Proceedings of the 2007 IEEE Symposium on
Foundations of Computational Intelligence (FOCI 2007)

Knowledge Representation and Reasoning in Conceptual Spaces

John T. Rickard[*]
Lockheed Martin Corporation
4637 Shoshone Drive
Larkspur, CO 80118
303/681-9245 (W)
503/214-6283 (F)
terry.rickard@lmco.com

Janet Aisbett.
Faculty of Science & IT
University of Newcastle
Callaghan, NSW, Australia
61 24929 6811
Janet.Aisbett@newcastle.edu.au

Greg Gibbon
Faculty of Science & IT
University of Newcastle
Callaghan, NSW, Australia
61 24929 6811
Greg.Gibbon@newcastle.edu.au

Abstract:

This paper presents a conceptual system in which concepts are defined by binary associations between properties. Properties are measurable membership functions, defined on sets equipped with a measure that are the disjoint domains of representation. Instances of concepts (observations) are sets of points from these domains. Requiring properties to be measurable enables their overlap to be precisely described. Similarity between concepts, and between observations and concepts, is naturally defined using fuzzy subsethood, and similarity and overlap are used to set attention in categorization tasks. This formulation therefore follows Gardenförs in recognizing the importance of property associations and of similarity in conceptual systems.

1 Introduction

Ordered sets of alphanumeric values are often used to describe physical objects, conditions, events, or examples of abstract concepts. This representational form is familiar as relational databases. The ordered sets can be thought of as points in a multidimensional feature space. A similarity measure imposed on such a space allows observations of like phenomena to be compared, or compared with groups of similar observations that define real world *concepts*, or to be compared with concepts that are defined theoretically, such as "below freezing level".

The importance of similarity as an element in human cognition was recognized by Gardenförs (2000) who noted that similarity was not naturally modeled in the associationist (neural net) and symbolic logic representation schemes which had played such a large role in artificial intelligence. He introduced the term *conceptual spaces* to describe a representation scheme based on domains equipped with geometrical properties that enable similarity to be modeled and computed in a natural way. Gardenförs saw conceptual spaces as lying midway between the symbolic and associationist approaches, and, because their points (entities) have measurable qualities, they supply grounding for symbols (Aisbett and Gibbon, 2001, Roy and Reiter, 2005).

[*] Corresponding author

1-4244-0703-6/07/$20.00 ©2007 IEEE

Recent applications of conceptual spaces include scene interpretation and robotic learning (Macaluso et al, 2005), geographical similarity (Schwering and Rabaul, 2005; Ahlquist, 2005), support for conversational agents (Agostero et al, 2005) and data mining (Lee, 2005).

Aisbett and Gibbon (2001) re-formulated Gardenförs' theory of conceptual spaces in the mathematical framework of metric spaces, where similarities are defined via distance. To interface a conceptual space both to symbolic and subconceptual representations, Aisbett and Gibbon divided conceptual space into *symbol* and *concept* subspaces, and provided dynamics by defining interactions between these subspaces. The dynamics were guided by an attention-setting mechanism, and the trajectories of states within the space were interpreted as solutions to reasoning tasks such as identification and categorization. Representation of composite concepts (for example, "a dark haired woman in a black dress") was accommodated via multiple copies of this basic notion of a conceptual space.

Rickard (2006) represented concepts by attributed graphs whose vertices correspond to properties and whose weighted edge strength matrix captured the correlations between properties, with the salience weight of each property as a node attribute. The similarity of two concepts was a fuzzy set theoretic measure based on the pairwise joint similarities of the correlations. This formulation followed Gardenförs in defining concepts in terms of properties and their correlations. Complex concepts could be represented directly because the graphs could involve multiple properties on the same domain. Rickard also described how individual observations might be represented in an analogous matrix format, and used this to describe the similarity of an observation to a concept in terms of its membership to the closest property on each domain used in defining the concept.

This paper unifies Aisbett and Gibbon's and Rickard's extensions of conceptual spaces by defining domains as sets equipped with a measure, properties as measurable membership functions on a domain, and concepts as sets of associations between properties.

2 Reformulation of Conceptual Spaces

2.1 Domains and Properties

Gardenförs' starting point for a conceptual space was a set of *dimensions* capable of describing the attributes of the information to be represented. Dimensions were organized into incommensurate *domains*. A natural *property* was a convex region in some domain, where, to extend the notion of convexity beyond real valued domains, Gardenförs stipulated that each domain Δ_i be equipped with a trivariate logical relationship he called "betweenness".

