freeman says little about language, and it may well turn out that, for this characteristically human and thus exceptional cognitive function, the brain does indeed use complex syntactic structure, in which case NLU would be well advised to use it as well. In the meantime, Freeman provides good cause for investigating sequentiality on its own.

3.2 Problems with Sequential NLU

Meaning Representation. However one understands 'meaning', it seems uncontroversial to say that the meaning of linguistic expressions has something to do with the way humans experience the material world. PSSH-based AI and NLU systems have formalized this relationship as a mapping from strings to states of the world, and have attempted to implement such mappings using explicitly-designed, system-internal representations of real-world domains of discourse. How best to represent (aspects of) the world and beliefs about the world is one of the main issues in PSSH-based AI, and a variety of formalisms for it exist, as noted earlier. It was also noted earlier, however, that this approach has only been moderately successful at best, and its relative lack of success is what has motivated alternative approaches to semantic interpretation and implementation over the last two decades or so, including the present one. The main trend in these alternative approaches has been based on a growing conviction of the importance of evolution in understanding cognition and the brain mechanisms that implement it [20] –that brains evolved as controllers for the bodies of creatures which had to survive in specific environments, and that the higher functions like language and reasoning are based on neural mechanisms developed for this purpose. As a consequence, recent 'embodied, embedded' cognitive science and AI have been much more interested than the PSSH tradition ever was in the interrelationship of mind, brain, body, and environment both as a means of explaining cognition and as a way of implementing cognitive functions (for example, [3,47,10,20,23,23,39,46,80]). With respect to meaning in particular [20], it is clear that, notwithstanding the concentration on linguistic meaning in the PSSH tradition, there is meaning apart from language –animals learn to behave meaningfully relative to their environments, and so do human babies before they have acquired language– and the idea that such meaning is the basis for linguistic meaning is at least plausible enough to merit further investigation. Consonant with this general trend, there has been a movement in AI and NLU away from attempting to map strings to explicitly designed representations which, inevitably, reflect a designer's analysis of what is significant in a task domain, and instead to relate language to system-internal states that are learned from interaction with the world without designer intervention. At its most ambitious, this approach aims to embed NLU systems in robotic agents that not only receive inputs from an environment via some combination of sensory transducers, but also interact with and change the environment by means of effectors. The aim is for the agents to develop internal states based on self-organization with the environment –'Concepts are thus the "system's own", and their meaning is no longer parasitic on the concepts of others (the system designer)' [28]– and are intended to learn the meanings of linguistic expressions from their use in environment-interactive situations.

Freeman's work is part of this trend in that it uses self-organizing integration of perceptions abstracted from sensory transducers as the mechanism of meaning implementation. What singles it out is that the proposed mechanism is based on empirical neuroscientific evidence; since the brain is the only mechanism currently known to implement linguistic meaning, the sensible approach is to attempt to mimic it.

Generation of (String, Meaning) Pair Lists. In Freeman's architecture, any given sensory input means the sequence of multisensory perceptions in whose generation it has participated. This includes linguistically relevant acoustic input sequences, and so, for a given string, there is a corresponding multisensory perceptual sequence or, put another way, a (string, meaning) pair. Because the architecture assumes availability of environmental inputs over time, there is no need to compile a (string, meaning) pair list explicitly. Such a list generates itself over time, as the system learns from its interactions with the environment.

Application to ANN-based NLU Design. Because Freeman is specific on the biological neural mechanisms involved in the implementation of intentionality and meaning, it is possible to carry these mechanisms directly over into ANN-based language engineering. Implementation of the architecture in Fig. 2 –including, crucially, the nonlinear dynamics among modules– involves a departure from the way that most ANN-based NLP has been done over the past decade or so [60], which, in essence, has been to implement explicitly-designed input-output mappings using feedforward or recurrent multilayer perceptrons and some variant of backpropagation. The required ANN architecture will have at least the following features:

– Processing units that output pulse trains in nonlinear response to input, not just numerical values which can be interpreted as representing an average firing rate as in currently-used ANN architectures, and that fire asynchronously relative to one another.
– Modules that contain excitatory and inhibitory neurons whose interaction generates oscillatory behaviour.
– Feedforward and feedback connection among modules to enable chaotic behaviour within the modules.
– An unsupervised, local learning mechanism such as the Hebbian to allow for self-organization within modules.

