COGNITIVE NEUROPSYCHOLOGY, 2003, 20 (3/4/5/6), 451–486

THE SIMILARITY-IN-TOPOGRAPHY PRINCIPLE: RECONCILING THEORIES OF CONCEPTUAL DEFICITS

W. Kyle Simmons and Lawrence W. Barsalou
Emory University, Atlanta, USA

Three theories currently compete to explain the conceptual deficits that result from brain damage: sensory-functional theory, domain-specific theory, and conceptual structure theory. We argue that all three theories capture important aspects of conceptual deficits, and offer different insights into their origins. Conceptual topography theory (CTT) integrates these insights, beginning with A. R. Damasio's (1989) convergence zone theory and elaborating it with the similarity-in-topography (SIT) principle. According to CTT, feature maps in sensory-motor systems represent the features of a category's exemplars. A hierarchical system of convergence zones then conjoins these features to form both property and category representations. According to the SIT principle, the proximity of two conjunctive neurons in a convergence zone increases with the similarity of the features they conjoin. As a result, conjunctive neurons become topographically organised into local regions that represent properties and categories. Depending on the level and location of a lesion in this system, a wide variety of deficits is possible. Consistent with the literature, these deficits range from the loss of a single category to the loss of multiple categories that share sensory-motor properties.

INTRODUCTION

Following certain insults to the brain, an individual may lose knowledge of some categories while retaining knowledge of others. Over the past 25 years, many clinical cases have been reported that demonstrate various patterns of category-specific deficits. Most commonly, patients lose knowledge of living things; in particular, animals. Less commonly, patients lose knowledge of nonliving things, such as manipulable artefacts. In some cases, patients lose a single category, such as living things. In others, they lose multiple categories, such as living things and musical instruments.[1]

In describing these deficits, researchers have used a variety of terms, sometimes referring to them as agnosias, and sometimes as *semantic deficits*. Whereas agnosia implies that a deficit reflects damage to a particular sensory-motor modality, semantic deficit implies damage to a higher-order conceptual representation. As will be seen, it is not always clear whether the loss of category knowledge

---

[1] "Category" will refer to a set of exemplars in the world, and "concept" will refer to a cognitive representation of the category in the brain.

Requests for reprints should be addressed to Lawrence W. Barsalou, Department of Psychology, Emory University, Atlanta, GA 30322, USA (Email: barsalou@emory.edu).

We are grateful to Alex Martin and Alfonso Caramazza for the opportunity to write this article. We are also grateful to Aron Barbey, Sergio Chaigneau, George Cree, Glyn Humphreys, Jay McClelland, and Ken McRae for helpful comments on earlier drafts. This work was supported by National Science Foundation Grants SBR-9905024 and BCS-0212134 to Lawrence W. Barsalou.

http://www.tandf.co.uk/journals/pp/02643294.html
DOI:10.1080/02643290342000032

reflects damage to sensory-motor modalities or to conceptual structures. Furthermore, particular theories tend to favour one term over the other, based on their particular assumptions about the conceptual system. To avoid terminology that favours one position, we adopt the term *conceptual deficit*, and use it inclusively in referring to deficits of both types. When known, the specific origins of a deficit should be clear from the surrounding text.

In this article, we develop a new account of conceptual deficits—conceptual topography theory (CTT)—that draws on A. R. Damasio's (1989) theory of *convergence zones*. To increase the explanatory power of convergence zone theory, we add the *similarity-in-topography (SIT) principle*. We then show that this revised theory is at least roughly compatible with the main theories of conceptual deficits in the literature. In particular, this theory predicts the particular findings that each theory explains most naturally, plus those that each theory alone cannot.

Theories of conceptual deficits

Three accounts of conceptual deficits currently dominate the literature: sensory-functional theory, domain-specific theory, and conceptual structure theory. We address each in turn.

