Available online at www.sciencedirect.com
SCIENCE DIRECT®
Brain and Cognition 61 (2006) 269–279
Brain and Cognition
www.elsevier.com/locate/b&c

Category-specificity can emerge from bottom-up visual characteristics: Evidence from a modular neural network

Tim M. Gale [a,b,c,*], Keith R. Laws [c]

[a] Department of Psychiatry, QEII Hospital, Welwyn Garden City, UK
[b] School of Computer Science, University of Hertfordshire, Hatfield, UK
[c] School of Psychology, University of Hertfordshire, Hatfield, UK

Accepted 3 February 2006
Available online 15 March 2006

Abstract

The role of bottom-up visual processes in category-specific object recognition has been largely unexplored. We examined the role of low-level visual characteristics in category specific recognition using a modular neural network comprising both unsupervised and supervised components. One hundred standardised pictures from ten different categories (five living and five nonliving, including body parts and musical instruments) were presented to a Kohonen self-organising map (SOM) which re-represents the visual stimuli by clustering them within a smaller number of dimensions. The SOM representations were then used to train an attractor network to learn the superordinate category of each item. The ease with which the model acquired the category mappings was investigated with respect to emerging category effects. We found that the superordinates could be separated by very low-level visual factors (as extracted by the SOM). The model also accounted for the well documented atypicality of body parts and musical instrument superordinates. The model has clear relevance to human object recognition since the model was quicker to learn more typical category exemplars and finally the model also accounted for more than 20% of the naming variance in a sample of 57 brain injured subjects. We conclude that purely bottom-up visual characteristics can explain some important features of category-specific phenomena.

Keywords: Category-specific; Self-organising map; Neural network; Bottom-up processing; Visual overlap; Superordinate

1. Introduction

Most reports of category specific agnosia describe patients with impaired recognition of living things (e.g., animals, birds, and fruit) relative to nonliving things (e.g., vehicles, clothing, and furniture); while the converse pattern is reported much less frequently (for reviews, see Capitani, Laiacona, Mahon, & Caramazza, 2003; Laws, 2005). Such cases have been very influential in current thinking about visual object processing and the organisation of semantic memory. One issue that has received considerable attention is the role of structural overlap in the emergence of category specific performance. In broad terms, structural overlap is the extent to which items from the same superordinate category share similar visual representations (e.g., to what extent do different examples of fruit look similar?). The prevailing view has been that living thing categories have greater structural overlap than nonliving thing categories, and that the emerging visual crowding effect predisposes living things towards recognition error (Gaffan & Heywood, 1993; Humphreys & Forde, 2001; Tranel, Logan, Frank, & Damasio, 1997). The essence of this account is that when the human visual system is presented with an item from a visually crowded category, the increased competition between stored representations has an inhibitory effect on the activation of a distinct structural description. By contrast, structural similarity is viewed as being potentially advantageous for other tasks where uniquely identifying

* Corresponding author. Fax: +44 0 1707 365169.
E-mail address: t.gale@herts.ac.uk (T.M. Gale).

0278-2626/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.bandc.2006.02.001

information is not required (e.g., superordinate classification).