Conclusion

Motivated by the lack of progress in broad-coverage natural language understanding systems based on the processing of complex syntactic structure, the foregoing discussion has proposed sequential processing as an alternative approach to implementation of string-to-meaning mapping. This is theoretically justifiable, but has substantial practical problems associated with it. Freeman's work on the neural basis of intentionality and meaning both supports the principle of sequential processing for implementation of cognitive functions, including language understanding, and also provides potential solutions to the associated problems. As such, it constitutes an attractive basis for development of an ANN-based NLU system.

Clearly, all this amounts to the merest sketch of a possible approach to NLU design. On the one hand, as Freeman himself makes clear, the multisensory perceptions that are synthesized and sequenced by the limbic system are a necessary basis for the implementation of meaning, but not a sufficient one; full linguistic meaning must involve processing of limbic output by the neocortex, and nothing has been said about that here. On the other, none of the fundamental NLU issues, such as the problem of string ambiguity and the roles of memory and reasoning in linguistic comprehension, have been addressed. Whether or not a sequential NLU system based on Freeman's ideas will be able to deal with these issues better than existing approaches, and perhaps succeed in general NLU, remains to be seen.

References

1. Allen J. (1995) Natural Language Understanding. Benjamin / Cummings, Redwood City, California
2. Bechtel W., Abrahamsen A.(1991) Connectionism and the Mind. Basil Blackwell
3. Beer R. (1995) A Dynamical Systems Perspective on Environment Agent Interactions. Artificial Intelligence 72:173-215
4. Beer R.(2000) Dynamical Approaches to Cognitive Science. Trends in Cognitive Sciences 4(3):91-99
5. Briscoe T.(1996) Robust Parsing. In [24]
6. Brooks R.(1991) Intelligence without Representation. Artificial Intelligence 47:139-59
7. Carbonell J., Hayes P. (1992) Natural Language Understanding. In [78]
8. Carroll J., Long D. (1989) Theory of Finite Automata. Prentice-Hall, Englewood Cliffs, New Jersey
9. Chalmers D. (1996) Minds, Machines, and Mathematics. Psyche 2. http://psyche.cs.monash.edu.au/psyche-index-v2.html
10. Chiel H., Beer R. (1997) The Brain has a Body: Adaptive Behaviour Emerges from Interactions of nervous System, Body, and Environment. Trends in Neurosciences 20:553-7
11. Chierchia G., McConnell-Ginet S. (2000) Meaning and Grammar: an Introduction to Semantics. 2nd ed. MIT Press, Cambridge, MA
12. Chomsky N. (1957) Syntactic Structures;. Mouton, s'Gravenhage
13. Chomsky N. (1959) On Certain Formal Properties of Grammars. Information and Control 1:91-112
14. Chomsky N. (1961) On the Notion 'Rule of Grammar'. Proceedings of the 12th Symposium in Applied Mathematics, American Mathematical Society

