

Application of recurrent networks to reasoning

Mauricio Iza*

University of Málaga (Spain)

Abstract: The logical ordering of events facilitates the processing of narratives. If readers are going to perform an inference in order to make easier their assimilation of an upcoming text, it seems likely that they will perform an inference closely related to the causal coherence of the narrative. Therefore, an appropriate approach to this issue should determine the contexts in which readers perform elaborative inferences. In this paper, we will show that a recurrent net can be able to reasoning and interact with a contextual module.

Key words: Connectionism; hybrid architecture; discourse processing; elaborative inference.

Título: Aplicación de redes recurrentes al razonamiento

Resumen: El orden lógico de los sucesos facilita el procesamiento de las narraciones. Si los lectores van a realizar una inferencia para facilitar su asimilación del texto que viene a continuación, parece probable que realizarán una inferencia estrechamente relacionada con la coherencia causal del a narración. Por tanto, un enfoque apropiado para este problema sería determinar los contextos en los que los lectores realizan inferencias elaborativas. En este trabajo, se muestra cómo una red conexionista recurrente puede ser capaz de razonar e interactuar con información contextual.

Palabras clave: Conexionismo; arquitectura híbrida; inferencia; procesamiento discurso.

1. Introduction

As happens in conversation, text's readers haven't direct access to the author's intended meaning. Being things so, readers have to undertake inference process in order to reach an interpretation of utterances and the connections among them. These inferences turn out to be of different kinds. Some of them, for instance, bridging inferences, have to be made necessarily for a text to be coherently interpreted. Other sorts of inferences are merely elaborative, in the sense that they aren't necessary in order to link the sentences in a text, even though their conclusions are usually correct.

Inferences enrich texts with additional information retrieved from long-term memory in such a way that what finally becomes stored is the input information plus information that wasn't initially stated but inferred during the reading process. There are two main issues about inference making that text comprehension researchers have been interested in. The first one is to determine *what* (and *when*) infer-

ences are made while a text is read, and whether they are encoded into a representation of its content. The second one concerns the organization of knowledge in the long-term memory, and how cues in a text access such a knowledge.

One of the most interesting attempts to deal with these issues is the *Construction-Integration Theory* (from now on "CI") proposed by Kintsch¹. CI is an architecture for comprehension that tries to account for a broad range of comprehension phenomena in everyday language. Its name "construction-integration" reflects a central assumption about the nature of the reader-text interaction in the text comprehension process: knowledge activation is locally guided and is therefore crude and imprecise, but a contextual integration process, in the connectionist manner, can nevertheless yield a coherent text representation. CI is thus a hybrid theory: Its construction component is rule-based and symbolic, but the integration process is of a connectionist sort. Nevertheless, this model has some shortcomings. For instance it leaves untouched the activation control issue, that is to say, why we perform certain inferences, and

* **Dirección para correspondencia:** Mauricio Iza.
Departamento de Psicología Básica. Universidad de Málaga. Campus de Teatinos. 29071 Málaga (España).
E-mail: iza@uma.es

¹ See Kintsch (1988) and Kintsch (1992).

ences, and not others, in a given situation. In the control of inferential processes, it is worthwhile to assume that a constrained number of inferences are made at any time, and that such inferences turn out to be central or relevant for comprehension. If this assumption is correct, it is necessary to take into account *when* elaborative inferencing has to be activated and *what* kind of inferences can be plausible in a given moment, depending on context.

As we will see below, predictive inferences constitute an illustrative case in the general controversy over elaborations, having received a significant attention in the literature. One reason for the increasing interest in predictive inferences is the close relation they bear with what we know concerning the causal structure of narratives. The logical ordering of events facilitates the processing of narratives. If readers are going to perform an inference in order to make easier their assimilation of an upcoming text, it seems likely that they will perform an inference closely related to the causal coherence of the narrative. Therefore, an appropriate approach to this issue should determine the contexts in which readers perform elaborative inferences. In this paper, we will try to show that a hybrid connectionist reasoner (Sun, 1994) can be able to interact with a contextual module. For this purpose, it becomes necessary to encode the possible contextual links between propositions and their degree of cohesion (coherence) within the discourse. Next, as happens in CI, we will also need such a module in order to contextually integrate the conclusions drawn from inference performance.

2. Psychological evidence for elaborative inferencing

A distressing question for theories of inference performance during reading is how to characterize both such a sort of inferences and the circumstances under which they are carried out. According to *minimalism theory*², some infer-

ences, i.e. those that are merely elaborative, aren't performed when a text is initially processed, but just if they are later required. This happens, for instance, when readers are asked for answering some question by an experimenter, or simply, when some question arises in the reader's mind. On the other hand, only those inferences that establish local cohesive links are made automatically.