In redefining the elements of conceptual space, we use membership functions that can be thought of as embodying the symbols which are the property labels.

<u>Definition 1.</u>

(a) A *domain* Δ_i is a set equipped with a measure m_i. A *property* is a member of an index set J_i (the property label) together with a measurable membership function $p_{ij} : \Delta_i \to [0,1]$, $j \in J_i$.

(b) The *intersection* of two properties p_{ij} and p_{ik} defined on the same domain Δ_i is the fuzzy set by the membership function $p_{ij} \cap p_{ik} : \mathbf{x} \to \min\left(p_{ij}(\mathbf{x}), p_{ik}(\mathbf{x})\right)$.

The assumption of measurability does not imply any topological structure on X itself, but allows us to talk about the *extent* of a property p_{ij}, defined as the Lebesgue integral $\int p_{ij}(\mathbf{x})dm_i$, and the *relative extent of overlap* B^i_{jk} of property k with respect to property j on the same domain,

$$B^i_{jk} = \frac{\int \min\left(p_{ij}(\mathbf{x}), p_{ik}(\mathbf{x})\right)dm_i}{\int p_{ij}(\mathbf{x})dm_i}. \tag{1}$$

This definition provides for a fuzzy notion of property membership, in which, for example, a particular person might have a high degree of membership for both "criminal" and "virtuous" properties, even though these properties may have little overlap. Combinations of properties to form new properties can be accomplished by functions on the basic property membership functions, such as conjunction or ordered weighted averaging (Yager and Kacprzyk, 1997). Hence, the membership function associated with a property is the fundamental descriptor of that property, rather than simply a region of space as in earlier works on conceptual spaces, including our own.

2.2 Concepts

We follow Gardenförs in recognizing that concepts are determined not only by their representative properties, but more importantly are determined through co-occurrences or associations between properties. Associations may be learned through experience, or may be theoretically determined, for example by defining a world in which all squares are red and 90% of circles are yellow. In Gardenförs' formulation, there was a tacit assumption that the properties defining a concept were disjoint, except possibly along shared boundaries. However, Rickard (2006) allowed that the set of properties that define a concept may contain multiple properties on any one domain, and we also allow this. In keeping with our fuzzy approach to property membership, we also allow that property self-associations need not be crisp. If a property is learned from experience, this reflects the fact that the average membership of the training set in the property may be less than unity.

Definition 2.

(a) A *concept* C is a named function $C : I(C) \times I(C) \to [0,1]$ where $I(C)$ is a finite subset of $\cup J_i$ containing the properties of C and where for all $a,b \in I(C)$, $C(a,a) > 0$ and $C(a, b) = 0$ whenever $C(b, a) = 0$.

The name of C is called the concept *label.* To reflect Rickard's notation, and standard matrix notation, we write $C(a,b)$ as C_{ab} and call it the *association* between the pair of properties $a,b \in I(C)$.

(b) A concept C' is *smaller* than C if $I(C') \subset I(C)$ and $C'_{ab} = C_{ab} \forall a,b \in I(C')$. We write $C' \subset C$.

The usual terminology of subconcepts and superconcepts has been avoided in (b) because a smaller property set may result in specialization, generalization or a combination. For example if C has mutually exclusive properties, say red or blue on the color domain, but C' only has the red property then C' is a specialization. If C has properties red and Arial on the color and font domains respectively, but C' only has the property red, then C' is a generalization.

It is well known that similarity judgments are dependent on context: thus, a cherry is more like an apple than a grape when the context is shape, but more like a grape than an apple when the context is size. Context can be defined by specifying domains, or properties within domains. In either case a context G can be considered to be a set of properties.

Allowance must be made for overlapping properties because, for example, one concept might be defined using properties "red" and "oval" and another as "dark red" and "circular". The second concept is really a subset of the first, yet if no allowance is made for the relationship between the properties, they might be taken to have zero similarity.

Definition 3

A *context* G is a set of properties.