15. Chomsky N., Miller G. (1965) Finitary Models of Language Users', Readings in Mathematical Psychology 2, ed. Bush R., Galanter E., Luce D. John Wiley, New York
16. Christiansen M. (1992) The (Non)-necessity of Recursion in Natural Language Processing. Proceedings of The 14th Annual Conference of the Cognitive Science Society, University of Indiana
17. Christiansen M., Chater N. (1999) Toward a Connectionist Model of Recursion in Human Linguistic Performance. Cognitive Science 23:157-205
18. Church K. (1980) On Memory Limitations in Natural Language Processing. TR MIT/CS/TR-45, Massachusetts Institute of Technology
19. Church K., Ejerhed E. (1983) Finite State Parsing. Papers from the Seventh Scandinavian Conference of Linguistics, University of Helsinki
20. Cisek P. (1999) Beyond the Computer Metaphor: Behaviour as Interaction. In [62]
21. Clark A. (1989) Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing. MIT Press, Cambridge MA
22. Clark A. (1993) Associative Engines: Connectionism, Concepts, and Representational change. MIT Press, Cambridge MA
23. Clark A. (1997) Being There. Putting Brain, Body, and World Together Again. MIT Press, Cambridge MA
24. Cole R., Mariani J., Uszkoreit H., Zaenen A., Zue V. (1996) Survey of the State of the Art in Human Language Technology. Centre for Spoken Language Understanding, Oregon Graduate Institute of Science and Technology. http://cslu.cse.ogi.edu/HLTsurvey/
25. Craig E. (ed) (1998) Routledge Encyclopedia of Philosophy. Routledge, London
26. Cruse D. (2000) Meaning in Language: an Introduction to Semantics and Pragmatics. Oxford University Press, Oxford
27. Dale R., Moisl H., Somers H. (eds) (2000) Handbook of Natural Language Processing. Marcel Dekker, New York
28. Dorffner G., Prem E. (1993) Connectionism, Symbol Grounding, and Autonomous Agents. Proceedings of the 15th Annual Conference of the Cognitive Science Society
29. Dorffner G. (ed) Neural Networks and a New Artificial Intelligence. International Thomson Computer Press, London
30. Dreyfus H. (1992) What Computers Still Can't Do. 2nd ed. MIT Press, Cambridge MA
31. Finlay J., Dix A. (1996) An Introduction to Artificial Intelligence. UCL Press, London
32. Fodor J., Pylyshyn Z. (1988) Connectionism and Cognitive Architecture: a Critical Analysis. Cognition 28:3-71
33. Freeman W. (1991) The Physiology of Perception. Scientific American 264 (2):78-85
34. Freeman W. (1992) Tutorial in Neurobiology: from Single Neurons to Brain Chaos. International Journal of Bifurcation and Chaos 2:451-82
35. Freeman W. (1994) Chaos in the CNS: Theory and Practice. Flexibility and Constraint in Behavioral Systems, ed. Greenspan R., Kyriacou C. John Wiley, New York
36. Freeman W. (1994) Qualitative Overview of Population Neurodynamics. Neural Modeling and Neural Networks, ed. Ventriglia F. Pergamon, Oxford
37. Freeman W. (1999) Consciousness, Intentionality, and Causality. In [62]
38. Freeman W. (1999) How Brains Make Up their Minds. Weidenfeld and Nicholson, London

