The Semantic Pointer Model of Similarity

1. Introduction

As early on as infancy, human beings engage in judgments of similarity. As infants these judgments are very simple, such as judging the similarity between different colours (Franklin & Davies, 2004). By the time we are adults, we engage in many highly abstract and complex judgments of similarity, some of which we may not even be aware of. Despite the strong presence of similarity judgments in our reasoning, an accurate model of similarity has yet to be agreed upon. In the last century, two dominant models of similarity[1] have been proposed by psychologists: the *geometric model* (Shepard, 1962) and the *contrast model* (Tversly, 1977). Both models have been criticized as deficient in many respects (Tversky, 1977; Markman & Gentner, 1993). I propose that neither traditional model on its own is sufficient and that in order to accurately capture the nature of similarity, aspects of both models are required. I further argue that with the use of semantic pointers, introduced by Chris Eliasmith (Forthcoming 2013), a superior account of similarity, what I call here the Semantic Pointer Model of Similarity (SPMS), can be given. This new model retains the attractive features of the traditional geometric model while matching the contrast model in its explanatory power. I begin by briefly outlining the geometric model and the contrast model individually as they have been traditionally presented, along with some of the major criticisms of each model. In the second section, I present an alternative metric model of similarity based on semantic pointers, the SPMS. In presenting the SPMS, I will address some of the major criticisms of the traditional geometric model of similarity (Tversky, 1997, Laakso & Cottrell, 2005) and show why the SPMS does not suffer from these problems.

2. Traditional Models of Similarity

On the geometric model, also known as the mental distance model, objects are represented as points in a mental space and the similarity between objects is taken as a measure of the distance between these points. As the distance between points in mental space decreases, the level of similarity between objects increases (Shepard, 1962; Thurstone, 1927; Torgerson, 1965). This mental distance is measured by a metric, which must satisfy the mathematical definition of a distance metric. This definition states that a metric distance function, d, assigns to every two points in a space a number, called their distance, in accord with the following axioms:

The triangle inequality: $d(x, z) \leq d(x, y) + d(y, z)$

Symmetry: $d(x, y) = d(y, x)$

Non-negativity: $d(x, y) \geq 0$

Minimality: $d(x, y) = 0$ if and only if $x = y$

Amos Tversky has famously criticized the geometric model on the grounds that its dimensional assumption has limited applicability (1977). Though Tversky believes that a geometric model may suffice for certain basic judgments of similarity, he argues that it cannot account for all judgments of similarity. More specifically, Tversky argues that the model suffers from problems with symmetry, minimality, and the triangle equality (1997, see also Laakso & Cottrell, 2005). These problems will be considered in detail further on in the paper.

On the contrast model, also called the feature-mapping model or the set theoretic model, objects are considered to have a collection of features and the similarity between two objects is determined by measuring the intersection of the objects' sets of features. On this model, famously advocated by Tversky (1977), the more features two objects have in common, the more similar they are to one another and the less features two objects have in common (or the more differences they have), the less similar they are to one another. Though Tversky's contrast model provides a way to avoid the problems attributed to the geometric model, it does not come without problems of its own (Markman & Gentner, 1993; Gentner & Markman, 1995). According to Dedre Gentner and Arthur B. Markman, the contrast model requires an assessor to examine the commonalities and differences in order to make a judgment of similarity, yet it defines no way to determine what counts as a commonality and what counts as a difference. Furthermore, the contrast model does not allow for certain features to be relevant in some judgments of similarity and not relevant in another. Allowing for certain features be relevant in some judgments of similarity and not relevant in others is crucial.

3. Breaking Traditions: The SPMS

Chris Eliasmith's theory of biological cognition introduces semantic pointers as a class of mental representation. In brief, the semantic pointer hypothesis states "higher-level cognitive functions in biological systems are made possible by semantic pointers. Semantic pointers are neural representations that carry partial semantic content and are composable into the representational structures necessary to support complex cognition" (Eliasmith, n.d.). For our purpose here, we can think of semantic pointers simply as compressed representations.[2] Take for instance, an image of a dog. Upon being presented with some form of perceptual input (in this case visual), some pattern of activity among a population of neurons in the brain is generated. This resulting pattern of activity is encoded and is referred to as a representation. The brain then continues to transform the representations, creating increasingly abstract statistical summaries of the original perceptual input. These abstract statistical summaries retain semantic information about the states they represent and are called semantic pointers. However, for those of you who have at some point encountered a dog, you know that when you think of a dog, you do not merely see a visual image. In fact, when you imagine a dog, you likely think of what a dog smells like, what a dog sounds like, what a dog feels like to touch, and a number of other things. Just as visual perceptual input can give rise to visual semantic pointers, other types of perceptual input can give rise to other types of semantic pointers. What is important to note here is the fact that semantic pointers can be further compressed along with other semantic pointers to form, what I call, *richer* semantic pointers. These richer semantic pointers can also be decompressed in order to retrieve the individual *constituent* semantic pointers. By making use of semantic pointers, I will now introduce the SPMS.

