ANALOGY MAKING IN THE TRIPLE MODEL

Maurice Grinberg
mgrinberg@nbu.bg
Vladimir Haltakov
vladimir.haltakov@gmail.com
Central and Eastern European Center for Cognitive Science, New Bulgarian University,
Montevideo 21, 1618 Sofia Bulgaria

ABSTRACT

This paper explores the analogy making capabilities of a recently introduced hybrid cognitive model, called TRIPLE. The model is compared to existing models of retrieval and mapping and its potential is demonstrated using two simple generic examples. The functionality of the model comes from the re-representation of symbolically encoded knowledge in distributed form. The distributed representation building mechanisms proposed can reflect taxonomic, relational, semantic or other structure of Long Term Memory (LTM). Retrieval and mapping are based on activation spreading, dynamic (activation dependent) similarity evaluation between target and base elements, and constraint satisfaction. The simulations show that without additional mechanisms the model performs well in simple typical examples of analogy making accounting well for structural correspondence.

INTRODUCTION

The model TRIPLE, introduced recently (see e.g. Grinberg & Kostadinov, 2008; Grinberg & Haltakov, 2009), has been created with two main goals: to have a cognitive model that can account for a wide range of cognitive processes and to achieve the computational performance and scaling needed for real-life applications.

A schematic presentation of the TRIPLE cognitive architecture is given in Figure 1. All the knowledge is organized as a semantic network in which there are concepts (including relations) and instances of concepts. The latter are used to represent the tasks (targets) and the episodic memory (bases). All knowledge, to become accessible to the architecture, must be expressed as sub-classes and/or instances of existing concepts in LTM (i.e. in symbolic form). All concepts and instances of concepts are considered to be nodes of a neural network with the task nodes as sources of activation. The LTM of the model follows the organization of LTM in the DUAL/AMBR architecture (Kokinov & Petrov, 2001). Whenever there is some input to the system activation is spread throughout LTM and relevant nodes start to become activated. The level of activation is considered to be a measure of the relevance of the activated knowledge to the input presented.

The TRIPLE cognitive model has three interconnected parts that function in parallel (see Figure 1). The Reasoning Engine (RE) is serial and coordinates the activities of the model and its modules on an event driven basis. RE is also responsible for transfer, evaluation and learning in analogy making. More details about this module can be found in Grinberg & Kostadinov, (2008).

The second module of TRIPLE is the so-called Similarity Assessment Engine (SAE) (see Figure 1). It is a connectionist engine, running as an independent parallel process (Grinberg & Haltakov, 2009). The main mechanisms in SAE are activation spreading in combination with similarity or correspondence assessment. The communication of SAE with RE is based on events related to the level of confidence for a match between the task and memory content. The description and exploration of SAE mechanisms, related to retrieval and mapping in analogy, are the focus of the present paper and will be discussed in detail later.

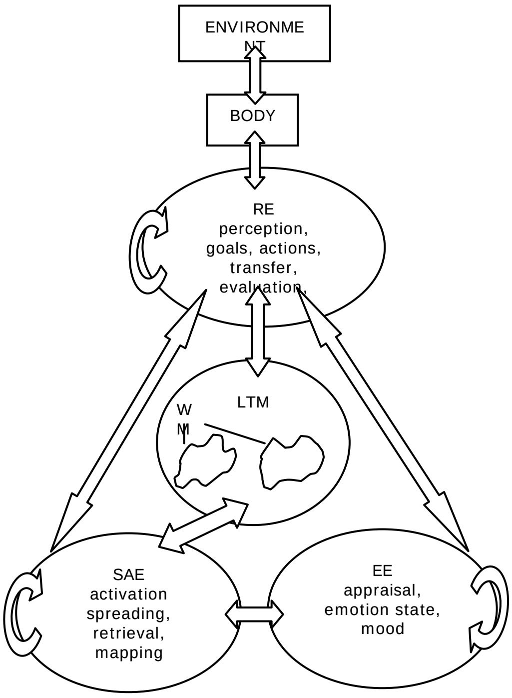

Figure 1. Structure of the TRIPLE architecture with its three main engines running in parallel: the reasoning engine (RE), the similarity assessment engine (SAE), and the emotion engine (EE) (see the text for details).

