

Using Support Vector Machines to Determine Linear Separability

Bram Van Rensbergen (bram.vanrensbergen@ppw.kuleuven.be)

Laboratory of Experimental Psychology, University of Leuven
Tiensestraat 102 – box 3711, B-3000 Leuven, Belgium

Eef Ameel (eef.ameel@ppw.kuleuven.be)

Laboratory of Experimental Psychology, University of Leuven
Tiensestraat 102 – box 3711, B-3000 Leuven, Belgium

Fons Van Assche (fonzva@gmail.com)

Laboratory of Experimental Psychology, University of Leuven
Tiensestraat 102 – box 3711, B-3000 Leuven, Belgium

Gert Storms (gert.storms@ppw.kuleuven.be)

Laboratory of Experimental Psychology, University of Leuven
Tiensestraat 102 – box 3711, B-3000 Leuven, Belgium

Keywords: Linear separability; support vector machines; machine learning.

Introduction

Most theories on categorization agree on a main principle: category learning involves creating a category representation, and categorizing an item involves comparing that item to the representations of different categories. The theories, however, disagree on the nature of these category representations. There are two main competing lines of thought on category representations: exemplar-based theories and prototype-based theories (Valentine, 1991).

Prototype-based theories argue that objects are stored based on how similar they are to a central prototype (Rosch, 1973). In contrast, exemplar-based theories reason that objects are encoded in their absolute structure, defined by their own properties only and unrelated to any abstract summary representation (Medin & Schaffer, 1978). Years of research on the nature of categorization has resulted in mixed results, with evidence for both approaches.

One example of an issue on which exemplar- and prototype based theories make different predictions is linear separability. Two categories are considered linearly separable if a linear function of attributes exists that perfectly separates their exemplars (Ruts, Storms, & Hampton, 2004).

According to prototype-based models, for any pair of linearly separable categories represented in a geometrical space, that space is divided into two half spaces by a linear function that defines the points which are equidistant towards both prototypes. An item is then categorized in the category with the closest prototype (in that geometrical space). Thus, category membership can be determined simply by looking at the distance to the prototype, making linearly separable categories relatively easy to learn. On the other hand, categories that are not linearly separable would

take considerably longer to master, as this simple strategy of deciding on the closest prototype would not be sufficient to determine category membership.

According to exemplar-based models, on the other hand, proximity to the center of the category plays no role of any kind (Ruts, Storms, et al., 2004). Thus, exemplar-based models predict that, all other factors kept constant, linearly separable categories are not easier to master than other categories.

As prototype-based and exemplar-based models make different predictions regarding linearly separable categories, we can use those categories to shed light on the mechanisms that underlie categorization. In order to do so, however, we have to determine which categories are linearly separable, and which are not. There has been surprisingly little research into this issue. Studies that do investigate this tend to assess linear separability by first obtaining a geometric representation of the exemplars using multidimensional scaling, and then analyzing this representation with visual inspection (Malt, Sloman, Gennari, Shi, and Wang, 1999) or log linear analysis (Ruts, Storms, et al., 2004).

The current research intends to expand on these previous studies in investigating which semantic categories are linearly separable and which are not. Compared to these existing studies, however, we will use a fundamentally different technique to assess linear separability: linear support vector machines (LSVMs).

A support vector machine (SVM) is a mathematical concept used for supervised pattern learning (Vapnik, 1982; Cortes & Vapnik, 1995). Presented with a set of input data and their corresponding classes, an SVM learns which data correspond to which class. Once trained, the machine can be used for classification; for any given input, it predicts the corresponding class.

SVMs transform the input vectors into a (usually) high dimensional feature space with the help of a kernel

function, and look for the hyperplane that separates the classes optimally in that feature space (Cortes & Vapnik, 1995). SVMs put no restrictions on the nature of the kernel, allowing both linear and nonlinear functions. As the goal of our study is to assess the extent of linear separability of categories, we need a classifier that produces linear boundaries; as such, we can only apply linear kernel functions.

Compared to assessing linear separability by analyzing a geometric representation obtained with multidimensional scaling, LSVMs hold several advantages. For one, there is no issue of choosing the optimal dimensionality, as LSVMs always use the maximal number of dimensions present in the data. Secondly, LSVMs make no assumptions about the nature of the distribution of the items. This is in contrast to many statistical criteria used to analyze the geometric representation obtained with multidimensional scaling, which do put restrictions on the distribution.

Method and Results

In a first study, we examined the linear separability of natural and artifact concepts. The idea was to teach a LSVM to use feature values to predict category membership, and then to examine which categories the LSVM could linearly separate from one another. We looked at six pairs of natural categories and five pairs of artifact categories. Each pair consisted of 61 to 85 exemplars, which were rated on 30 to 51 features. We found that LSVMs are very efficient at using feature values to predict to which class an item belongs. Prediction accuracy was high both for natural classes (up to 100% accuracy) and for artifact classes (up to 97.07% accuracy). Additionally, we found that some of the natural categories were linearly separable and some were not, and that none of the artifact categories could be considered linearly separable.

A second study again examined the linear separability of natural and artifact concepts, but this time in a multiclass environment. We made use of two datasets, one comprising 129 animals divided over five natural categories, described by 764 features, and the other containing 166 artifacts divided over six artifact categories, described by 1295 features. We found that multiclass LSVMs could efficiently use these feature values to predict category membership: Prediction accuracy was high for both natural classes (up to 98.78% accuracy) and artifact classes (up to 99.29% accuracy). We found that all natural categories were linearly separable from one another, except for the fish and mammal categories, and that most of the artifact categories were linearly separable from one another as well.

In our third study, we investigated whether LSVMs can use similarity ratings to linearly separate the different names people give to types of movement. We examined two datasets: one with data from English-speaking students, and one with data from Dutch speaking students, each containing the dominant name and similarity ratings for 24 video clips depicting movement. We found that LSVMs

could use similarity ratings to predict the name participants give to a type of movement, with a maximal predictive accuracy of 95.25% for the English dataset, and 79.7% for the Dutch dataset. Additionally, we found that for both datasets, some of the categories were sufficiently linearly separable from one another, and some were not.

Conclusion

We demonstrated that linear support vector machines can be used efficiently to determine the relative linear separability of semantic concepts. We showed how LSVMs can use feature values or similarity ratings to predict category membership, and how we can use the LSVMs misclassifications to determine the extent of the linear separability of the tested categories.

Acknowledgments

The first author was funded by research grant OT/10/024.

References

- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3), 273-297.
- Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, 40(2), 230-262.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207-238.
- Rosch, E. (1973). Natural categories. *Cognitive Psychology*, 4, 328-350.
- Ruts, W., Storms, G., & Hampton, J. (2004). Linear separability in superordinate natural language concepts. *Memory & Cognition*, 32(1), 83-95.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology*, 43(2), 161-204.
- Vapnik, V. (1982). *Estimation of dependences based on empirical data*. New York: Springer-Verlag.