The Semantic Pointer Model of Similarity is a metric model of similarity in that it takes similarity to be the distance between semantic pointers in some n-dimensional mental comparison space. Like with the traditional geometric model of similarity, the shorter the distance between two semantic pointers, the more similar the two things represented by the semantic pointers are and the farther the distance, the less similar they are. The key difference here between the semantic pointer model of similarity and the traditional geometric model of similarity is the comparison spaces. In the traditional geometric model of similarity, all judgments of similarity are made in the same mental comparison space. However, in the SPMS the comparison space changes depending on the semantic pointers being used in the judgment and with the context of the judgment. What this means is that the same two objects can have different measures of similarity between them in different spaces (or in different contexts). Think back to our sugar and salt example. For each judgment, similarity based on appearance and similarity based on taste, the comparison space changes. This allows for two objects to have different degrees of similarity between them in different contexts. By limiting similarity to one comparison space, we assume that all judgments of similarity are equally complex and this is simply not the case. In order to better explain the details of the SPMS, let us consider some of the original objections to the geometric model raised by critics.

4. Problems for Similarity

Like Tversky, Aarre Laakso & Garrison Cottrell have raised objections to the geometric model by targeting problems with symmetry, minimality, and the triangle inequality (Tversky, 1977; Laakso & Cottrell, 2005).[3] I will consider each of these three problems here and show why they do not exclude human similarity judgments from satisfying the mathematical definition of a metric.

4.1 The Triangle Inequality

Tversky (1977) and Laakso & Cottrell (2005) argue that human semantic similarity violates the triangle inequality condition of the mathematical definition of a metric. Consider Laakso & Cottrell's example of a banana, a boomerang, and an apple.[4] On a traditional geometric model, when the similarity distance is measured between all three objects, the distance between the apple and the banana (which is short due to their both being fruits) and the distance between the banana and the boomerang (which is short due to their shapes) added together is shorter than the distance between the apple and the boomerang. For this reason, critics argue that the triangle inequality condition is not satisfied. However, this argument is based on the problematic assumption that similarity judgments are all made within the same comparison space.

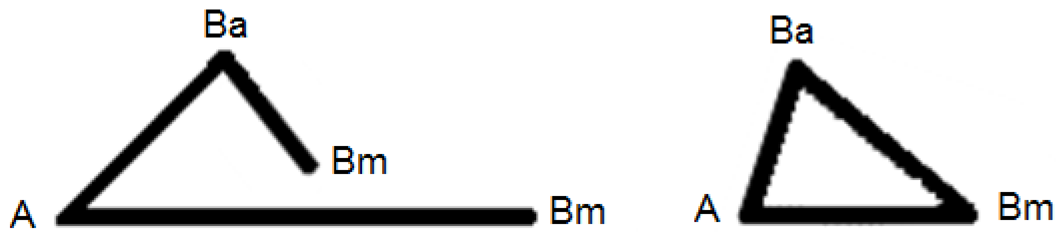

Figure 1: (A=Apple, Ba=Banana, Bm=Boomerang) The image on the left represents the example of similarity judgments between an apple, a banana, and a boomerang made in comparison spaces. The image on the right represents the same similarity judgment when made in the same comparison space.

On the SPMS, when one measures the similarity between the apple and the banana, he does so in a different context (that is, in different comparison space) than when he measures the similarity between the banana and the boomerang. Specifically, when he measures the distance between the apple and the banana to be short, he does so with respect to their both being fruits. When one measures the distance between the banana and the boomerang to be short, he does so with respect to their curved shape. Though it may be tempting to compare these measurements of similarity directly, one has to keep in mind that a distance metric is only defined within a particular comparison space. For this reason, measurements of similarity made within different comparison spaces cannot be directly compared. In this case and all other cases where the triangle inequality appears to have been violated, the comparison space changes as you judge the similarity between the three objects, and so the three measurements of similarity cannot be directly compared and do not violate the triangle inequality.

