Acting on Conceptual Spaces in Cognitive Agents

Agnese Augello[3], Salvatore Gaglio[1,3], Gianluigi Oliveri[2,3], and Giovanni Pilato[3]

[1] DICGIM- Università di Palermo
Viale delle Scienze, Edificio 6 - 90128, Palermo - ITALY
[2] Dipartimento di Scienze Umanistiche - Università di Palermo
Viale delle Scienze, Edificio 12 - 90128, Palermo - ITALY
[3] ICAR - Italian National Research Council
Viale delle Scienze - Edificio 11 - 90128 Palermo, Italy
{gaglio.oliveri}@unipa.it
{augello,pilato}@icar.pa.cnr.it

Abstract. Conceptual spaces were originally introduced by Gärdenfors as a bridge between symbolic and connectionist models of information representation. In our opinion, a cognitive agent, besides being able to work within his (current) conceptual space, must also be able to 'produce a new space' by means of 'global' operations. These are operations which, acting on a conceptual space taken as a whole, generate other conceptual spaces.

1 Introduction

The introduction of a cognitive architecture for an artificial agent implies the definition of a conceptual representation model. Conceptual spaces, used extensively in the last few years [1] [2] [3], were originally introduced by Gärdenfors as a bridge between symbolic and connectionist models of information representation. This was part of an attempt to describe what he calls the 'geometry of thought'.

If, for the sake of argument, we accept Gärdenfors paradigm of conceptual spaces, and intend to avoid the implausible idea that a cognitive agent comes with a potentially infinite library of conceptual spaces, we must conclude that a cognitive agent, besides being able to work within his (current) conceptual space, must also be able to 'produce a new space' by means of 'global' operations. These are operations which, acting on a conceptual space taken as a whole, generate other conceptual spaces.

We suppose that an agent acts like an experimenter: depending on the particular problem he has to solve, he chooses, either consciously or unconsciously, what to observe and what to measure. Both the environment and the internal state of the agent, which includes his intentions and goals, affect the manner in which the agent perceives, by directing the focus of its measurements on specific objects.

In this work we focus on operations that can be performed *in* and *on* conceptual spaces in order to allow a cognitive agent (CA) to produce his conceptual representation of the world according to his goals and his perceptions.

In the following sections, after a background on Conceptual Spaces theory, we introduce such operations and we discuss an example of the way they come to be applied in practice.

2 Conceptual spaces

In [4] and [5] we find a description of a cognitive architecture for modelling representations. This is a cognitive architecture in which an intermediate level, called 'geometric conceptual space', is introduced between a linguistic-symbolic level and an associationist sub-symbolic level of information representation.

According to the linguistic/symbolic level:

> Cognition is seen as essentially being *computation*, involving symbol manipulation. [4]

whereas, for the associationist sub-symbolic level:

> Associations among different kinds of information elements carry the main burden of representation. *Connectionism* is a special case of associationism that models associations using artificial neuron networks [4], where the behaviour of the network as a whole is determined by the initial state of activation and the connections between the units [4].

Although the symbolic approach allows very rich and expressive representations, it appears to have some intrinsic limitations such as the so-called 'symbol grounding problem,' [4] and the well known A.I. 'frame problem'.[5] On the other hand, the associationist approach suffers from its low-level nature, which makes it unsuited for complex tasks, and representations.

Gärdenfors' proposal of a third way of representing information exploits geometrical structures rather than symbols or connections between neurons. This geometrical representation is based on a number of what Gärdenfors calls 'quality dimensions' whose main function is to represent different qualities of objects such as brightness, temperature, height, width, depth.

Moreover, for Gärdenfors, judgments of similarity play a crucial role in cognitive processes. And, according to him, it is possible to associate the concept of distance to many kinds of quality dimensions. This idea naturally leads to the conjecture that the smaller is the distance between the representations of two given objects the more similar to each other the objects represented are.

---

[4] How to specify the meaning of symbols without an infinite regress deriving from the impossibility for formal systems to capture their semantics. See [6].

