Combining Representational Domains for Computational Creativity

Agnese Augello, Ignazio Infantino, Giovanni Pilato, Riccardo Rizzo, Filippo Vella
Istituto di Calcolo e Reti ad Alte Prestazioni
Consiglio Nazionale delle Ricerche
sede di Palermo, ed. 11
e-mail: name.surname@cnr.it

Abstract

The paper describes a combinatorial creativity module embedded in a cognitive architecture. The proposed module is based on the focus of attention model proposed by and is implemented using Self Organising Map (SOM) neural networks.

Introduction

Creativity is mainly perceived as a high level cognitive characteristic, which should always be referring to a conceptual space, whether it is conceived to explore or to transform such space (Boden 2009). One of the components of creativity is an associative memory capable of restoring an incomplete sensory input stimulus by adjusting focus of attention.

A cognitive model for creativity based on the ability of adjusting focus of attention has been proposed in (Gabora 2002). According to this model a variable focus of attention, while pointing the basic idea, also collects other concepts that are parts of the stream of thought. The focus of attention can be considered as a basic idea, a *framework* that drives the creative process which is connected to the analytical mode of thought.

At the same time another basic component of the cognitive model proposed by Gabora is the associative memory. By means of associations between different concepts and completion mechanisms, new and surprising results can emerge (Bogart and Pasquier 2013). This kind of creative process can be bound to the process that Boden calls combinatorial creativity which is related to making unusual combinations, consciously or unconsciously generated, of familiar ideas (Boden 2009).

The Arcimboldo painting can be a good example for clarifying what we intend.The painting of a human figure presumes a very precise *framework* that is constituted by figure details, as nose, eyes, lips, rules and relative positioning and all the other details that made a human figure. The attention focus is what we use to “navigate” on the *framework*, what is pointing at the details of the figure, that can be substituted with elements belonging to another domain (as in the painting in fig. 1) exploiting the associative memory. We believe that, during the creative process of imagining the painting, the attention is relaxed and other images, searched in another domain, come in mind and take the place of the original parts of the human figure.

We consider completion operation in a very large meaning. The basic point of the combinatorial creativity is to mix together parts coming from different sources. In this sense completion is a way to enrich a *framework* with new items in order to obtain new combinations.

In our opinion it is possible to have robust fusion algorithms and completion through the combination of various models of neural networks: an example of such an approach is described in (Thagard and Stewart 2011) that allows emphasising associations useful to generate creative ideas by simple vector convolution. The importance of associative mechanisms is also underlined by neurobiological models of creativity, many of which are based on the simultaneous activation and communication between brain regions that are generally not strongly connected (Heilman, Nadeau, and Beversdorf 2003).

In this paper we illustrate an approach aimed at supporting the execution of an artificial digital painter (Augello et al. 2013b) (Augello et al. 2013a). The approach is exploited by the Long Term Memory (LTM) module of the cognitive architecture presented in (Augello et al. 2013b) and reported in fig. 2. The proposed approach is based on a multilayer mechanism that implements an associative memory based on Self Organizing Maps (SOMs) (Kohonen, Schroeder, and Huang 2001) and it is capable to properly mix elements belonging to different domains.

Figure 1: A detail from Spring (1563), an Arcimboldo painting (Image from Wikipedia).

Architecture

In (Augello et al. 2013b) we defined the mechanisms to support creativity in a cognitive framework. In this work we use the same architecture (see fig. 2) but we adopt a new version

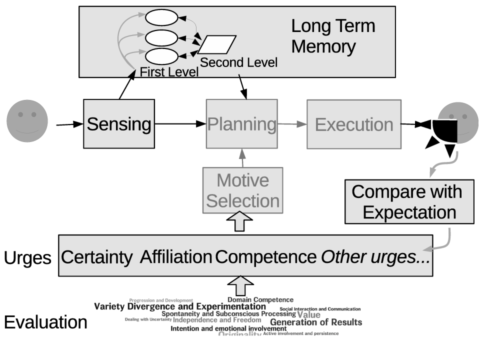

Figure 2: The general cognitive framework used for the proposed system. Light grey blocks are neglected in this implementation.