The third important part of the architecture is the Emotion Engine (EE) (see Figure 1) which is based on the FAtiMA emotional agent architecture (Dias & Paiva, 2005). This module is intended to generate emotions from a subjective appraisal of events and determine the emotional state and mood of the model. EE, similarly to SAE, runs in parallel and influences various parameters of the model like the volume of WM, the speed of processing, etc.

In this paper, we will try to demonstrate that the SAE can be regarded as a standalone distributed model of retrieval and mapping based on connectionist principles – activation spreading, distributed representations, similarity assessment of vectors, connectivity and constraint satisfaction. The possibility to have such type of processing in a model with localist symbolic encoding of knowledge lies in the re-representation mechanisms proposed here, which in our belief produce distributed representations which are much closer to the ones underlying analogy making.

We will also consider how SAE compares to existing models of retrieval and mapping using similar approaches and finally, in order to explore its capabilities, apply the model to two schematic generic cases of analogy making.

THE SIMILARITY ASSESSMENT ENGINE

General Approach

The SAE can be considered as a standalone connectionist model which receives from RE some task as input, presented in symbolic form (as instances of concepts existing in LTM). It returns similarity evaluations (or correspondence probabilities) about the task elements and some elements in LTM, typically part of some episodic knowledge. In the case of analogy making, the task will be to map a target to an appropriate base. The output of SAE is further processed by RE.

The general idea behind the SAE is related to an inverse problem like approach to symbolic representation of knowledge. Such an approach is related to the assumption that symbolic knowledge is the expression of knowledge represented in a distributed fashion and activated in order to become available. If we consider knowledge in the form of semantic networks, ontologies, or texts we can try to 'reconstruct' at least partially some of the possible distributed representation underlying such structured symbolic representations. A well known example is the Latent Semantic Analysis (LSA), in which words are given distributed representation in terms of their co-occurrences in texts, paragraphs, and sentences (Landauer and Dumais, 1997). The representation of each word is made in a space with dimension number given by the number of texts, paragraphs, or sentences taken as a basis for the representation. If we count the co-occurrences of the word with any of the other in a specific text or corpus of related texts we can have the representation of each word in terms of all other words which again results in a distributed representation. It has been shown (see Kintsch et al., 1999; Yamscar et. al., 2003) that such an analysis can give rich information about the semantic relatedness of the words and be used for retrieval from memory.

Thus, using the idea of co-occurrence in a context as a signature for relatedness, we can construct distributed representations of knowledge elements by using various relevant contexts. For instance, a taxonomy of the type 'a kitchen chair' → 'sub-class-of' → 'chair' → 'sub-class-of' → 'furniture' → etc., can be seen as a distributed representation of 'kitchen-chair' over the members of the hierarchy 'chair', 'furniture', etc., by giving a weight proportional to the distance in terms of hierarchical connections, i.e. 1, 1/2, 1/3 , etc.

Using the episodic knowledge, where instances of various concepts mix, a concept or its instance can be represented by the number of times its instances participate in a relation across all episodes in a specific role. For instance, 'chair' can be represented in terms of the relations and actions in which it participates in a specific role like 'being ON the floor', 'being seated ON', 'being BEHIND a table', etc. which will give a distributed representation of objects and instances of objects over all relations and actions. Vice versa, the relations and actions can be decomposed in terms of the elements taking part in them and the corresponding frequency. As far as the whole knowledge of the cognitive architecture is represented in the semantic net (or ontology) any distributed representation of knowledge elements, obtained by the method outlined above, will be in terms of elements of the same semantic network, e.g. objects in terms of relations and actions, objects in terms of objects, etc.