[5] Having to give a complete description of even a simple robot's world using axioms and rules to describe the result of different actions and their consequences leads to the 'combinatorial explosion' of the number of necessary axioms.

According to Gärdenfors, objects can be represented as points in a conceptual space, and concepts as regions within a conceptual space. These regions may have various shapes, although to some concepts—those which refer to natural kinds or natural properties[6]—correspond regions which are characterized by convexity.[7]

For Gärdenfors, this latter type of region is strictly related to the notion of prototype, i.e., to those entities that may be regarded as the archetypal representatives of a given category of objects (the centroids of the convex regions).

3 A non-phenomenological model of Conceptual Spaces

One of the most serious problems connected with Gärdenfors' conceptual spaces is that these have, for him, a phenomenological connotation. In other words, if, for example, we take, the conceptual space of colours this, according to Gärdenfors, must be able to represent the geometry of colour concepts in relation to how colours are given to us.

Now, since we believe that this type of approach is bound to come to grief as a consequence of the well-known problem connected with the subjectivity of the so-called '*qualia*', e.g., the specific and incommunicable quality of my visual perception of the rising Sun or of that ripe orange etc. etc., we have chosen a non phenomenological approach to conceptual spaces in which we substitute the expression 'measurement' for the expression 'perception', and consider a cognitive agent which interacts with the environment by means of the measurements taken by its sensors rather than a human being.

Of course, we are well aware of the controversial nature of our non phenomenological approach to conceptual spaces. But, since our main task in this paper is characterizing a rational agent with the view of providing a model for artificial agents, it follows that our non-phenomenological approach to conceptual spaces is justified independently of our opinions on qualia and their possible representations within conceptual spaces

Although the cognitive agent we have in mind is not a human being, the idea of simulating perception by means of measurement is not so far removed from biology. To see this, consider that human beings, and other animals, to survive need to have a fairly good ability to estimate distance. The frog unable to determine whether a fly is 'within reach' or not is, probably, not going to live a long and happy life.

Our CA is provided with sensors which are capable, within a certain interval of intensities, of registering different intensities of stimulation. For example, let us assume that CA has a visual perception of a green object h. If CA makes of the measure of the colour of h its present stereotype of green then it can, by means

[6] Actually, we do not agree with Gärdenfors when he asserts that:

> Properties... form a special case of concepts. [4], chapter 4, §4.1, p. 101.

[7] A set S is *convex* if and only if whenever $a, b \in S$ and c is between a and b then $c \in S$.

of a comparison of different measurements, introduce an ordering of gradations of green with respect to the stereotype; and, of course, it can also distinguish the colour of the stereotype from the colour of other red, blue, yellow, etc. objects. In other words, in this way CA is able to introduce a 'green dimension' into its colour space, a dimension within which the measure of the colour of the stereotype can be taken to perform the rôle of 0.

The formal model of a conceptual space that at this point immediately springs to mind is that of a metric space, i.e., it is that of a set X endowed with a metric. However, since the metric space X which is the candidate for being a model of a conceptual space has dimensions, dimensions the elements of which are associated with coordinates which are the outcomes of (possible) measurements made by CA, perhaps a better model of a conceptual space might be an n-dimensional vector space V over a field K like, for example, $\mathbb{R}^n$ (with the usual inner product and norm) on $\mathbb{R}$.

Although this suggestion is very interesting, we cannot help noticing that an important disanalogy between an n-dimensional vector space V over a field K, and the 'biological conceptual space' that V is supposed to model is that human, animal, and artificial sensors are strongly non-linear. In spite of its cogency, at this stage we are not going to dwell on this difficulty, because: (1) we intend to examine the 'ideal' case first; and because (2) we hypothesize that it is always possible to map a perceptual space into a conceptual space where linearity is preserved either by performing, for example, a small-signal approach, or by means of a projection onto a linear space, as it is performed in kernel systems [7].