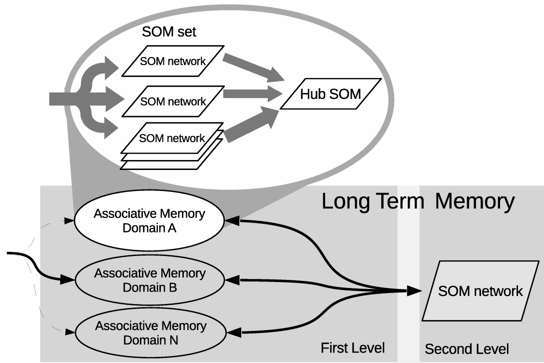

Figure 3: The overall schema of the proposed architecture for the Long Term Memory module (LTM).

of LTM (Long Term Memory) that implements an associative mechanism described in details below.

As said before, one of the basic components is an associative memory capable of restoring an incomplete sensory input stimulus.

Completion is guided by context: when we interpret fuzzy or confused handwritten characters, we use associations with memorised handwritten characters, then we complete or rebuild the input, so that the most common association are made using objects of the same context. Objects coming from the same domain are probably represented by the same features and share the same concept space that was described in Gärdenfors (Gärdenfors 2004). Associations can also involve objects from different contexts in a more “creative” way. In this case the original context is discarded and objects come from different domains.

According to these considerations we have built a multilayer mechanism that allows to connect memory locations related to a single domain. We have also built another layer that is used to connect memory locations with a more general association mechanism that allows to make associations that go beyond the domain. This second upper layer will be used when the original domain is discarded, for example when we want to find other solutions or we want to mix different domains. The kind of associations made at the second level will be the associations made when the focus of attention is relaxed and associative connections can be made even outside of a specific domain. The structure we propose is represented in fig. 3.

Input from sensors are sent to the proper domain at the first level and they are memorised or completed when necessary. The second level contains the associations among different domains that will be further explained in the following paragraphs.

The associative memory module that we propose is inspired by the work in (Morse et al. 2010) and is implemented using a Self Organising Map (SOM) neural networks (Kohonen, Schroeder, and Huang 2001).

Self Organising Maps are neural networks constituted by a single layer of neural units usually organised in a 2D grid. After a successful training phase each neural unit ideally approximates the centroid of an input pattern cluster and the neighbour units represent similar values. This way each neural unit corresponds to a sort of average pattern for a cluster of inputs.

The architecture proposed in (Morse et al. 2010), is made by multiple SOM, each one receiving inputs from a different sensory modality. In our architecture the SOM array, in the upper part of fig 3, receives inputs from different features extracted from the same sensory input, so that a SOM of the set can have colour features from image, another image boundaries, another one texture information and so on. The values of the SOMs are collected by the *hub–SOM* that synthetically represents the object gathering the representations of the different SOMs. This process is sketched in fig. 4, where different features are substituted by different parts of the image.

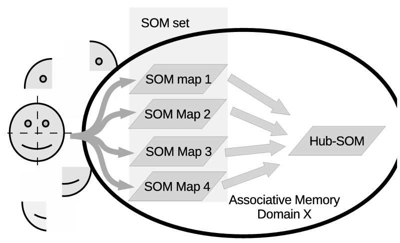

Figure 4: The associative domain memory training.

While the SOM set and the hub–SOM constitute the associative module for a domain there is also another SOM, named second level SOM, where the association among different domains takes place.

The information from the domain modules, in this second level, are represented using more general features. For example if a domain is used to memorise images of trees and one of the SOM in the array in fig.3 memorises the shape of the leafs, the second level SOM can use the dimension of the bounding box[1] as a feature. When we want to mix together objects from other domains we can consider objects that have the same bounding box. The substitution will

[1] the bounding box is the rectangle surrounding an image detail

be driven by the second level SOM whose aim is to faithfully reflect the general structure of the image. A substitution according the bounding box dimensions is a simple criterion but a more general set of features that could also be employed. This second level SOM will implement the spreading of the attention focus because it will mix objects from different domains and group them just considering very rough characteristics.

The next two subsections will explain how we can implement the effects of a variable focus of attention: with a narrow focus we can obtain simple completion inside the same domain while with the spread of the focus we can recover objects from different domains.

Completion in the same domain

Completion in the same domain is the simplest form of completion. For example, let us assume that a domain is trained to memorise simple images, and imagine that, inside this domain there are SOMs that memorise very specific parts of the image: we can think that each SOM memorises a quadrant of the image or, when representing faces, segments of human faces. In this case the basic components would be eyes, lips, noses and so on, memorised along with their positions, in different SOMs. The hub–SOM takes into account all the positions of the components. This is sketched in fig. 4.