The *similarity* $s(C^1, C^2)$ of two concepts C^1 and C^2 in the context G is given by

$$s\left(C^1,C^2\right)=\frac{\sum_{a,b}\min\left(C^1_{ab},C^2_{ab}\right)}{\sum_{a,b}\max\left(C^1_{ab},C^2_{ab}\right)} \tag{2}$$

where the sum is over all pairs of elements in $I\left(C^1,C^2;G\right)\equiv\left(I\left(C^1\right)\cup I\left(C^2\right)\right)\cap G$.

If a is a property in domain Δ_r involved in concept C^1 but not in concept C^2, and b is a property in domain Δ_s involved in concept C^2 but not in concept C^1, then the associations C^1_{ab} and C^2_{ab} are not defined. In this case, we specify C^1_{ab} to be

$$C^1_{ab}=B^r_{aa*}B^s_{bb*}C^1_{a*b*} \tag{3}$$

where $a*$ and $b*$ are chosen so that if where B^r_{xy} is defined as in (1) then $B^r_{aa*}\ge B^r_{aa'}$ for all $a'\in I_r\cap I\left(C^1\right)$, $B^s_{bb*}\ge B^s_{bb'}$ for all $b'\in I_s\cap I\left(C^1\right)$, and likewise for C^2_{ab}.

In other words, when the association between two properties a and b in $I(C^1, C^2)$ is not defined for the concepts being compared, we specify it as the associations between properties which maximally overlap a and b in their respective domains.

2.3 Observations

Observations are conventionally described as points in a feature space, that is, as ordered sets of attribute values. The following definition allows an observation to have an arbitrary number of values in any one domain, and hence to have multiple disjoint properties.

Definition 4.

(a) An *observation* $\mathbf{o}$ is a collection of sets of points $\mathbf{o}_i$ from domains $\{\Delta_i\}_{i\in K}$, where the set $\mathbf{o}_i$ is called the values of the observation in the i^{th} domain.

(b) The membership $p_{ij}(\mathbf{o})$ of an observation $\mathbf{o}$ in a property p_{ij} is defined as $\max\{p_{ij}(y):\ y\in\mathbf{o}_i\}$.

As noted, properties may be defined using functions of basic property membership functions, such as conjunctions. The membership of an observation $\mathbf{o}$ in a conjunction of properties such as "red and blue" is defined as $\min\{p_{ij}(\mathbf{o}), p_{ik}(\mathbf{o})\}$. Exclusive properties such as "blue only" are defined as $\min\left\{p_{ij}(\mathbf{o}), \min_k\left[1-p_{ik}(\mathbf{o})\right]\right\}$ where, for example, j indexes the color "blue" and the indices k run over a set of color properties that are disjoint from "blue". A property membership function corresponding to "blue only" would map to zero any observation that possesses any of the colors other than "blue" to a full degree. Likewise a property membership function requiring "blue and yellow" would map an observation possessing only yellow to zero.

Next, we need a way of identifying observations as being instances of concepts, or at least being "like" a concept. A distinguishing characteristic of Rickard's (2006) definition was that it attempted to allow an observation to have full similarity to a concept as long as it fitted a subtype of that concept, no matter how unusual that subtype was as an exemplar of the concept. Rickards's definition was problematic when concepts had subtypes defined using different domains. For example, if a concept includes all red objects, and all square objects, and no red objects are square, then an observation of a red object (that was therefore not square) would only have similarity 1/2 to the concept. The definition below explicitly decomposes a concept into subtypes before comparing it with an observation.

Definition 5

The *subtypes* of a concept C are the maximal concepts smaller than C that have no zero associations. That is, if C' is a subtype, then $C'_{ab}>0\ \forall a,b\in I(C')$ and if C'' is any concept that is smaller than C then either $C'\not\subset C''$ or $C''_{ab}=0$ for some $a,b\in I(C'')$.

The *similarity of an observation* $\mathbf{o}$ *to a concept* C in the context G is:

$$s(C,\mathbf{o})=\max_{C'\text{ a subtype of }C}\left(\frac{\sum_{a,b\in I(C')\cap G}\min\left(C'_{ab},o_{ab}\right)}{\sum_{a,b\in I(C')\cap G}C'_{ab}}\right) \tag{4}$$

where for $a\in J_r\cap I(C)$ and $b\in J_s\cap I(C)$

$$o_{ab}=\min\left(\max_{j\in I_r\cap I(C)}p_{rj}(\mathbf{o}),\ \max_{j\in I_s\cap I(C)}p_{sj}(\mathbf{o})\right). \tag{5}$$

When a domain Δ_n is used in defining a property of C but the observation does not include a value for this domain, it is given membership value 0 for all properties on Δ_n.