There is a wide agreement concerning the existence of several general kinds of inferences that can be performed in reading. Some of them are required just for preserving textual coherence -for example, referential coherence and causal coherence. Almost all researchers agree with the claim that these necessary inferences are routinely executed during the encoding of a text (Garrod *et al.*, 1990; McKoon & Ratcliff, 1986). Instead, there are other kinds of inferences that are not so required for establishing text coherence. These inferences are usually considered as *elaborative* inferences. This kind of inferences include, among others, *instrumental* inferences (i.e., inferences that would supply a typical instrument for a verb, for instance, inferring hammer for pounding a nail), *semantic* inferences (i.e., inferences that add contextually appropriate shades of meaning to a concept, for instance, emphasizing round for tomato in the girl rolled a tomato), and *predictive* inferences (i.e., inferences that would convey information about the likely outcome of a described event). Nevertheless, the frequency and conditions under which readers carry out elaborative inferences still remain as open questions. Research in the field has been usually aimed at showing whether or not some particular type of information is inferred. On the other hand, it has been argued recently that inferences are not necessarily encoded into the memory representation of a text in an all-or-none fashion, proposing instead that they can be encoded in different degrees. Several experiments examining the encoding degree of in-

ence will not be constructed. Also, this proposal suggests that elaborative inferences are often minimally encoded (McKoon & Ratcliff, 1992)

² This view suggests that if an inference is not required to give coherence to the information in a text, then the infer-

ferences about predictable events and contextually defined categories have provided evidence in this sense. For example, McKoon & Ratcliff (1989) show that inferences about the most typical tokens of such categories are encoded at a high level during reading, and that the content of the inference is made up of information relating properties of the most typical exemplar with textual information.

For this reason, McKoon & Ratcliff (1989) modified their general framework concerning minimal inferences: If a specific inference is provided by easily available general knowledge from long-term memory, then it will be built even if it is not required for coherence. However, this modification holds just for the most easily available information; other inferences such as the ones about predictable events, about default values in schema representations, and about the instruments of verbs, are usually encoded partially or not encoded at all. Given that, the crucial determinants of inference making turn out to be, on the one hand, the establishment of local coherence and the presence of knowledge readily available favouring elaboration, on the other. In any event, a comprehensive representation of a text must involve elaborative (minimalist or non-minimalist) inferences just because texts do not describe situations completely³. There are quite different views about how much elaborative inferencing goes on, and many mental models theorists do seem to believe that there is a lot of elaborative inferencing activity⁴. However their views on this issue tend to favour the minimalist approach, in the sense that they don't think many elaborative inferences are made as a matter of course (Garnham & Oakhill, 1992; Stevenson,

1993). Furthermore, the role of implicit knowledge doesn't need to be necessarily considered as inference (accessibility).

To sum up, elaborative inferences use operational processes on premises based on evidence in the input and over inferences retrieved from long-term memory (Stevenson, 1993). There are reasons to believe that elaborative inferences can be made during comprehension, but only very slowly and without being fully represented (if at all) in a mental model. For this reason, researchers employ parallel architectures in order to implement this kind of processing. It can be expected that only relevant inferences are executed, because a connectionist system constrains inferences through its interconnected network of excitatory and inhibitory links (Kintsch, 1992). However, as Sanford (1990) has pointed out, relevance itself turns out to be very difficult to define, and it is also very difficult to establish a procedure for them. In the case of elaborative inferencing, there can be made some progress in situations of strongly biasing contexts (Garrod *et al.*, 1990), due to the fact that, in these cases, inferences are constrained to those that are relevant to the given context. Finally, there are evidence favouring the hypothesis that inferences of all types (including elaborative activity) are more frequent with main characters than with secondary ones (see Garrod & Sanford, 1990). This amounts to mapping into the background knowledge that serves the perspective of the main character. A knowledge based account of comprehension, assuming restriction or preference, would be to recruit a single-perspective knowledge, which we would normally acquire as a result of our experience⁵.

³ Keefe & McDaniel's (1993) results can be taken as supporting the view that predictive inferences are temporarily drawn and then deactivated. In the presence of a more difficult material, deactivation of predictive inferences may either be delayed or prevented.

⁴ It is true that many mental model theorists have emphasized not just constructive processing, which is an essential part of text comprehension, but on-line elaborative inference making. However, on-line elaboration is not an essential part of a mental models theory of text comprehension. Hence McKoon & Ratcliff's (1992) attitude.