In this way, the knowledge of the model can be represented as a multi-dimensional space, spanned by all elements in LTM (plus the instances in episodic memory). Using the various types of connections among the elements, various representations can be built based on sub-spaces of the full knowledge space. These distributed re-representations of the knowledge elements can be used to evaluate the similarity or correspondence between them taking into account various criteria and goals.

The re-representation, described above, relies actually on connections among the elements of symbolically represented knowledge. These connections and their corresponding weights characterize the taxonomic, relational, episodic and semantic structure of the knowledge domain of interest. Each type of connectivity can have its own weight matrix and its own activation spreading mechanism.

The activation spreading mechanism dynamically relates the two components of SAE (the set of weight matrices and distributed representations of LTM elements). It gives the relevance of each LTM element with respect to a task (given that in TRIPLE the task elements are the sources of activation similarly to Kokinov & Petrov, 2001) and simultaneously the relevance of each component in the distributed representations of this element. It is obvious that this relevance is highly dynamic (depends on the activation patterns at a specific moment in time) and context sensitive depending on the weights of the distributed representation sets. Thus, even the similarity, based on a single distributed representation set, changes over time due to the change of the components' activations.

For instance, from a taxonomic point of view, 'kitchen-chair' and 'kitchen-table' can be highly dissimilar if only the lowest nodes in the taxonomy are activated, i.e. 'chair' and 'table', respectively. But when 'furniture' gets active they will have something in common and will become more similar as activation spreads higher in the taxonomy.

The correspondence of target and base elements being context dependent, there is always a competition between candidates for retrieval and mapping. The best mapping candidate can depend on superficial or relational similarity, or on structural correspondence constraints. In SAE, the mapping process is modeled by parallel activation spreading, similarity assessment using the distributed representation discussed above, and a constraint satisfaction process, taking into account the connectivity of the elements. No explicit structural correspondence mechanisms is assumed nor used. Our goal was to obtain structural correspondence as emergent from connectionist mechanisms and the connectivity of the knowledge elements.

In our opinion, the complementary representations approach is a more realistic approach than the canonical representation, proposed in Gentner et al. (1993), which account explicitly for the similarity between two relations or verbs. These distributed representations differ from the MAC/FAC 'content vectors' (see e.g. Forbus et. al., 1994) in that they characterize each element of the target and the base, and not the target and base as episodes as wholes. Moreover, the distributed representations, we propose, are not random, as assumed in Eliasmith & Thagard (2001) and their relation to a pure symbolic representation can be traced back. Similar methods have been used in the concept maps' and case-based reasoning literature (e.g. Leake et al., 2002).

Implementation

In the present implementation there is only one activation spreading mechanism (using the connections in LTM) and two distributed representations. The first distibuted representation is based on taxonomic relations of the type 'instance-of' and 'sub-class-of'. The second takes into account the participation and role played in relations and actions. The aim of the present implementation is to explore the capabilities of TRIPLE in retrieval and mapping. Therefore, the taxonomy-based distributed representation accounts for superficial and class belonging similarities while the relation- and action-based one accounts for similarities in the relational structure.

As described above, retrieval from memory is based on the level of correspondence of the task elements to knowledge in LTM. The measure of correspondence chosen is element-by-element similarity between target and LTM elements, measured on the basis of the two distributed representations by normalized dot product. The whole process involves three sub-processes running in parallel – activation spreading, correspondence assessment and constraint satisfaction.

Activation spreading takes place following the equation:

$$\mathbf{a}^{t} = f_{a}(\mathbf{W}_{LTM+T}\mathbf{a}^{t-1}) \qquad (1)$$

where **a** is a vector, corresponding to the activation, represented in the space of all nodes of LTM and the target nodes; f_a is a standard activation spreading function which keeps activity in the range [0,1]; $\mathbf{W}_{LTM+T}$ is a weight matrix with elements corresponding to the weights of links between the nodes in LTM (chosen equal to 1); t is the current iteration.