Consider the same three objects, only this time all the measurements of similarity will be taken with respect to fruitness (on the same subspace). The distance between the apple and the banana will again be short, as they are both fruits. Now, the distance between the banana and the boomerang will not be short as it was before, as a banana is a fruit and a boomerang is not a fruit.

Likewise, the distance between the apple and the boomerang will not be short, as an apple is a fruit and again, a boomerang is not a fruit. In this case, if the distance between the apple and the banana is added to the distance between the banana and the boomerang, it is not immediately obvious and it seems highly unlikely that it would be shorter than the distance between the apple and the boomerang (see Figure 1). By taking all the similarity measurements between the three objects in the same comparison space, the triangle inequality is not violated.[5]

4.2 Minimality

The second problem presented by Tversky (1977) and Laakso & Cottrell (2005), is the problem of minimality. Laakso & Cottrell cite a study (Podgorny & Garner, 1979) in which subjects were found to judge the letter S to be more similar to itself than the letter W is to itself. In this study, subjects were shown images of letters and asked to make similarity judgments about the letters they had just seen. Reaction time was used as the measure of similarity, where longer reaction times indicated a lower degree of similarity and shorter reaction times indicated a higher degree of similarity. The two conclusions of the study that Laakso & Cottrell focus on concern intraobject similarity. The study concluded that (1) all objects are not equally similar to themselves and (2) when subjects judge letters as similar-dissimilar rather than as same-different, reaction time to say "similar" to identical objects supports the proximity constraint which requires that an object be most similar to itself.[6] From the conclusions of the Podgorny & Garner study, Laakso & Cottrell claim that the symmetry condition of the mathematical definition of a metric is violated. However there appears to be more going on here than Laakso & Cottrell recognize.

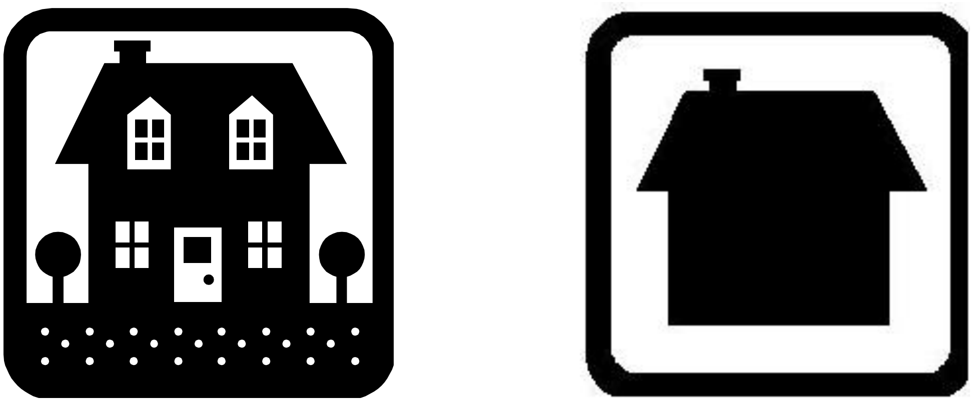

Figure 2: Both images depict a house, though the image on the left is considerably more complex than the image on the right.

It may be that reaction time as a measure of similarity can be useful in certain situations, if a certain level of control in the circumstances of the judgment is attained. However, without this control it appears that there may be outside factors interfering with the possibility of using reaction time to gage similarity. Consider the images in Figure 2. Even though the two images both depict a house, one is considerably more complex than the other. Imagine we ask someone to make a similarity judgment about whether or not the image on the left is similar to itself, and then to make another similarity judgment about whether or not the image on the right is similar to itself. The reaction time required to judge the image on the left as similar to itself will be longer simply because there is more to look at. However, this does not mean that the image on the left is any less similar to itself than the image on the right is to itself. Furthermore, there is a significant amount of time that it takes for a semantic pointer to go through a transformation in the brain. So, depending on the size and nature of the input, it can take longer for certain semantic pointers to be compressed or decompressed than it takes for others. Again, perhaps reaction time is not the best measure of similarity in all cases.