In order to compare the observation with a subtype, it is treated as a concept involving each of the properties of that subtype. The association function for this new concept allocates the same association to each pair of properties from a given pair of domains, and this association is based on the highest membership value that $\mathbf{o}$ has for a subtype property on each of the domains.

The set of subtypes is defined iteratively by identifying an association $\widehat{C}_{ab} = 0$ in the concept $\widehat{C}$ then replacing it by the smaller concepts defined on $I(\widehat{C}) - \{a\}$ and $I(\widehat{C}) - \{b\}$. The process ends when all associations are non-zero.

The subtypes allow "criss-cross" categories to be separated, in a way analogous to nearest neighbor formulations of similarity (which depend on a distance measure on the domains). Criss-cross categories, however, confuse classifiers based on prototypes or which take mean values in a domain. A "criss-cross" situation occurs, for example, when one type of class or category is composed of instances having property *a* on domain Δ_r and property *c* on domain Δ_s, and also of instances having property *b* on Δ_r and property *d* on Δ_s; and a second class or category is composed of instances having property *b* on Δ_r and property *c* on Δ_s, and also of instances having property *a* on Δ_r and property *d* on Δ_s. The two properties on each domain are assumed to be mutually exclusive.

3 Categorization and attention setting

The tasks of interest to us are identification or categorization of an observation. Suppose therefore the task is to find the concept amongst the set of alternatives $\{C^1, C^2, \ldots C^g\}$ which the observation best represents, in the sense of equation (4). An observation $\mathbf{o}$ can simply be compared with each of the concepts in $\{C^1, C^2, \ldots C^g\}$ using the similarity measure proposed in (4), which compares observations with subtypes of each concept in the current context. Should the highest resulting similarity exceed an *a priori* threshold, then the observation is said to be an instance of the corresponding concept or category.

The subsethood calculation in (4) could be weighted according to the importance of the properties in discriminating the concepts, as computed in the attention setting phase, or through *a priori* knowledge.

However, it is well known that classification performance is improved when only attributes which carry discriminatory information are used in assessing the similarity of an observation to a target class. This is because noise on observations in non-discriminatory domains can distort the similarity rating (Holte, 1993). So to improve performance, the context should be domains or properties which are likely to be helpful to the task at hand.

Attention is used to prioritize the properties of a concept C^i according to how different they are from properties in the alternative concepts, *or* how different their associations with other properties of C^i are from corresponding associations in the alternative concepts. Essentially, we are only interested in properties which have little overlap with properties of other concepts, or else are part of property pairs with very different associations yet high overlap. To identify properties with little overlap, we look for properties *j* for which the best match overlap $B^i_{jj(n)}$ is small. To identify different associations between the property pairs we look for very small ratios

$\dfrac{\min\left(C_{js}^{m}, C_{j(n)s(n)}^{n}\right)}{\max\left(C_{js}^{m}, C_{j(n)s(n)}^{n}\right)}$, which we want to occur in conjunction with high overlap on the property s or, equivalently, with small value of $\left(B_{ss(n)}^{r}\right)^{-1}$ for s defined on Δ_r.

We are now in a position to construct a fuzzy ambiguity measure $\zeta(j,m)$ of property j with respect to concept C^m for the task of discriminating C^m from other C^n, $n \in \{1 \ldots g\}$, $n \neq m$. A small value of this measure will indicate that property j has high discriminating power (i.e., low ambiguity) due to its low overlap with properties involved in other concepts in the same domain and/or its distinctly different association with properties in other domains with respect to other concepts. An appropriate expression for this measure is

$$\zeta(j,m) = \min_{n}\left(\min\left(B_{j\,j(n)}^{i}, \min_{s \in I(C^m)} \left(\frac{\min\left(C_{js}^{m}, C_{j(n)s(n)}^{n}\right)}{B_{s\,s(n)}^{r} \max\left(C_{js}^{m}, C_{j(n)s(n)}^{n}\right)} \right) \right) \right), \tag{6}$$

where the outer minimum is over $n \in \{1, 2, \ldots g\} \setminus \{j\}$.