The distributed representations of any memory or task element are presented in matrices, called similarity matrices (denoted **X**). For instance, in the case of taxonomic similarity the rows of **X** are the concepts and instances of concepts including the target. The columns are only the LTM concepts and the matrix elements depend on the distance from a given row element following taxonomic connections. The elements to be compared are represented as rows in **X** and the normalized dot product of rows gives a correspondence or similarity measure for any two elements, weighted by the current activation:

$$S_{ij}^{t} = a_i^t a_j^t \sum_k a_k^t X_{ik} X_{jk}, \qquad (2)$$

where a^t_i and a^t_j are the activities of the elements being compared at the time t; a^t_k are the activities of the elements which form the 'basis' of the distributed representation of elements i and j. The matrix elements X_{ik} can be the weights of the connections between the corresponding elements or be evaluated using a distance function (see the discussion bellow). The inclusion of activity in the similarity evaluation makes eq. (2) dynamic and dependant only on the most active elements in LTM and in the input.

The general similarity matrix is a superposition of particular similarity matrices:

$$\mathbf{S}^t = \sum_s c_s \mathbf{S}^t_s \tag{3}$$

where

$$\sum_s c_s = 1.$$

In the present implementation we make use of two similarity matrices. The first is based on taxonomic relations like 'instance-of' and 'is-a' relations, and the second on co-participation in relations and actions taking into account the role. Under these restrictions eq. (3) reads:

$$\mathbf{S}^t = (1 - c_s)\mathbf{S}^t_{ts} + c_s \mathbf{S}^t_{rs} \tag{4}$$

where $\mathbf{S}^t$ is the resultant similarity matrix and $\mathbf{S}^t_{ts}$ and $\mathbf{S}^t_{rs}$ are the taxonomy based and the relation based similarity matrices, respectively. The upper index 't' denotes time dependence as the similarity matrices depend on time through the time dependence of the activation patterns and via their definition by eq. (2). The parameter c_s determines the mixture of the two similarity matrices. As discussed in the examples in the next section, if the taxonomy based similarity matrix dominates the retrieved base is more superficially similar to the target and when the relation based matrix dominates the base is more structurally similar to the target.

Once the similarities have been updated (at each time step), they become external input into a Constraint Satisfaction Network (CSN) based on the IAC model of McClleland & Rumelhart, (1981). Our approach and parameterization are similar to the well-known uses of CSN methods in analogy research (Thagard & Holyoak, 1995; Kokinov & Petrov, 2001). However, our implementation differs in the 'activation' used in CSN. In our case, this is the similarity (or relevance probability) that plays the role of activation (see Ming Mao et al., 2008 for a similar approach in ontology mapping research). Thus, CSN takes place in a sub-network where similarity is considered to be a special kind of activity standing for probability of mapping between two elements characterized by a so-called correspondence hypothesis. Interaction and competition take place among such correspondence hypotheses based on the constraint to have only one mapping between a target element and a base element and the requirement to have consistent mappings between connected elements, i.e. elements which are connected themselves to correspond to connected elements). In this way, each correspondence hypothesis supports any other non-competing correspondence hypothesis. In other words the correspondence hypotheses for connected elements will support each other if they do not map both to the same element (e.g. 'chair'→'table' and 'on'→'in-front-of', coming from episodes in which 'the chair is on the floor' and 'the table is in-front-of the window'). The latter is implemented by a connectivity matrix for LTM which contains weights inversely proportional to the distance between two elements, excluding taxonomy relations like 'instance-of' and 'sub-class'.