⁵ If elaborations are going to be made, they are more likely to be related to foregrounded concepts, because these concepts are easier to be repeated and become part of the overall text representation. On the other hand, this raises the interesting question of when and how a new entity is considered more topical than the previous topic entity.

3. The Construction-Integration Theory

The significance of the sentences in a discourse is composed of its propositional content (there is information given in the sentence that determines what can be described as the propositional content expressed by such sentences) and by what may be inferred from the rest of the text in the course of its interpretation. Therefore, comprehension results in a representation of the linguistic input (a propositional representation) that is then converted into a mental model of the text through the use of inferences based on non-linguistic knowledge (Stevenson, 1993). The processes CI focuses on are how the text propositions activate knowledge and how an integrated representation of text and knowledge is reached. The basic assumption is that knowledge is activated by means of local, associative processes, without the benefit of guidance from control structures such as frames and schemata. Thus, contextually relevant as well as irrelevant knowledge is activated, requiring a constraint satisfaction process -the integration process- in order to reject the contextually irrelevant or contradictory material. The upshot of the above processes is a text representation, which can serve as a basis for action. This representation is stored in episodic text memory, which is involved in experimental tasks such as recall and summarization. Finally, CI is concerned with how comprehending a text alters the comprehender's knowledge base or long-term memory. In a nutshell, CI deals with learning properly, as distinguished from, for instance, text recall.

The representations generated by CI in the comprehension processes are holistic structures. An associative network of elements (i.e., propositions, concepts) is constructed in order to represent the meaning of a discourse or situation as a whole, but at the same time, the elements derive part of their meaning from the structure in which they are embedded. The final network is based on principles of constraint satisfaction: The elements of the network affect

each other positively or negatively until a balanced state is achieved, in which the positively interrelated core of the network dominates and outliers are rejected.

It is expected from the theory that, as each proposition is constructed, it becomes integrated into the growing network. In most simulations, however, all propositions corresponding to a single input cycle are constructed first and then integrated simultaneously. Its outcome represents the content of working memory at the end of a cycle. As the process shifts to the next cycle, two things happen. First, the whole structure that has been generated is deleted from working memory and stored in secondary memory, as an episodic text representation. Second, the most highly activated propositions remain activated in working memory during the processing of the next input sentence. That is to say, at the end of each cycle, the focus of attention shifts to a new input, exception made in the case of the most important propositions of the previous cycle, which continue under focus.

Episodic text memory consists of the accumulated results of each processing cycle. For each processing cycle, there is a coherence matrix, summarizing the interrelations among all the elements constructed during that cycle. The outcome of comprehension modifies this coherence matrix in that the element c_{ij} -the relation between two nodes i and j in the network- is adjusted by the final action values a_i and a_j of the elements i and j so that the corresponding element in the episodic text representation becomes $e_{ij} = c_{ij} \times a_i \times a_j$. Thus, if the activation value of either element is 0 after the integration process, it is deleted from episodic text memory.

Learning in the CI model is defined in terms of the effects that the episodic text memory has over the long-term memory network. How this phenomenon has to be exactly conceptualized is still under investigation. Nevertheless, Kintsch (1992) offers two possibilities: one possibility is that the episodic memory network is simply merged with the long-term

memory network (maybe with some decay factor); another possibility is that the links in the LTM net are updated with information from episodic memory.

3.1. Limitations

A common objection usually raised against any account of language processing which assumes the continuous use of knowledge-based inference, is that drawing inferences is computationally costly and hence must be minimized if the processor is to operate as fast as it is required (Shastri & Ajjanagadde, 1993). If all possible inferences were being drawn all the time, many of these will turn out to be inappropriate by the time the end of the sentence is encountered. Consequently, the argument for carrying out less costly computation first and minimizing the use of inference is a reasonable one.

We have suggested elsewhere (Ezquerro & Iza, 1993, 1995) that attentional mechanism (i.e., focus of attention) can perform this double function, mainly in relation to inference control and anaphora resolution, giving thus more consistency to the model. In this concern, it is worthy to note that in the working memory proposed by Kintsch, the use of context -the relevance and salience of information which restricts memory retrieval- in order to disambiguate between alternative results of memory retrieval, and to select between the results of the construction processes, is closely related to the notion of a current "focus" (Grosz, 1986). In this sense, the connectionist approach seems to be closer to the *real* nature of the construction process. Connectionist models proposed for reflexive and commonsense reasoning try to perform this process automatically and within some biologically motivated values (capacity bounds). For instance, Rohwer (1992, 1993) presents an apparently novel hybrid representation of phasic (Shastri & Ajjanagadde, 1993) and tensor-product representations, which retains the desirable properties of each one.