4.3 Symmetry

Studies of similarity have shown that human beings reliably judge less prominent things to be more similar to more prominent things than the reverse (Holyoak & Gordon, 1983). Similarly, people judge their friends to be more similar to themselves than they are to their friends. Based on this evidence, Laakso & Cottrell argue that the symmetry condition of the mathematical definition of a metric is violated. However, this criticism, like the one made with regard to the triangle inequality, is based on the idea that similarity judgments are all made in the same comparison space. By adopting the SPMS and the idea that comparison spaces can change with context, these problems disappear.

Consider Laakso & Cottrell's example of China and North Korea, China being the more prominent thing and North Korea being the less prominent thing. When a human makes the similarity judgment that North Korea is more similar to China than China is to North Korea, they appear to do so precisely because for them, China is a more prominent thing and North Korea is a less prominent thing. That is, the assessor knows[7] more things about China than they know about North Korea, meaning the assessor's representation (semantic pointer) of China is richer (has more constituent semantic pointers) than their representation of North Korea has. Given that the assessor knows more about China than North Korea, it is clear that the semantic pointer for China will be richer (that is, be made up of more constituent semantic pointers) than the semantic pointer for North Korea. On the SPMS, making this kind of similarity judgment is not as simple as measuring one distance between two distinct points. Rather, when making a similarity judgment of this type, one is examining the constituent semantic pointers of two richer semantic pointers and identifying the constituent semantic pointers of one that match the constituent semantic pointers of the other.

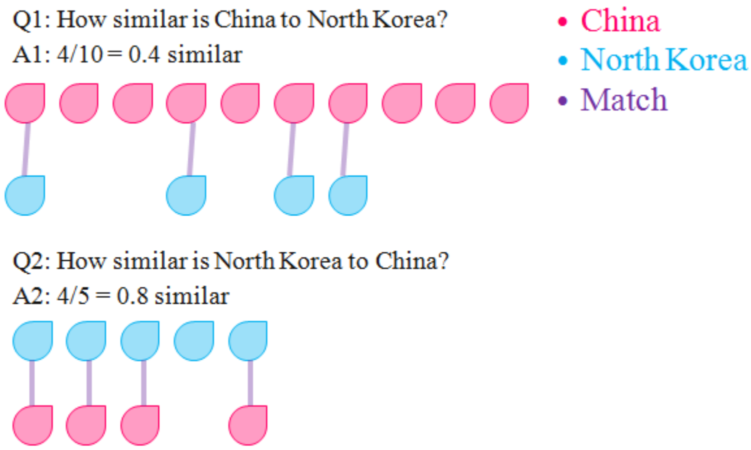

Figure 3: In the first question, the 10 things known about China are compared to what is known about North Korea and 4/10 matches are found. In the second question, the 5 things known about North Korea are compared to what is known about China and 4/5 matches are found. Thus it appears that North Korea is more similar to China than China is to North Korea.

To make this clear, imagine that a participant knows 10 things about China and 5 things about North Korea, as shown in Figure 3. When asked to judge the similarity between two semantic pointers, the context of the judgment can be changed depending on how the question is asked. In the first question, the context is set by China because it is introduced first. We can imagine what the participant is doing as lining up the 10 things she knows about China and checking them against what she knows about North Korea. In the first case, the participant confirms that 4 out of the 10 things she knows about China match what she knows about North Korea. For example, she may know that both China and North Korea are Asian countries (match), but she may also know that China's national currency is the Yuan Renminbi but not know what North Korea's national currency is (no match). In the second case, North Korea sets the context, and so we can imagine what the participant is doing as lining up the 5 things she knows about North Korea and checking them against what she knows about China. In this case, the participant confirms that 4 out of the 5 things she knows about North Korea match what she knows about China. This leads to the counter-intuitive judgment that North Korea is more similar to China than China is to North Korea.

The reason it seems as though a less prominent thing is more similar to a more prominent thing than the other way around is because there are matches between the constituent semantic pointers of each object's representation; there are constituent semantic pointers of each object's representation that are similar. Without these matches, these counter-intuitive similarity judgments are not made. For example, if I tell you that Flipsharps: have 4 legs, like cotton candy, speak Latin, are scared of the dark, and can breathe underwater, and I also told you that Snoolits: are allergic to gold and fly only during the autumn, you do not have the intuition that Snoolits are more similar to Flipsharps than Flipsharps are to Snoolits. Furthermore, the more one knows about the less prominent thing, the less likely these counter-intuitive judgments become. This accounts for why it is some people find these judgments to be so counter-intuitive in the first place. For example, a Canadian citizen who knows a roughly equal amount about Canada and the United States, may find it wildly counter-intuitive to claim that Canada may be more like the United States than the reverse (and vice versa). However, if a person from Egypt were to know a fair amount about the United States and very little about Canada, chances are a large portion of what they do know about Canada is similar to what they know about the United States, and so for them, a judgment such as Canada is more similar to the United States than the United States is to Canada, may not be so wild.