Formally, the implementation of this mechanism is done as follows. A matrix H is initialized with dimensions equal to the similarity matrix between task and LTM elements. Then it is updated, after each cycle of activation spreading and similarity calculation, by the formula:

$$H^t_{ij} = dH^{t-1}_{ij} - c_{inhib} H^{t-1}_{ij,inhib}(H^{t-1}_{ij} - H_{\min}) + (c_{excit} H^{t-1}_{ij,excit} + S^t_{ij})(H_{\max} - H^{t-1}_{ij}), \tag{5}$$

where i refers to a target element, and j to a LTM base element. Every cell (i,j) of the matrix **H** represents the relevance probability for a correspondence hypothesis (or similarity) between the elements i and j. The elements of **H** remain in the interval [-1,1] (Rumelhart & McClelland, 1981). The terms $H_{\mathrm{ij,\ inhib}}$ and $H_{ij,\ excit}$ are the inhibitory and the excitatory contributions to H_{ij}, respectively. The quantity d is a decay parameter. The parameters c_{inhib} and c_{excit} are parameters controlling the contribution of $H_{\mathrm{ij,\ inhib}}$ and $H_{ij,\ excit}$, respectively.

The quantities $H_{ij,\ \mathrm{inhib}}$ and $H_{ij,\ \mathrm{excit}}$ are calculated by using the following expressions:

$$H_{ij,inhib} = \sum_{t \neq i} H_{it} + \sum_{m \neq j} H_{mj} \tag{6}$$

$$\mathbf{H}_{excit} = \mathbf{C}_T \mathbf{H} \mathbf{C}_{LTM} \tag{7}$$

$$C_{ij} = \begin{cases} (1/n_{ij})^m, & i \neq j \\ 0, & i = j \end{cases} \tag{8}$$

and the indexes t and m refer to 'target' and 'memory'. The quantity C_{jm} and C_{it} account for the closeness in terms of number of connections of any memory element m to the memory element j, and of any target element t to the target element i. The connectivity weights are calculate using eq. (8), where n_{kl} is the distance between the two nodes k and l in terms of the number of connections to be followed in order to reach p from k (e.g. 1 – if there is one direct connection between them, 1/2 – for two connections, etc.) and m defines the metric.

EXPLORATION OF THE MODEL CAPABILITIES

In this section, two examples of simulations are presented using the SAE. There are chosen with the special purpose to explore the capabilities of the model in simple generic situations and have more of a proof of concept than of realistic applications. However, we consider them as necessary first steps in the systematic exploration of the model potential and limitations.

The examples consist of one or two targets and one or two bases. The goal is to probe different retrieval and mapping scenarios based on the change of the parameters – c_{inhib} and c_{excit}, and c_s (see eqs. (4) and (5), respectively). The variation of c_{inhib} and c_{excit} controls primarily the CSN dynamics but c_{excit} is expected to play a role in structural alignment process via the connectivity matrices which contain structural information. The excitatory connections tend also to keep the bases together and suppress the blending of bases.

However, it is the parameter c_s which controls which type of similarity measure is used taxonomy based or relation structure based. The larger the value of this parameter is, the more structural correspondence is expected in retrieval and mapping.

Due to the small LTM used and the few bases, the typical number of CSN cycles is about 10 and is kept constant in all examples. LTM includes a simplified concept structure for supporting the targets and the bases built with instances of those concepts. The decay CSN parameter d was always taken to be 0.9. The CSN parameters used are: c_{inhib} = 0.15 and c_{excit} = 0.1. In Example 1, we demonstrate how the change in the parameter c_s (see eq. (4)), changes the type of mapping preserving the structural alignment. In Example 2, we focus on the sensitivity to structural correspondence in cases with a larger complexity.

Example 1

In this example (see Figure 2), we have only one base episode in LTM with the following structure: ('table-base-1' 'on' 'floor-base-1') and ('book-base-1' 'on' 'table-base-1'). The target is similar in structure: ('note-1' 'on' 'book-1') and ('letter-1' 'on' 'note-1').

As expected (see Figure 2a), when we start to use a large relational component in the similarity matrix S – c_s = 0.9, the model maps the base to the target using the relational information involved and the connectivity matrices of the base and the target and gives the mappings: 'book-base' to 'letter', 'table-base' to 'note' and 'floor-base' to 'book'. Note that 'book' from the target is not mapped to the 'book-base' of the base. This is due to the different roles they have in the corresponding 'on' relations – in the base the 'book' is 'on' the 'table' and in the target something is 'on' the 'book' (the 'note').