Concerning learning, on the other hand, connectionist models for script-based text processing deal naturally with the interaction

between episodic and long-term memory. After all, people make predictions about what will happen next, on the basis of the particular episodes they have found in memory as well as on the basis of the generalisations drawn from similar experiences (Schank, 1982). In this concern, we will discuss whether distributed representations can bring some insights in order to define more *complete* connectionist architecture for comprehension.

4. A connectionist hybrid model

Finding a way of enabling connectionist systems to perform high level inference which symbolic AI approaches exhibit has been usually relayed upon localist representations. In this concern, it is well known that connectionist reasoning systems impose the constraint that just one rule may fire at time (Touretzky & Hinton, 1988). Some systems as ROBIN, which was designed for addressing the problem of ambiguity resolution using evidential knowledge (Lange & Dyer, 1989), constitutes a notable exception in this sense. However, the shortcomings of this system call for a distributed connectionist system in order to support more expressive power.

CONSYDERR (Sun, 1994) is an architecture composed of a hierarchical (acyclic) localist network (CL), and a hierarchical (acyclic) distributed network (CD). The two networks are composed of assemblies of nodes, with separate top-down and bottom-up inter-level connectivity (or path) and phasic interaction inter-level dynamics, thus giving up a three-phase cycle. The conceptual level of this structure contains primitive knowledge statements, or concepts. This level is composed of a set of nodes, or processing elements, representing these concepts in the domain. In order to express composite knowledge statements, or rules, these nodes are connected by links from antecedents to consequents of a rule. The microfeature level contains nodes, each one representing a fine-grained element (a microfeature) pertaining to the meaning field of the concepts represented in the conceptual level.

On the other hand, each node in the concept level is connected to all the relevant microfeature nodes in the microfeature level; once a concept node is activated, the related microfeature nodes will be activated subsequently from inter-level connection, and vice versa. Links in the concept level are diffusely replicated in the microfeature level by multiple links between two sets of microfeature nodes representing an antecedent and a consequent of a rule respectively.

Links between nodes, both in CL and CD, represent the rule strength measures (the link weights are defined as the corresponding rule strengths). Similarity measures, on their part, are implemented with CD representations. Sun joins the localist network (CL) with this distributed network (CD), by linking each node in CL representing a concept to all nodes representing the same concept in CD, assigning them the appropriate weights.

The cross-level links (*td* and *bu*) are moderated by a latch mechanism. The rule links in CL (labelled as *r*) are (diffusely) duplicated in CD (labelled as *h*). The interaction of the two components goes on in fixed cycles: first the latch opens for allowing the activation of CL nodes to flow into the corresponding CD nodes. Next, the two parts begin to settle down independently and simultaneously. Finally the latch opens allowing the activation of nodes in CD to flow back into CL in order to be combined with the activation of the corresponding CL nodes.

This way, CD is a connectionist network with distributed representation, corresponding roughly to reasoning at the subconceptual level. Concepts and rules are diffusely represented by sets of feature nodes overlapping each other. The overlapping amount of two sets of nodes representing two different concepts is proportional to the degree of similarity between these two concepts. In this *similarity-based representation*, units can be features, perceptual primitives, and internal goals or affect states. On the other hand, concepts are "defined" in terms of their similarity to other concepts in these primitive representations.

4.1. Performance of the model

Because of its massively parallel architecture, the model is able to perform efficiently simple forward-chaining reasoning in a parallel fashion. The results will be activated spontaneously following the activation of the initial conditions (together with other information). Nevertheless, it should be noted that Sun's system does not deal with backward chaining and goal-directed inference, although Sun & Waltz (1991), admit that goal components may be added to rules, so that reasoning turns out to be more "goal-directed". Notwithstanding, the model offers several extensions closely related to our aims. A simple extension may be to add a thresholding mechanism to the inter-level interaction, like in the intra-level case, so that a CD node (or a CL node) is activated during the top-down (or bottom-up) phase, only if the received top-down (or bottom-up) activations exceed a threshold. This serves to make the inter-level interaction non-linear. It should be noted that the similarity measure used can be easily generalized by adding different weights, functions, and/or thresholds to the inter-level links and the nodes involved; this can be useful for complex situations requiring more complex similarity measures (e.g. non-linear combination features). Even, an intermediate level of "hidden" nodes can also be added in order to provide more complex mapping capabilities.