4 Operating *in* and *on* Conceptual spaces

If our model of a conceptual space is, as we have repeatedly said, an n-dimensional vector space V over a field K, we need to distinguish between operating *in* V and operating *on* V. If we put $V = \mathbb{R}^n$ (over $\mathbb{R}$), then important examples of operations *in* $\mathbb{R}^n$ are the so-called 'rigid motions', i.e. all the functions from $\mathbb{R}^n$ into itself which are either real unitary linear functions[8] or translations.[9] Notice that if f is a rigid motion then f preserves distances, i. e. for any $v, w \in \mathbb{R}^n$, $d(v, w) = d(f(v), f(w))$. Examples of rigid motions which are real unitary linear functions are the θ-anticlockwise rotations of the x-axis in the x, y-plane.

To introduce operations *on* V, where V is an n-dimensional vector space over a field K, we need to make the following considerations. Let CA be provided with a set of measuring instruments which allow him to perform a finite set of measurements $M = \{m_1, \ldots, m_n\}$, and let $\{V_i\}_{i \in I}$ be the family of conceptual spaces— finite-dimensional vector spaces over a field K—present in CA's library.

[8] A linear function $f : \mathbb{R}^n \to \mathbb{R}^n$ is *real unitary* if and only if it preserves the inner product, i.e. for any $v, w \in \mathbb{R}^n$, we have $f(v) \cdot f(w) = v \cdot w$.

[9] The function $t : \mathbb{R}^n \to \mathbb{R}^n$ is a *translation* if and only if there exists a $v \in \mathbb{R}^n$ such that, for any $w \in \mathbb{R}^n$, we have $t(w) = w + v$.

If we assume that c is a point of one of these conceptual spaces, the coordinates $c_1, c_2, \ldots c_n$ of c represent particular instances of each quality dimension and, therefore, derive from the set of n measures performed by the agent on the subset of measurable elements. We, therefore, define two operations $\times$ and π *on* $\{V_i\}_{i \in I}$ such that: (1) $\times$ is the *direct product* of vector spaces, that is:

1. $V_i \times V_j = \{< v_i, v_j > | \ v_i \in V_i \ \text{ and } \ v_j \in V_j\}$;
2. for any $< v_{i,1}, v_{j,1} >, < v_{i,2}, v_{j,2} > \in V_i \times V_j$, we have: $< v_{i,1}, v_{j,1} > + < v_{i,2}, v_{j,2} > = < v_{i,1} + v_{i,2}, v_{j,1} + v_{j,2} >$
3. for any $k \in K$ and $< v_i, v_j > \in V_i \times V_j$, we have that: $k < v_i, v_j > = < kv_i, kv_j >$;

clearly, $V_i \times V_j$, for any $i, j \in I$, is a vector space, and

$$\dim (V_i \times V_j) = \dim V_i + \dim V_j;$$ [10]

and (2) π_i is the *projection* function onto the i-th coordinate space, i.e. $\pi_i(V_i \times V_j) = V_i$ and $\pi_j(V_i \times V_j) = V_j$, for $i, j \in I$. Obviously, we have that $\pi_i(V_i \times V_j)$ and $\pi_j(V_i \times V_j)$ are vector spaces, and that

$$dim\ \pi_i(V_i \times V_j) = dim\ V_i.$$

Now, with regard to the importance of the operator $\times$, consider that if we have the vector space $\mathbb{R}^3$, over the field $\mathbb{R}$, whose dimensions do not include time, we cannot then form the concept of velocity; and if the dimensions of the vector space $\mathbb{R}^3$, over the field $\mathbb{R}$, do not include colour, we cannot form the concept of red block. It is by producing, by means of $\times$, the right type of finite dimensional vector space that we make possible to formulate within it concepts such as velocity, red block, etc. The $\times$ operation on finite vector spaces has, to say it with Kant, an ampliative function. The relevance of π is, instead, all in its analytic rôle of explicating concepts by drawing attention to the elements belonging to a given coordinate space.

At each moment CA, instead of relying on the existence of a potentially infinite library of conceptual spaces, if necessary, individuates new dimensions following the procedure briefly illustrated on p. 3-4, and builds the current conceptual space suitable for the tasks that it has to accomplish by performing operations *on* the conceptual spaces which are already available.