39. Freeman W., Nunez R. (1999) Restoring to Cognition the Forgotten Primacy of Action, Intention, and Emotion'. In [62]
40. Gardner H. (1985) The Mind's New Science. A History of the Cognitive Revolution. Basic Books, New York
41. Gazdar G., Pullum G. (1985) Computationally Relevant Properties of Natural Languages and their Grammars. New Generation Computing 3:273-306
42. Gazdar G., Mellish C. (1989) Natural Language Processing in LISP. Addison-Wesley, Wokingham, UK
43. Haugeland J. (1985) Artificial Intelligence: the Very Idea. MIT Press, Cambridge MA
44. Hopcroft J., Ullman J. (1979) Introduction to Automata Theory, Languages, and Computation. Addison Wesley, Wokingham, UK
45. Horwich P. (1998) Meaning. Clarendon Press, Oxford
46. Hurley S. (1998) Consciousness in Action. MIT Press, Cambridge MA
47. Johnson M. (1987) The Body in the Mind: the Bodily Basis of Meaning, Imagination, and Reason. University of Chicago Press, Chicago
48. Kelso J. A. S. (1995) Dynamic Patterns: the Self-organization of Brain and Behavior. MIT Press, Cambridge MA
49. Kramer B., Mylopoulos J. (1992) Knowledge Representation. In [78]
50. Krawer S., Tombe L. (1981) Transducers and Grammars as Theories of Language. Theoretical Linguistics 10:173-202
51. Langendoen D. (1975) Finite State Parsing of Phrase-structure Languages and the Status of Readjustment Rules in Grammar. Linguistic Inquiry 6:533-54
52. Langendoen D., Langsam Y. (1984) The Representation of Constituent Structure for Finite State Parsing. Proceedings of the Conference for Computational Linguistics, COLING 84
53. Langendoen D., Postal P. (1984b) The Vastness of Natural Languages. Blackwell, Oxford
54. Lappin S. (1996) The Handbook of Contemporary Semantic Theory. Blackwell, Oxford
55. Larson R., Segal G. (1995) Knowledge of Meaning. MIT Press, Cambridge, MA
56. Lochbaum K., Grosz B., Sidner C. (2000) Discourse Structure and Intention Recognition. In [27]
57. Luger G., Stubblefield W. (1998) Artificial Intelligence. Structures for Complex Problem Solving. Addison Wesley Longman, Harlow, UK
58. McClelland J., Rumelhart D. (1986) Parallel Distributed Processing. Explorations in the Microstructure of cognition. MIT Press, Cambridge MA
59. Moisl H. (1992) Connectionist Finite-state Natural Language Processing. Connection Science 2:67-91
60. Moisl H (2000) NLP Based on Artificial Neural Networks: Introduction. In [27]
61. Newell A., Simon H. (1976) Computer Science as Empirical Enquiry: Symbols and Search. Communications of the Association for Computing Machinery 19:113-26
62. Nunez R., Freeman W. (1999) Reclaiming Cognition. Imprint Academic, Thorverton, UK
63. Partee B., ter Meulen a., Wall R. (1990) Mathematical Methods in Linguistics. Kluwer, Boston
64. Penrose R. (1989) The Emperor's New Mind. Oxford University Press, Oxford
65. Penrose R. (1994) Shadows of the Mind. Oxford University Press, Oxford
66. Pereira F. (1996) Sentence modeling and parsing. In [24]
67. Petrick S. (1992) Parsing. In [78]

68. Poesio M. (2000) Semantic analysis. In [27]
69. Port R., van Gelder T. (1995) Mind as Motion. Explorations in the Dynamics of Cognition. MIT Press, Cambridge MA
70. Pulman S. (1986) Grammars, Parsers, and Memory Limitations. Language and Cognitive Processes 1:197-225
71. Pulman S. (1996) Semantics. In [24]
72. Pylyshyn Z. (1986) Computation and Cognition. MIT Press, Cambridge MA
73. Pylyshyn Z. (1992) Cognitive science. In [78]
74. Reich P. (1969) The Finiteness of Natural Language. Language 45:831-43
75. Roche E., Schabes Y. (1997) Finite-state Language Processing. MIT Press, Cambridge MA
76. Saeed J. (1997) Semantics. Blackwell, Oxford
77. Samuelsson C., Wiren M. (2000) Parsing Techniques. In [27]
78. Shapiro S. (1992) Encyclopedia of Artificial Intelligence. 2nd ed. John Wiley and Sons, New York
79. Shieber S. (1985) Evidence against the Context Freeness of Natural Language. Linguistics and Philosophy 8:333-43
80. Skarda C. (1999) The Perceptual Form of Life. In [62]
81. Skarda C., Freeman W. (1987) How Brains Make Chaos in order to Make Sense of the World. Behavioral and Brain Sciences 10 (2):161-95
82. Steels L., Brooks R. (1995) The Artificial Life Route to Artificial Intelligence. Lawrence Erlbaum, Hove, UK
83. Thelen E., Smith L. (1994) A Dynamic Systems Approach to the Development of Cognition and Action. MIT Press, Cambridge MA
84. van Gelder T. (1998) Cognitive Architecture: What choice do we have? In Constraining cognitive theories: issues and opinions, ed. Pylyshyn Z. Ablex, Norwood NJ