With $c_s = 0.1$, as seen from Figure 2b, the mapping is dramatically changed, due to the predominance of the taxonomic part in the similarity matrix **S** (see eq. (4)). In this case, the 'book-base' from the base becomes similar to the 'book' in the target and they are mapped together while the remaining elements follow this change – the mappings are reversed and 'letter' is mapped to 'floor,' and 'note' to 'table'. Thus, the structural correspondence is preserved in this simple case, due only to the connectivity matrices and the similarity measures.

In the first case, the 'book' from the target is not mapped to the 'book-base' from the base, because they have different roles in the structure of the episodes. In the second case, the 'book-base' and 'book' are very similar and determine the mappings by imposing connectivity constraints to the rest of the target and base elements.

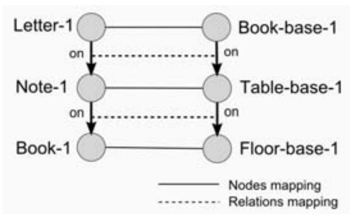

a) Alignment ($c_s = 0.9$)

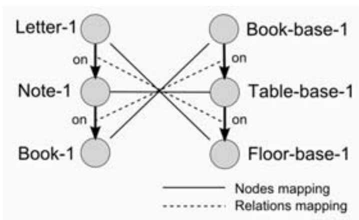

b) 'Cross' mapping ($c_s = 0.1$)

Figure 2. Mapping between base and target for two values of c_s: a) ($c_s = 0.9$) and b) ($c_s = 0.1$).

The connectivity matrices encapsulate the structure of the episode. The excitation part in the CSN (see eq. (7) and (8)) relies on the $\mathbf{C}_T$ and $\mathbf{C}_{LTM}$ matrices to determine the mutual support among the hypotheses in the target and in the base respectively. In the example, the hypothesis for similarity between ['note-1' → 'table-base-1'] is supported by the hypothesis ['letter-1' → 'book-base-1'] with strength 1/4 because the distance between 'letter-1' and 'note-1' in $\mathbf{C}_T$ is 1/2 and the distance from 'book-base-1' to 'table-base-1' in $\mathbf{C}_{LTM}$ is also 1/2 in LTM.

The excitation and inhibition mechanisms contribute essentially to the consistency and the correctness of the mappings for both cases of mapping in this example (see Figure 2, a) and b)).

In the case when the $c_s = 0.9$ (Figure 2a), the correspondence hypothesis ['book-1' → 'book-base-1'] which has a strong taxonomic similarity, has a smaller influence due to the small contribution in eq. (4) (factor '1- c_s'). The structural similarity between 'note-1' and 'table-base-1' due to their participation in the corresponding 'on' relations, are dominant in this mapping and determine the mapping displayed in Figure 2a.

In the case, with $c_s = 0.1$ (see Figure 2b), the correspondence hypothesis ['book-1' → 'book-base-1'] has a very strong similarity, due to the large weight of the taxonomy component, and decides the final mapping. The structural similarity between 'letter-1' and 'book-base-1', and between 'book-1' and 'floor-base-1' is relatively high and contributes to the winning mapping ['note-1' → 'table-base-1'], and this determines the final mapping shown in Figure 2b.

Example 2

This example describes two bases and two targets with similar objects in them, but in different spatial relations (see Figure 3).

The two bases and the two targets contain similar objects but in different spatial relations implying different relational structure. The first target (Figure3c) is structurally similar to the first base (Figure 3a) and the second target (Figure 3d) resembles structurally the second base (Figure3b). In this situation, with strong relational similarity (c_s = 0.9), the model is expected to map ‘Garage 1’ to ‘Room 1’ and ‘Garage 2’ to ‘Room 2’.