5 A case study

We assume that CA is located on and can move around the floor of a room where objects of different type, size and color may be found. His sensors allow CA to obtain information concerning some of the characteristics of the surrounding environment and of some of the objects in it. When CA moves around the room, the perspective from which he views the objects present in the environment changes.

[10] $dim(V_i)$ is the dimension of the vector space V_i.

Of course, on the assumption that CA can tell from its receptors whether a given point of the floor of the room on which he is focussing is 'occupied' or not, it follows that CA is capable of performing tasks — like 'coasting around' the objects placed on the floor of the room — which do not require the use of conceptual spaces. But, on the other hand, there are tasks which require the use of systems of representation, such as conceptual spaces, which allow CA to build faithful representations (models) of the environment, etc.

Every time CA focuses its attention on something, CA identifies, *via* his receptors, the quality dimensions necessary for the representation of the object of interest and creates a specific current conceptual space individuating the regions (concepts) belonging to it.

To see this, assume that on the floor of the room where CA is there are two discs D_1 and D_2, and that CA's task consists in comparing in size D_1 with D_2. The initial current conceptual space V_0 of CA can be the vector space $\mathbb{R}^2$ (on $\mathbb{R}$) with the conceptual structure $\mathcal{C}_0$. CA is at the origin of the two axes of V_0 and the conceptual structure $\mathcal{C}_0$ associated to V_0 is $\mathcal{C}_0 = \{$FRONT (F), BACK (B), LEFT (L), RIGHT (R)$\}$. Here F, B, L, R are the *primitive* regions of V_0. (From now on, instead of talking about the conceptual space V_0 with structure $\mathcal{C}_0$, we shall simply consider the conceptual space $(V_0, \mathcal{C}_0)$.)

Note that the terms we use to refer to the *primitive* regions of V_0 are just a *façon de parler*, i.e., our way of describing the conceptual structure of the conceptual space of CA. In fact, we assume that the conceptual activity of CA is sub-linguistic.

CA can perform algebraic operations internal to the conceptual space which are mainly set operations given that the regions of V_0 are sets of points of V_0. The elementary operations defined on such regions are: $\cup, \cap, C_A^B$ (where $A \subseteq B$ and A and B are regions). Such operations applied to our primitive regions F, B, L, R allow us, for example, to individuate regions of particular importance such as the y-axis which can be characterized as the set of points $y \in C_{L\cup R}^{V_0}$, the x-axis as the set of points $x \in C_{F\cup B}^{V_0}$, the minimal region $\{0\}$, where 0 is the origin of the x and y axes as $C_{L\cup R}^{V_0} \cap C_{F\cup B}^{V_0} = \{0\}$, $F \cap R = \{(x, y) \mid 0 < x \text{ and } 0 < y\}$ (the first quadrant of $\mathbb{R}^2$), $L \cap R = \emptyset$, etc. As we have already seen at the very beginning of §3, another important class of operations internal to $(V_0, \mathcal{C}_0)$ are what we there called 'rigid motions'.

At this point we need to notice that $(V_0, \mathcal{C}_0)$ is a genuine conceptual space irrespective of the logic (first-order, second-order) used in studying it, because there is a difference between what CA does in constructing $(V_0, \mathcal{C}_0)$ and what the mathematician does in studying the properties of $(V_0, \mathcal{C}_0)$.

At the end of the exploration of the room on the part of CA, the current conceptual space will be $(V_1, \mathcal{C}_1)$, where V_1 is exactly like V_0 apart from the fact that a finite portion of it now models the room representing, for instance, within the conceptual structure of V_1 the sets of points corresponding to D_1 and D_2 by including within $\mathcal{C}_1$ the corresponding regions.

The task set to CA can now be accomplished within $(V_1, \mathcal{C}_1)$. In fact, CA can, without knowing what a circle, a disc, etc. are, translate D_1 onto D_2 and *vice versa*. (Remember that a translation is a rigid motion within $(V_1, \mathcal{C}_1)$.)