Other interesting extension is the mechanism for taking context into account, especially for context-sensitive relationships involving rules and similarities (cf. Collins & Michalsky, 1989). The mechanism in CONSYDERR (Sun, 1994) is a feed-forward network that takes current contexts (for instance, a query) as input, and produces two types of signals for modulating feature nodes: *enable* and *disable*. The disabled feature nodes will have activations equal to 0 and therefore will not participate in similarity matching. The actual disabling occurs at the inter-level links connecting these feature nodes to their corresponding concept nodes and adjusting thus accordingly the inter-level weights. This way, similarity measures take con-

text into account and produce more accurate results. According to Sun, this feedforward network can be structured in a similar way as CL; say, each link may represent some *context rules* for deciding whether or not a feature is relevant in a given context. Anyway, more complex structures, such as backpropagation networks, are also permissible.

As remaining issues we have learning of context rules, hierarchical structuring of context, and more elaborate interactions of context and reasoning. All of them have to be accomplished within reasonable time and space constraints (Sun, 1994). As we have argued previously, a first step in this direction could be to consider this context as the focus shifting mechanism. Karen (1992) proposes a distributed recurrent model for dealing with focus shifting mechanism. This network (based both on the work of Elman (1988) and Jordan (1986)) is intended to handle locally-controlled phenomena wherein surface features and minimal semantic information act together constraining the selection of topic entities. Even though it is not intended for addressing other processes controlled by the discourse structure on a more global level, and it does not use background knowledge as a source of information, its interaction with CONSYDERR could explain how the different activation of (focused) entities affects reasoning performance, that is to say, how inference is controlled during anaphora resolution. For instance, if complex types of causal knowledge could result in certain *somehow emphasised* relationships. In any event, we should take into account as well the role that the discourse topic plays on inference during discourse processing.

5. A proposal of architecture

Global assumptions of discourse comprehension involve coherence, for instance, roles and fillers, and incrementality, that is to say, that a contextually significant representation is built up quickly and continuously during the comprehension process. In so far as many aspects of the situation are represented during dis-

course comprehension, there exist the possibility for an inference to connect what is currently under interpretation to the mental model of the discourse.

As we have seen above, a connectionist system can take into account these inferential processes during the construction of a mental model of a text. In particular, CONSYDERR has some interesting capacities for reasoning performance during text processing, mainly, in what concerns elaborative inferences. It uses rules when conditions are mentioned and manipulated explicitly; and it uses similarity matching when it turns out that none of the relevant conditions is mentioned explicitly (being thus suggestive of a holistic process). Also, we have seen that this system can be extended by adding a contextual or attentional module, implemented within a back-propagation network, that could intervene in both construction and integration processes during comprehension.

While the reasoning process, the system involves the production of a large number of dynamic facts as intermediate results that would be represented in the working memory. These facts, nevertheless, are only potentially relevant, and will remain hidden, or quickly decay, unless they turn out to be relevant for answering some "query", or for providing an explanation. In this concern, we have pointed out that only a small number of facts/entities in the working memory would happen to be relevant during text processing as they receive an attentional focus.

Within such contextual module, the possible contextual links between the propositions and their degree of coherence within the discourse have to be encoded. As happens in the CI model, this should play a role during the construction processes, as well as during the integration processes. Such a module is needed in order to construct inferences contextually (i.e., the control of activation), and also to integrate the conclusions drawn in the inference process (i.e., the inhibition of contextually irrelevant activated knowledge).

In a nutshell, by proposing the interaction between different connectionist modules, we

are suggesting connectionist architecture for comprehension. Concretely, we are suggesting how a contextual module could joint a perceptual mechanism (a connectionist parser) and the reasoning mechanism. A possible candidate for embodying this contextual module is a recurrent network. The reason is that such a network is able to maintain temporal and contextual dependencies. In particular, we have in mind the use of parallel chains, that can model feature sets with disparate times scales, and looped networks that can model long-term dependencies between hidden states.

5.1. Construction processes

Starting from the inputs coming from the perceptual module (parser), our architecture should perform construction processes in a similar way as the construction process proposed by Kintsch. Nevertheless, contrarily to Kintsch's proposal, we think that these processes involve the performance of a context module and a reasoner module. Both modules have been discussed above in relation to the role of attention as inference controllers⁶.

According to Sun, this contextual module can be integrated within a feedforward associative BP network. However, the model of context we have in mind is not just concerned with context application for restricting memory retrieval when solving interpretation problems. It also should provide a computational mechanism for gradually accumulating, and combining, the influence of different kinds of information contributing to context (such as recency of mention, subject area, and syntactic marking of focus) as the processing of a text progresses.