If a part of an image is missing, only some of the SOMs can recall the corresponding memory locations and help to reconstruct the memorised image: one, or more, SOMs will not answer because they do not have any input. The hub–SOM accomplishes the task to recall the necessary memory locations from the SOMs that do not have any input, in order to put together all the pieces of the image.

This procedure is depicted in fig.5: the missing piece of image causes a failure of recalling in SOM Map 4, so that the hub–SOM, containing the reference of the whole, outputs the address of the location of the SOM Map 4 and recall the missing piece.

Completion in a different domain

When completion is obtained using “parts” or memories that are outside the domain of the original image, or input, we are making an association that is not causal. This can happen when the recalled part is used to obtain memory contents from other, different domains. In this case the associations are the ones memorised in the second layer SOM, i.e. an association that corresponds to features of different kind.

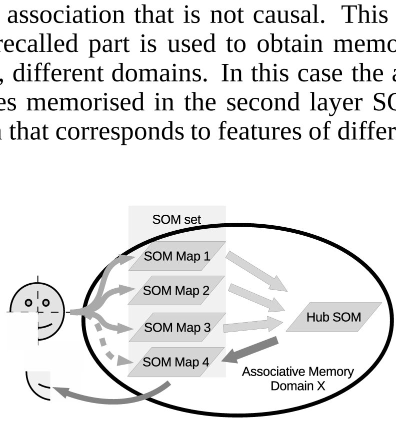

Figure 5: The completion procedure in the domain

In fig. 6 the whole process is sketched : the missing part is recalled as said before; however, in this case, it is not sent to the output but it is sent to the second level SOM where it is used for recalling objects from different domains.

The recovered information is used as a reference in order to obtain the missing part that is sent to the second level SOM. This signal excites a unit of the second level SOM and its output is sent back to all the associative memory of the other domains. Each domain answers with a list of the excited units that point out to a set of signal corresponding to the memorised objects. As indicated in fig 6 all these objects are proposed as substitution of the missing part. At this point the completion proposed from the original domain is again used as a reference: all the proposed substitution are compared to the original completion and the most similar one is chosen as a substitute. This mechanism is implemented in the box “Implementation with Expectation” in fig 6.

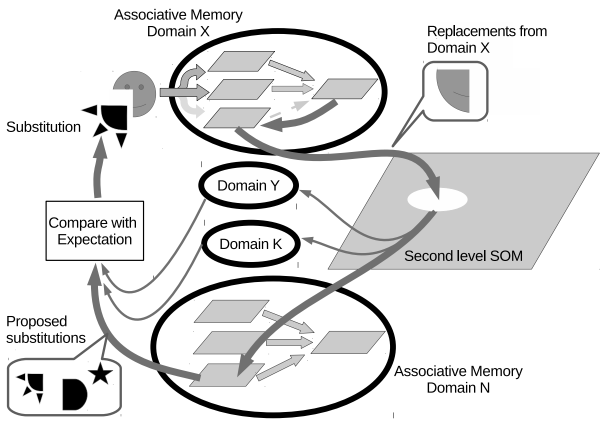

Figure 6: The completion procedure outside the domain.

Some Experimental Results and Conclusive Remarks

The experiments were mainly performed to evaluate the effectiveness mechanism of the replacement of some parts of the images by the associative memory previously described. We have chosen a *face domain* that allows an immediate recognition and a *leaves and flowers domain* in order to resemble the effect of the Arcimboldo style. A sample of the images in our dataset is given in fig.7. The system was trained using 113 grey-scale images of faces and 100 images of leaves and flowers. Each image is 100×100 pixel in order to maintain a manageable size of the neural architecture.

In order to reproduce the completion mechanism and to partially simulate the mechanism of focus of attention, each quarter of the image has been memorised in a different map of the array of SOMs (see fig.4). We have tried a quad tree decomposition and the learning process described above. An example is reported in fig.9.