The innermost parenthetical expression in (6) measures the ambiguity between the associations of property j to property s in concept C^m and of property $j(n)$ to property $s(n)$ in concept C^n, inversely weighted by the overlap between properties s and $s(n)$. Thus a small value of the association ratio combined with a large value of the overlap $B_{s\,s(n)}^{r}$ indicates that property j has low ambiguity with respect to property s, notwithstanding a high overlap between s and $s(n)$.

A property may be shared by multiple concepts. So finally, we rank a property by its minimum ambiguity over all concepts in the set, that is, by $\min_{m} \zeta(j,m)$, which yields the overall efficacy of property j relative to the set $\{C^1, C^2, \ldots C^g\}$. Ideally we would only need to consider properties for which $\min_{m} \zeta(j,m) = 0$. These include properties of a concept which have no overlap with any property on a domain that is used in defining at least one other concept (in which case an observation having this property cannot represent the second concept, but may represent the first). The properties for which $\min_{m} \zeta(j,m) = 0$ also include any property involved in a zero association for which there is a non-zero association between the best matching properties in some other concept. In this case, an observation exhibiting both properties cannot represent the first concept but may represent the second.

Having ranked properties, the cutoff for attention is dependent in part on the time available for processing. In the absence of time constraints, such a cutoff would only be used to remove properties on domains from the task context which are not involved with any other concept in $\{C^1, C^2, \ldots C^g\}$, or which have similar associations in all the concepts and so are not useful in discriminating between them.

4 Conclusion

This paper combined two independently developed extensions of Gardenförs' original formulation of conceptual spaces into a single, more capable knowledge representation and inferencing framework. The new theory was less restrictive than the metric space formulation of Aisbett and Gibbon (2001) yet enhanced Rickard's (2006) formulation by explicitly describing properties. The domains underlying the representation were simply sets equipped with a measure, which allowed for notions of extent and overlap of properties which were defined as measurable membership functions on the domains. Additional elements, including subtypes, were introduced to handle reasoning difficulties identified in our earlier formulations. We also introduced the notion of the ambiguity of properties in a given categorization task, which enabled attention to be focused upon properties having the greatest discrimination power.

Research underway is evaluating the comparative performance of the new theory in categorization. Future research will examine learning aspects, since this is a critical requirement of all useful reasoning systems.

References

Aisbett J.and Gibbon, G. (2001) A general formulation of conceptual spaces as a meso level representation, *Artificial Intelligence*, 133, 189-232.

Agostaro R., Augello A., Pilato G., Vassallo G., Gaglio S. (2005) A conversational agent based on a conceptual interpretation of a data driven semantic space, in *Advances In Artificial Intelligence, Lecture Notes In Artificial Intelligence* 3673, 381-392.

Ahlqvist O. (2005) Using uncertain conceptual spaces to translate between land cover categories *International Journal Of Geographical Information Science* 19, 7, 831-857.

Gardenförs, P. (2000) *Conceptual Spaces: the Geometry of Thought,* Cambridge, MA: MIT Press.

Holte, R. (1993) Very simple classification rules perform well on most commonly used datasets Machine Learning, 3, 11, 63-91.

Lee I. (2005) Data mining coupled conceptual spaces for intelligent agents in data-rich environments, in *Knowledge-Based Intelligent Information And Engineering Systems, Pt 4, , Lecture Notes In Artificial Intelligence* 3684, 42-48.

Macaluso I., Ardizzone E., Chella A., Cossentino M., Gentile A., Gradino R., Infantino I., Liotta M., Rizzo R., Scardino G. (2005) Experiences with CiceRobot, a museum guide cognitive robot, in *Advances In Artificial Intelligence, Lecture Notes In Artificial Intelligence* 3673, 474-482.

Rickard J. (2006) A concept geometry for conceptual spaces, *Journal of Fuzzy Optimization and Decision Making,* vol. 5, 311-329.

Rickard J. and Yager R. (2006) Hypercube graph representations and fuzzy measures of graph properties, to appear in *IEEE Trans. Fuzzy Systems*.

Roy D. and Reiter E. (2005) Connecting language to the world, *Artificial Intelligence* 167, 1-2, 1-12

Schwering A. and Raubal M. (2005) Spatial relations for semantic similarity measurement, in *Perspectives In Conceptual Modeling, Lecture Notes In Computer Science* 3770, 259-269.

Yager, R. and J. Kacprzyk (1997) *The Ordered Weighted Averaging Operators: Theory and Applications*, Boston: Kluwer Academic Publishers.