Concerning the construction processes within the reasoner module, we have suggested that this contextual module should have the following two ideal properties as output: (i) it should deal with the representation that a given proposition specifies according to the current

context. On the one hand, this only means dealing with the focusing degree of entities that come from the perceptual module in a more global way (i.e., a focus shifting). On the other side, however, the representation specified by the proposition itself can be transformed (i.e., into another proposition); (ii) it should decide whether or not an inference can be made and what type of inference the reasoner module can perform. For instance, if the context is so highly restricted that an elaborative inference can be extracted. In this case, the activation value of variables and arguments within propositions sent to the reasoner module should have a role in controlling inference performance. On the contrary, we could consider that the flow of activations always occurs and we should be nearer to Kintsch's knowledge activation position. Anyway, surely this is not the case: at the beginning of reading we are strongly attached to explicit information and the reasoner module could not properly enter to work.

5.2. Integration processes

As mentioned above, our contextual module should contain a structure of the elaborated text representation. This structure would be similar to the network of interrelated propositions proposed in CI. Its role would be to determine the construction processes within the reasoner module, as well as their later integration within the mental model of the text.

According to CI, relations among propositions are considered to be of several types. For instance, argument repetition (i.e., sharing the same referent) or more specific semantic relations such as causality, that has different activation values depending on its type. In this concern, our contextual network should integrate information coming from the construction processes by reaching a balanced state. As a consequence, strongly interconnected propositions will have a higher activation depending on the text context. This process should happen to be in a single step, representing the output the more salient propositions of the content of the working memory (i.e., about two) . Such highly

⁶ The reason for such a consideration lies in avoiding the *explosion* of knowledge that arises in Kintsch's construction process.

activated propositions of context would be under attentional focus when the system goes on to the next cycle (i.e., a new sentence). Thus, within our architecture, we need some way of incorporating the value of these contextual representations on the perceptual module (for instance, by using some input units of the parser for representing the previous context of the text) during the processing of the following sentence. As discussed above, such contextual information turns out useful in order to integrate new information (i.e., for solving some anaphora cases). Anyway, it seems conceptually difficult to discover some way of encoding contextual information within a network for the process of comprehension. Simplifying things, we believe that this network can be a simple associative network that receives as input a propositional representation (p), coming from the parser or the reasoner, giving as output the contextually processed proposition (p'), that should go to the reasoner module, and contextually salient information, that should go to the perceptual module. Such associative network could learn through training (experience) how things happen in the world, similarly to the connectionist models for knowledge-based text processing. Ideally, it would also be profitable for this network to be able to abstract rules, through different paths or patterns that pass over it, that should be finally established within the CL component of the reasoner system⁷.

5.3. Limits of learning

An interesting issue is how the two modules could learn to control inferences during text processing through interconnections between the reasoner (interlevel connections) and the contextual module. This way, some cases of elaborative inferences, such as conceptual, predictive and instrumental inferences, could be simulated within this composite architecture. Furthermore, incomplete processes in elaborative inferencing would depend upon its distrib-

uted representation. According to Sanford (1990), a possible solution for dealing with such processes could be that the relation was constrained but uninstantiated (e.g., that only some microfeatures were minimally activated). Nevertheless, in the case of inferences and anaphora resolution in highly constrained contexts, such an account seems a simplified solution given the real nature of inference control during text processing.

It is reasonable to suppose that, in the control of inferential processes, only a constrained number of inferences are made at any time, and that these inferences happen to be central or relevant for the discourse. Therefore, it is necessary to take into account when elaborative inferencing should be activated and what kind of inferences should be plausible in a given moment. These two features seem to be intimately related to the current context.

According to Sun (1994), this contextual module can be integrated within a feedforward associative BP network. However, the model of context we want to develop is not just concerned with context application for restricting memory retrieval when solving interpretation problems. Instead, it also should provide a computational mechanism for gradually accumulating, and combining, the influence of different kinds of information contributing to context (such as recency of mention, subject area, and syntactic marking of focus) as the processing of a text progresses. That is to say, learning of context rules, hierarchical structuring of contexts and more elaborate interaction of contexts and reasoning.