Each SOM in the array has a size of 20×20 units and is trained with segments in the same position of the image using the fast training procedure described in (Rizzo 2013),

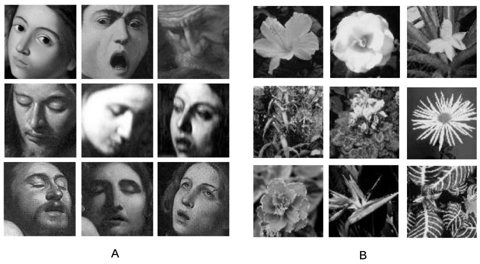

Figure 7: An example of the images in the faces domain (A) and leafs and flowers (B) domain.

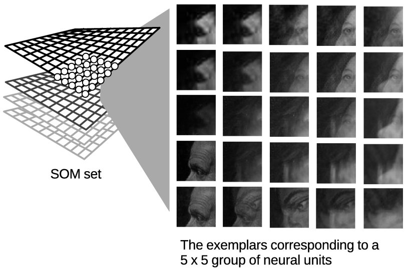

Figure 8: The map in the array of SOMs after the training are used as memory units.

and the result is shown in fig. 8. At the end of the process it is possible to train the hub–SOM submitting the images of the training set to the SOM array; for each SOM we will get the two digits coordinates on the neural units array, of the most similar exemplar (often called best matching unit or b.m.u.). These coordinates are submitted to the hub–SOM that learns this 8 digits image coding and, after training, will be able to rebuild the correct coding for each image.

This kind of representation is too precise to be used also at higher level, were we want to "mix together different things". At higher levels we want a representation that captures just some of the characteristics of the images, for example colour masses, boundaries and shapes, and so on. For this reason we used the Haar and Gabor features, which contain less information.

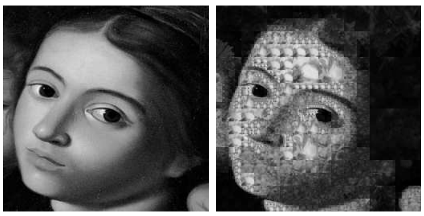

Figure 9: Final artwork obtained by our approach

Conclusion

The preliminary experimental results show that the proposed associative memory module is promising for the implementation of a sort of combinatorial creativity mechanism. Future works will regard the modelling of artist's behaviour and motivation, the choice of domains during the completion process, and the evaluation of both creative process and produced artworks, according to the literature works (Pease and Colton 2011) (Colton and Wiggins 2012) (Jordanous 2012).

Acknowledgment

This work has been partially supported by the PON01_01687 - SINTESYS Research Project.

References

Augello, A.; Infantino, I.; Pilato, G.; Rizzo, R.; and Vella, F. 2013a. Binding representational spaces of colors and emotions for creativity. *Biologically Inspired Cognitive Architectures* 5:64–71.

Augello, A.; Infantino, I.; Pilato, G.; Rizzo, R.; and Vella, F. 2013b. Introducing a creative process on a cognitive architecture. *Biologically Inspired Cognitive Architectures* 6:131–139.

Boden, M. 2009. Computer Models of Creativity. *AI Magazine* 23–34.

Bogart, B., and Pasquier, P. 2013. Context Machines: A Series of Situated and Self-Organising Artworks. *Proceedings of the fourth conference on Creativity & cognition - C&C '02* 46(2):114–122.

Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In *ECAI'12*, 21–26.

Gabora, L. 2002. Cognitive mechanisms underlying the creative process. *Proceedings of the fourth conference on Creativity & cognition - C&C '02* 126–133.

Gärdenfors, P. 2004. *Conceptual spaces: The geometry of thought.* The MIT Press.

Heilman, K. M.; Nadeau, S. E.; and Beversdorf, D. O. 2003. Creative innovation: possible brain mechanisms. *Neurocase* 9(5):369–379.

Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4(3):246–279.

Kohonen, T.; Schroeder, M. R.; and Huang, T. S., eds. 2001. *Self-Organizing Maps*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 3rd edition.

Morse, A. F.; de Greeff, J.; Belpeame, T.; and Cangelosi, A. 2010. Epigenetic Robotics Architecture (ERA). *IEEE Trans. on Autonomous Mental Development* 2(4):325–339.

Pease, A., and Colton, S. 2011. Computational creativity theory: Inspirations behind the face and the idea models. In *ICCC 2011*.

Rizzo, R. 2013. A new training method for large self organizing maps. *Neural Processing Letters* 37(3):263–275.

Thagard, P., and Stewart, T. C. 2011. The aha! experience: Creativity through emergent binding in neural networks. *Cognitive science* 35(1):1–33.