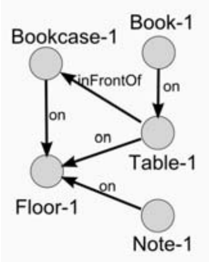

a) Base – Room 1

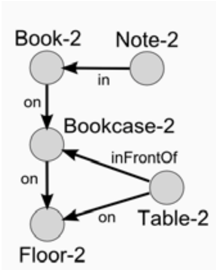

b) Base – Room 2

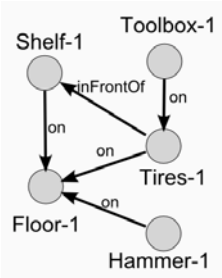

c) Target – Garage 1

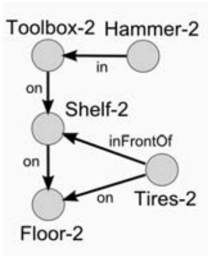

d) Target – Garage 2

Figure 3. Example 2: two base episodes a) ‘Room 1’ and b) ‘Room 2’, and two targets: c) ‘Garage 1’and d) ‘Garage 2’.

Target – ‘Garage 1’	Base – ‘Room 1’
Floor-1	Floor-base-1
Shelf-1	Bookcase-base-1
Tires-1	Table-base-1
Toolbox-1	Book-base-1
Hammer-1	Note-base-1
(Tires-1 - inFrontOf - Shelf-1)	(Table-base-1 - inFrontOf - Bookcase-base-1)
(Tires-1 - on - Floor-1)	(Table-base-1 - on - Floor-base-1)
(Shelf-1 - on - Floor-1)	(Bookcase-base-1 - on - Floor-base-1)
(Toolbox-1 - on - Tires-1)	(Book-base-1 - on - Table-base-1)
(Hammer-1 - on - Floor-1)	(Note-base-1 - on - Floor-base-1)

Table 1a. Retrieval and mapping for target ‘Garage-1’ (Figure3b).

Target – ‘Garage 2’	Base – ‘Room 2’
Floor-2	Floor-base-2
Shelf-2	Bookcase-base-2
Tires-2	Table-base-2
Toolbox-2	Book-base-2
Hammer-2	Note-base-2
(Tires-2 - inFrontOf - Shelf-2)	(Table-base-2 - inFrontOf - Bookcase-base-2)
(Tires-2 - on - Floor-2)	(Table-base-2 - on - Floor-base-2)
(Shelf-2 - on - Floor-2)	(Bookcase-base-2 - on - Floor-base-2)
(Toolbox-2 - on - Tires-2)	(Book-base-2 - on - Table-base-2)
(Hammer-2 - on - Floor-2)	(Note-base-2 - on - Floor-base-2)
(Tires-2 - on - Floor-2)	(Table-base-2 - on - Floor-base-2)
(Shelf-2 - on - Floor-2)	(Bookcase-base-2 - on - Floor-base-2)
(Toolbox-2 - on - Tires-2)	(Book-base-2 - on - Table-base-2)
(Hammer-2 - on - Floor-2)	(Note-base-2 - on - Floor-base-2)

Table 1b. Retrieval and mapping for target ‘Garage-2’ (see Figure 3d).

This is indeed the case, as seen from Tables 1a and 1b. The mappings shown in Table 1, are the mappings with highest relevance as measured by the **H** matrix (see eq. (5)). In this example, the retrieval and mapping outcome, show once again a good structural correspondence evidenced by the dependence of retrieval and mapping on the structure of the target presented.

DISCUSSION AND CONCLUSION

In the paper, the capabilities of the TRIPLE SAE as a model of analogy retrieval and mapping were presented. The model relies on multifaceted re-representations of symbolic knowledge and builds two complimentary structures: connections between knowledge elements and distributed representations. The connections matrices can be based on relations or associations between the knowledge elements or be derived on the basis of semantic relatedness (e.g. by LSA or related methods). The similarity matrices can be constructed by using the connections among the elements of LTM. In the paper, these mechanisms were illustrated by using the taxonomic connections and statistics based on participation of objects in relations. Both similarity matrices were modulated by the current activation, spread over all connections in LTM. Thus, the similarity between the target and LTM episodes is dynamic, context sensitive and evolving during mapping and retrieval.