In this concern, Jordan (1992) considers how to specify targets by sets of constraints, rather than as particular vectors. Jordan distinguishes two classes of constraints: *configurational constraints*, which define regions of output space in which an output vector must lie, and *temporal constraints*, which define relationships between outputs produced at different points in time. These constraints generally arise from the knowledge of the structural or dynamical properties of the environment, and involve a two-phase learning process. In the first phase, the

⁷ Similar to the learning role played by the episodic memory in CI.

weights of an auxiliary network are adjusted in such a way as to approximate the unknown constraint function. In the second phase, the weights of the auxiliary network are held fixed and errors are transformed by propagating them backward through the auxiliary network and down into the primary network where the weights are changed. This two-phase process can be used to implement a large class of nonlinear constraints⁸. This way, implicit constraints on the solutions found by the learning algorithm are provided by the "sequential network" (Jordan, 1986) and the topology of the set of input vectors introduced during training. Usually, these constraints are considered to help to generalize from the training data to novel inputs; they can play a related role in the choice of the appropriate output vectors for indeterminate training inputs. Alternatively, if such implicit constraints yield unsatisfactory results, additional explicit configurational constraints may be added in order to optimize the result. Finally, the form of an output vector can be designed to be sensitive to a context larger than the one that defines the configurational constraints. This can be done by enlarging the scope of the optimization problem in such a way as to include other classes of constraints, such as the temporal constraints.

On the other hand, incremental learning, in a variety of forms, has loomed too large in recently published attempts to expand the horizons of connectionist knowledge acquisition. Despite some significant individual differences, a common thread affects all such approaches. What they all add to standard connectionist learning is an increased opportunity to discover transformation factors: processing episodes, which re-configure the statistical task, involved in learning a given mapping. Modularization, incremental memory expansion, batched and

sequenced training and the use of highly multi-layered architectures is all means to the achievement of this common, statistically intelligible goal. The underlying trick is always the same: to maximise the role of the achieved representations, minimizing, thus, the space of the subsequent search. Nevertheless, if we consider that inference involves the existence of some semantic justification, like discourse focus, in order to control which operations constitute a correct inference, it should be required from a connectionist reasoner system, not only to be able to represent structured knowledge, but also to be able to operate upon such representations in a systematic way, so that it would dynamically generate new representations of structured objects [learning]. Moreover, presumably human beings perform inferences jumping steps in the inference chain (due to previous experiences) thus shortening both the inferential chain length and the required time⁹.

6. Summary

It may be possible, with the help of some additional mechanism such as focus, to design a system capable of performing controlled reasoning. For this purpose, a recurrent net can keep traces of the history and respond to the current context, and it can infer context based on various cues as well. Such a system will be capable of (i) representing incoming information and making predictions based on this information by using its long-term knowledge, and (ii) generating explanations for, and testing the consistency of, incoming information by referring to its long-term knowledge.

⁹ In order to develop versatile agents that learn in situated contexts and to generalize the resulting knowledge to different environments (a task not yet achieved in the existing work for combining these two types of knowledge), Sun (in prep.) explores the possibility of learning both types of knowledge in a hybrid connectionist architecture, based on the two-level architecture proposed above. The architecture integrates reactions, rules, learning, and decision-making in a unified framework, and structures different learning components (including Q-learning and rule induction) in a synergistic way in order to perform on-line and integrated learning.

⁸ The technique of using an auxiliary network as a model of an unknown constraint function can be extended in a number of directions. Jordan & Rumelhart (1992) provide a general discussion of the role of learned internal models in supervised learning problems and show how to utilize such models in the control of an unknown dynamical system.

Such an architecture should have into account, on the one hand, evidence of incomplete processes during the performance of a range of cognitive tasks such as comprehension, learning and reasoning; involving a limited capacity system for the temporary holding and manipulation of (pragmatic) information (Baddeley, 1986; Sanford, 1990). On the other hand, a possible way to inhibit the explosion of inferences within this working memory, at least for some tasks, can be to allow attentional focusing constraint inference (e.g., feature enhancement and suppression), providing a system with

knowledge in a more structured form. However, the model proposed by Sun (1994) doesn't deal with backward inference. Thus, an unsolved question is how to decide between performing a backward or forward inference during reading, for instance, if the context module could select the type of reasoning to be performed in a given situation. This issue has not been seriously addressed yet in the literature and it is plain that we will need to combine both forward and backward reasoning. Anyway, the solution proposed by Ajjanagadde (1994) takes a step in such direction.