However, there is a task that CA cannot accomplish within a 2-d conceptual space, and this is: placing D_1 on top of D_2. To represent the situation CA needs a 3-d conceptual space, i.e., a vector space $X = \mathbb{R}^3$ (over $\mathbb{R}$) together with the appropriate conceptual structure $\mathcal{C}$. Of course, here X is obtained by means of the direct product of $\mathbb{R}^2$ by $\mathbb{R}$.

An interesting application of projection is the following which relates to a 3-d task that can be accomplished by means of a projection onto a 2-d conceptual space: seeing whether a given sphere lying on the floor fits into a cubic box placed next to it. Once again, our agent does not know what a sphere or a cube are, but can find a way of representing and solving the problem in a 2-d conceptual space by considering whether or not a maximum circle of the sphere can fit into a face of the cubic box.

6 Conclusions

In this paper we have introduced global operations which allow cognitive agents to build and rearrange their conceptual representations as a consequence of their perceptions and according to their goals.The proposed operations provide the agent with the capabilities to focus on and represent, in a proper current conceptual space, specific aspects of the perceived environment.

In order to evaluate the correctness of our proposal, we intend to produce a simulation environment within which to test on an artificial agent the efficiency of the model put forward

Acknowledgements

This work has been partially supported by the PON01_01687 - SINTESYS (Security and INTElligence SYSstem) Research Project.

References

1. A. Chella, M. Frixione, and S. Gaglio. A cognitive architecture for artificial vision. Artif. Intell., 89:73111, 1997.
2. Alessandra De Paola, Salvatore Gaglio, Giuseppe Lo Re, Marco Ortolani: An ambient intelligence architecture for extracting knowledge from distributed sensors. Int. Conf. Interaction Sciences 2009: 104-109.
3. HyunRyong Jung, Arjun Menon, Ronald C. Arkin. A Conceptual Space Architecture for Widely Heterogeneous Robotic Systems. Frontiers in Artificial Intelligence and Applications, Volume 233, 2011. Biologically Inspired Cognitive Architectures 2011, pp. 158 - 167. Edited by Alexei V. Samsonovich, Kamilla R. ISBN 978-1-60750-958-5
4. Gärdenfors, P.: 2004, Conceptual Spaces: The Geometry of Thought, MIT Press, Cambridge, Massachusetts.

5. Gardenfors, Peter (2004). Conceptual spaces as a framework for knowledge representation. Mind and Matter 2 (2):9-27.
6. S. Harnad. The symbol grounding problem. Physica D, 1990.
7. Bernhard Scholkopf and Alexander J. Smola. 2001. Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. MIT Press, Cambridge, MA, USA.
8. Janet Aisbett and Greg Gibbon. 2001. A general formulation of conceptual spaces as a meso level representation. Artif. Intell. 133, 1-2 (December 2001), 189-232. DOI=10.1016/S0004-3702(01)00144-8 http://dx.doi.org/10.1016/S0004-3702(01)00144-8
9. Augello, A., Gaglio, S., Oliveri, G., Pilato, G. (2013). An Algebra for the Manipulation of Conceptual Spaces in Cognitive Agents. Biologically Inspired Cognitive Architectures, 6, 23-29.
10. Chella A., Frixione,M. Gaglio,S. (1998). An Architecture for Autonomous Agents Exploiting Conceptual Representations. Robotics and Autonomous Systems. Vol. 25, pp. 231?240 ISSN: 0921-8890.
11. Carsten Kessler, Martin Raubal - "Towards a Comprehensive Understanding of Context in Conceptual Spaces" - Workshop on Spatial Language in Context - Computational and Theoretical Approaches to Situation Specific Meaning. Workshop at Spatial Cognition, 19 September 2008
12. Parthemore, J. and A. Morse (2010). "Representations reclaimed: Accounting for the co-emergence of concepts and experience", Pragmatics and Cognition, 18 (2), pp. 273-312.