In this model of retrieval and mapping, we have assumed that the elaborate semantic nets or ontologies are the expression of knowledge which is acquired in a connectionist, distributed way. So, this is an attempt to make an inverse problem type of analysis and derive distributed representations from the resulting localist symbolic representations.

The various possible distributed representations – taxonomic, relational, semantic-proximity based, or other – complement each other and give the power of the model. They represent different aspects of the target and LTM and could give a variety of analogies, ranging from superficial similarities up to distant analogies depending on the context and the state of the model. At the same time, this functionality is obtained just by spreading of activation and dynamic 'similarity' assessment (typically scalar products weighted by the activation of the LTM elements). If successfully applied to realistic cases, the latter would make the model computationally effective and amenable to efficient computational implementation (using standard linear algebra packages) which potentially ensures the scalability in large size memories.

While the full capabilities and limitations of this approach need to be carefully explored and compared to existing models of analogy making, in this paper, we demonstrated based on simple simulations the considerable potential of the model in analogy making. In these simple examples, the manipulation of three parameters of the model – the mixture of taxonomic and relational similarity measures, and the parameters of the constraint satisfaction network – allowed to demonstrate that the model can account for various phenomena in analogy retrieval and mapping like superficial and structural correspondence.

It should be stressed that this model is just a component of the TRIPLE and its interaction with the reasoning and emotion engines remains to be investigated but will in any case only improve its performance. Elaboration and discussion of such examples will be presented in future publications.

Acknowledgments

This work was partially supported by the project RASCALLI (FP6, EC). We would like to acknowledge the fruitful discussions with Dr. Hristo Laftchieff who contributed to an earlier version of the work presented here.

REFERENCES

Dias, J., & Paiva, A. (2005). Feeling and Reasoning: A Computational Model for Emotional Characters. EPIA. LNCS 3808, 127-140.

Eliasmith, C. & Thagard, P. (2001). Integrating Structure and Meaning: A Distributed Model of Analogical Mapping. Cognitive Science, 2001, 25, 245-286.

Forbus, K., Gentner, D., & Law, K. (1994). MAC/FAC: A model of similarity-based retrieval. Cognitive Science, 19, 141–205.

Gentner, D., Rattermann, M., & Forbus, K. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. Cognitive Psychology, 25, 524–575.

Grinberg, M., & Kostadinov, S. (2008). The Triple Model: Combining Cutting-Edge Web Technologies with a Cognitive Model in an ECA. Proc. International Workshop on Agent Based Computing V (ABC'08), Wisla, Poland

Grinberg, M., & Haltakov, V. (2009). The TRIPLE Model: Connectionist Aspects. Proc. WIRN'09, Vietri sul Mare, Italy, to appear in LNCS, Springer.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. Cognitive Science, 13, 295–355.

Kintsch, W., Patel, V. L., & Ericsson, K. A., (1999) The Role of Long-term Working Memory in Text Comprehension, Psychologia, 42, pp.186-198.

Kokinov, B. & Petrov, A. (2001). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In: Gentner, D., Holyoak, K., Kokinov, B. (eds.) The Analogical Mind: Perspectives from Cognitive Science, Cambridge, MA: MIT Press.

Landauer, T. K., & Dumais, S. (1997). A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. Psychological Bulletin 104, 211-240.

Leake, D., Maguitman, A., & Cañas, A. (2002). Assessing Conceptual Similarity to Support Concept Mapping. Mapping, Proc. of the Fifteenth FLAIRS, AAAI Press, 168-172.

Ming Mao, Yefei Peng, & Spring M. (2008). Neural Network based Constraint Satisfaction in Ontology Mapping. AAAI 2008, 1207-1212.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. Psychological Review , 88, 375-407.

Ramscar, M., & Yarlett, D. (2003). Semantic grounding in models of analogy: an environmental approach. Cognitive Science, 27, 41–71.