References

- Ajjanagadde, V. (1994). Unclear Distinctions lead to Unnecessary Shortcomings: Examining the rule vs fact, role vs filler, and type vs predicate distinctions from a connectionist representation and reasoning perspective. In *Proceedings of the AAAI-94*.
- Baddeley, A. (1986). *Working Memory*, London: Oxford University Press, Oxford Psychology Series. No. 11.
- Collins, A., & Michalsky (1989). The Logic of Plausible Reasoning: A Core Theory. *Cognitive Science*, 13, 1-49.
- Elman, J.L. (1988). Finding structure in time. *CRL Technical Report 8801*, Center for Research in Language, University of California, San Diego.
- Ezquerro, J. & Iza, M. (1993). Reflexive Reasoning, Focus Theory and Anaphora Resolution, *Logic Seminar Report LPHS-EHU-02.3*. Departamento de Lógica y Filosofía de la Ciencia, Universidad del País Vasco/EHU.
- Ezquerro, J. & Iza, M. (1995). A hybrid architecture for text comprehension: Elaborative inferences and attentional focus. *Pragmatics & Cognition*, 3(2), 247-279.
- Garnham, A., & Oakhill, J. (1992). Discourse processing and text representation from a "Mental Models" perspective. *Language and Cognitive Processes*, 7 (3/4), 193-204.
- Garrod, S., O'Brien, E.J., Morris, R.K., & Rayner, K. (1990). Elaborative Inferencing as an Active or Passive Process. *J.E.P.: Learning, Memory and Cognition*, 16, 250-257.
- Garrod, S., & Sanford, A.J. (1990). Referential processes in reading: Focusing on roles and individuals. In D.A. Balota, G.B. Flores d'Arcais and K. Rayner (Eds.), *Comprehension Processes in Reading* (pp. 465-486). Hillsdale, NJ: LEA.
- Grosz, B. (1986). A theory of discourse structure. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, pp. 432-441.
- Jordan, M. L. (1986). Serial order: A parallel distributed processing approach, *ICS Report 8604*, Institute for Cognitive Science, UCLA.
- Jordan, M.L. (1992). Constrained Supervised Learning. *Journal of Mathematical Psychology*, 36, 396-425.
- Jordan, M.L., & Rumelhart, D.E. (1992). Forward Models: Supervised Learning with a Distal Teacher. *Cognitive Science*, 16, 307-354.
- Karen, L.F. (1992). Identifying of topical entities in discourse: a connectionist approach to attentional mechanisms in language. In N. Sharkey (Ed.), *Connectionist Natural Language Processing. Reading from Connection Science*, Chapter 13, Oxford: Intellect.
- Keefe, D.E., & McDaniel, M.A. (1993). The time course and durability of predictive inferences. *Journal of Memory and Language*, 32, 446-463.
- Kintsch, W. (1992). A cognitive architecture for comprehension. In H.L. Pick, P. van den Broek and D. C. Knill (Eds.), *Cognition: conceptual and methodological issues* (pp. 143-163). Washington: American Psychological Association.
- Lange, T., & Dyer, M. G. (1989). Frame selection in a connectionist model of high-level inferencing. *Technical Report UCLA-AI-89-05*, Computer Science Department, UCLA.
- McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 12, 82-91.
- McKoon, G., & Ratcliff, R. (1989). Semantic Associations and Elaborative Inference. *JEP: Learning, Memory and Cognition*, 15, 326-338.
- McKoon, G. & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, 99, 440-466.
- Rohwer, R. (1992). A representation of representation applied to a discussion of variable binding. In M. Oaksford and G. Brown (Eds.), *Proceedings Neurodynamics and Psychology Workshop*, Bangor, Wales.
- Rohwer (1993). A representation of representation applied to a discussion of variable binding. In M. Oaksford & G. Brown (Eds.), *Neurodynamics and psychology*. New York: Academic Press.

- Sanford, A.J. (1990). On the nature of text driven inference. In D.A. Balota, G.B. Flores d'Arcais and K. Rayner (Eds.), *Comprehension Processes and Reading* (pp. 515-533). Hillsdale, NJ: LEA.
- Schank, R.C. (1982). *Dynamic memory*. Cambridge: Cambridge University Press.
- Shastri, L., & Ajanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, 16 (4).
- Stevenson, (1993). *Language, Thought and Representation*. Chichester: John Wiley & Sons.
- Sun, R. (1994). *Integrating Rules and Connectionism for Robust Commonsense Reasoning*. New York: John Wiley.
- Sun, R. (in prep.). A hybrid connectionist model for integrating reactions, rules and on-line learning.
- Sun, R., & Waltz, D. (1991). A neurally inspired massively parallel model of rule-based reasoning. In B. Soucek and the IRIS Group (Eds.), *Neural and Intelligent Systems Integration* (pp. 341-381). New York: Wiley.
- Touretzky, D.S., & Hinton, G.E. (1988). A distributed connectionist production system. *Cognitive Science*, 12, 423-466.

(Artículo recibido: 31-1-2002, aceptado: 3-3-2003)