A critical factor in the debate about visual or structural overlap concerns how overlap is quantified. Several studies have produced measures of within-category visual overlap using largely the same source data, i.e., the Snodgrass and Vanderwart (1980) corpus of monochrome line drawings, but with markedly different predictions. Humphreys, Riddoch, and Quinlan (1988) derived a measure of contour overlap (CO) by placing line-drawn exemplars from the same superordinate category on top of each other, overlaying a grid, and calculating average overlap between pictures as a function of the amount of contour in each picture (at a gross visible level). This method produced greater CO for living things. Using a somewhat different approach, Tranel et al. (1997) measured the number of pixels falling within the maximal silhouette overlap (i.e., the common category silhouette when items from the same category were rotated) for five different categories. The greatest overlap occurred for fruits/vegetables, followed by vehicles, animals, musical instruments and, finally, tools/utensils; a pattern that does not fully correspond with the performance profiles of most category-specific patients. Since the Snodgrass and Vanderwart pictures have been used in the vast majority of category-specific studies published to date (reviewed in Laws, 2005), it is particularly important to examine the role of any variables that are strongly associated with this corpus. Recently, Laws and Gale (2002) derived a measure of Euclidean overlap (EO) by taking standardised, digitised versions of the Snodgrass and Vanderwart line-drawings and calculating the average Euclidean distance between each item and its within-category associates, i.e., the actual physical overlap. Although no overall living versus nonliving difference emerged, body parts and musical instruments were notable for their atypical profiles when compared to other living and nonliving categories respectively. Furthermore, when these two categories were excluded from the comparison, nonliving things emerged as showing greater pixel overlap (Laws, Gale, Frank, & Davey, 2002). Although the predictions made by EO run contrary to the notion that living things display greater visual crowding, this measure has nonetheless predicted behavioural data in several object processing tasks, including naming error rate (Laws & Gale, 2002) and naming latency for both degraded and non-degraded stimuli (Laws, Leeson, & Gale, 2002).

Turning specifically to the categories of musical instruments and body parts, these superordinates are renowned in the category-specific literature (Barbarotto, Capitani, & Laiacona, 2001; Parkin & Stewart, 1993) because they tend to produce anomalous response profiles. In particular, patients who have marked difficulties in naming living things often have similar problems with musical instruments, yet show spared naming for body-parts. The converse dissociation is rare but, arguably, this must be contextualized within the general scarcity of reported nonliving deficits. Whatever the pattern of association between musical instruments, body parts, and other categories, the trend emerging from three decades of patient research is that musical instruments tend to cause recognition difficulties in the majority of category-specific cases whereas recognition of body parts is usually spared (for a review, see Barbarotto et al., 2001; Capitani et al., 2003). This profile of impairment has proved difficult to account for within most models of semantic memory. For example, the sensory functional theory (SFT) of Warrington and Shallice (1984) proposes a multi-modal semantic store comprising perceptual and functional subsystems. SFT posits that living and nonliving things are predominantly specified by visual/perceptual and functional/associative properties respectively and that damage to one of these sub-systems may lead to emergent category-specific impairments. This account is appealing on grounds of parsimony but it is difficult to envisage why musical instruments would associate with living things under this account and not, for example, other nonliving categories such as tools (both tend to be held in the hand, both have very specific functions, both have parts directly related to their specific function, both require training in use, and so on). So, although SFT may account for a broad living vs. nonliving dissociation, it does not comfortably explain some of the more subtle patterns in patient data. Similarly, the organised unit content hypothesis (OUCH) model proposed by Caramazza, Hillis, Rapp, and Romani (1990) proposes that, relative to nonliving things, living things tend to share a greater number of semantic features (e.g., has eyes, has legs, has a tail, etc.) and that these correlated features tend to be represented in adjacent neural substrate. This would leave living thing categories more vulnerable to localised neural damage because they are prone to catastrophic loss of supporting features/knowledge. Similar damage would impact less on nonliving things because their properties are not highly correlated and hence are dispersed more widely in substrate. Again, while this account may predict a living vs. nonliving dissociation, it is difficult to see how body parts and musical instruments come to associate with living and nonliving categories respectively. Arguably, many musical instruments share few physical or functional features (for example, compare violin, drum, flute, and piano), so it is difficult to explain why this category is typically impaired alongside living thing categories where shared features are more numerous.

Several connectionist models of category specificity have been discussed in the neuropsychological literature. Typically, these simulate 'emergent property' accounts (e.g., SFT and OUCH: see Caramazza, 1998), whereby category specific impairments are proposed to emerge following damage to functional sub-systems that are delineated along dimensions other than living and non-living categories (e.g., Devlin, Gonnerman, Andersen, & Seidenberg, 1998; Farah & McClelland, 1991; Tyler, Moss, Durrant-Peatfield, & Levy, 2000). However, while these models have all demonstrated a living versus nonliving dissociation after simulated lesion damage, none of them were set-up to specifically examine the differential contributions of specific