5. Conclusion

Finding an accurate model for similarity is by no means an easy task. The geometric model and the contrast model account for much of similarity's complex nature. However, on their own they are insufficient. The SPMS not only resolves the objections to the traditional geometric model, it also avoids the objections to the contrast model all together. By adopting the SPMS we allow judgments of similarity to be made in any number of comparison spaces and we account for counter-intuitive results of studies on human similarity judgments. By conceiving of similarity as the distance between semantic pointers in some mental comparison space, we are able to account for the vast range of complexities in similarity judgments.

Notes

[1] [1] The models of similarity that this paper is concerned with are being considered as models of similarity for human judgments of similarity. Whether or not there is some kind of similarity that exists independently of the mind or independently of human judgments of similarity are not concerns of this paper.

[2] For the purpose of this paper, many of the technical details surrounding neuronal activity and semantic pointers are left aside here. For a more detailed account of semantic pointers, see Eliasmith (Forthcoming).

[3] Tversky and Laakso & Cottrell both point to problems with the triangle inequality, minimality, and symmetry for very similar reasons. I will focus on the objections raised by Laakso & Cottrell in this paper.

[4] In Tversky's example we are asked to consider three countries: Jamaica, Cuba, and Russia. Though I have chosen here to discuss Laakso & Cottrell's example of the banana, boomerang, and apple, both examples make the same argument and can be responded to in the same way.

[5] It may be the case that when comparing some objects there is an overall sense of similarity, but it may also be the case that the subject doing the comparison is not able to make these judgments of overall similarity. In any case, answering that question is beyond the scope of this paper and does not affect the claims currently being made.

[6] This conclusion is rather interesting with regard to the argument put forth by Laakso & Cottrell. Laakso & Cottrell claim based on results from a second study, that subjects judged the letter M to be more similar to the letter H than to itself. However, the results of the first study cited by Laakso & Cottrell (Podgorny & Garner, 1979) do not support and in fact contradict this claim. Based on their study, Podgorny & Garner conclude that there is a proximity constraint on similarity judgments which requires that an object be most similar to itself.

[7] The epistemic issues of whether or not these features of the representation are known or believed, or are justified or not justified are not important here. For ease of discussion, a particular constituent semantic pointer of an object's representation can be considered to be a particular 'thing known' about the object.

References

Eliasmith, C. (Forthcoming 2013). *How to build a brain: A neural architecture for biological cognition*. Oxford: Oxford University Press.

Eliasmith, C. (n.d.) *How to Build a Brain*. Retrieved April/24, 2013, from http://www.nengo.ca/build-a-brain

Franklin, A., & Davies, I. R. L. (2004). New evidence for infant colour categories. *British Journal of Developmental Psychology*, 22(3), 349-377.

Gentner, D., & Markman, A. B. (1995). Similarity is like analogy: Structural alignment in comparison. In C. Cacciari, & D. Gentner (Eds.), *Similarity: In language, thought and perception* (pp. 111). Turnhout: Brepols.

Gilmore, G. C., Hersh, H., Caramazza, A., & Griffin, J. (1979). Multidimensional letter similarity derived from recognition errors. *Perception and Psychophysics*, 25, 425-431.

Holyoak, K. J., & Gordon, P. C. (1983). Social reference points. *Journal of Personality and Social Psychology*, 44, 881-887.

Laakso, A., & Cottrell, G. W. (2005). Churchland on connectionism. In B. L. Keeley (Ed.), *Paul Churchland* (pp. 113). New York: Cambridge University Press.

Markman, A. B., & Gentner, D. (1993). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, 32(4), 517-535.

Podgorny, P., & Garner, W. R. (1979). Reaction time as a measure of inter-intraobject visual similarity: Letters of the alphabet. *Perception and Psychophysics*, 26(1), 37-52.

Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function, I. *Psychometrika*, 27(2), 125-140.

Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, 34(4), 273-286.

Togerson, W. S. (1965). Multidimensional scaling of similarity. *Psychometrika*, 30(4), 379-393.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327-352.