Sensory-functional theory. The first attempt to explain conceptual deficits was the sensory-functional theory proposed by Warrington and her colleagues (e.g., Warrington & McCarthy, 1987; Warrington & Shallice, 1984), and later implemented computationally by Farah and McClelland (1991). Other formulations of this view include Humphreys and Forde's (2001) HIT theory, McRae and Cree's (2002) property profile theory, and Damasio's (1989) convergence zone theory.[2] Sensory-functional theory assumes that knowledge of a specific category is located near the sensory-motor areas of the brain that process its instances. Thus knowledge of living things is stored near sensory areas (especially vision), because sensory mechanisms dominate the processing of these categories. Conversely, knowledge of manipulable artefacts is stored near motor and somatosensory areas, because mechanisms involved in the functional use of artifacts dominate processing. Given the proximity of category knowledge to sensory-motor mechanisms, damage to a sensory-motor system may produce deficits in the specific categories that rely on them. Thus, a deficit for living things may arise from damage to brain regions that process sensory information, whereas a deficit for manipulable artefacts may arise from damage to regions that implement functional action.[3]

An important prediction from sensory-functional theory is that conceptual deficits should often include multiple categories. Because conceptual deficits arise from damage to a particular sensory-motor system, any category that relies on that system for processing its instances may suffer. Because a given sensory-motor system almost always serves multiple categories, more than one category should suffer if the system is damaged. Clearly other factors are relevant, such as how strongly one property type is correlated with other property types, and how dominant a given modality is for processing a category (e.g., Farah & McClelland, 1991). Regardless, sensory-functional theory generally predicts that multiple categories should show deficits when a given sensory-motor system is damaged. When sensory areas are damaged, multiple categories whose processing relies heavily on sensory analysis should suffer. When motor areas are damaged, multiple categories whose processing relies heavily on motor execution should suffer.

---

[2] Readers may find it confusing that an elaborated version of A. R. Damasio's (1989) convergence zone theory integrates all three theories, while being an instance of sensory-functional theories. In its original form, convergence zone theory is closest to sensory-functional theories because of the central roles that sensory-motor areas play in representation. Adding the SIT principle, however, changes the theory's character such that it exhibits central properties of domain-specific and conceptual structure theories as well.

[3] A major issue in the sensory-functional literature concerns the specific nature of functional information. McRae and Cree (2002) address this issue extensively. Barsalou, Sloman, and Chaigneau (in press) provide a detailed theory of the conceptual content that underlies function.

Domain-specific theory. According to the domain-specific theory, knowledge of a category is not distributed across the sensory-motor systems that process its instances. Instead knowledge of a category is stored in a circumscribed brain system that solely represents the category downstream from sensory-motor processing (e.g., Caramazza, 1998; Caramazza & Shelton, 1998; also see Capitani, Laiacona, Mahon, & Caramazza, 2003-this issue). According to this view, evolutionary contingencies led to at least three dissociable brain systems that represent animals, plants, and conspecifics, respectively. One further representational system of a more generic variety represents knowledge about categories having less evolutionary significance, such as artefacts.

Most importantly, the domain-specific theory predicts that conceptual deficits should typically affect a single category. When damage to the evolutionary system for a category occurs, a deficit in processing that category should follow, but not a deficit for others. In contrast to the sensory-functional theory, deficits should typically not occur for different categories that share the same type of core information. Instead deficits should be observed for individual types of categories.

Conceptual structure theory. According to this final view, all categories are represented in a single conceptual system (e.g., Caramazza, Hillis, Rapp, & Romani, 1990; Tyler & Moss, 2001; Tyler, Moss, Durrant-Peatfield, & Levy, 2000).[4] Prior to developing the domain-specific theory, Caramazza and his colleagues proposed the organized unitary content hypothesis (OUCH), which Tyler, Moss, and colleagues developed further into the conceptual structure theory. Similar to OUCH, the conceptual structure theory proposes that knowledge of objects is represented in a single semantic space. Within this continuous semantic space, the structure of categories arises from: (1) property correlations, (2) distinctive properties, and (3) interactions between property correlations and distinctive properties. Together, property correlations and property distinctiveness create a similarity metric within semantic space, such that the information for a particular category tends to lump together in a specific subregion. If the instances of a category share many correlated properties, its instances lump tightly in semantic space, whereas a category with diverse instances is distributed more broadly. Furthermore, as a category's distinctive properties increase, so does its distance to other categories.

The conceptual structure account proposes that three factors determine how susceptible a category is to brain injury. First, a category is more robust when its instances share many correlated properties. For example, different instances of mammals, such as dog, horse, and lion, share many properties, such as *four legs*, *fur*, and *bear live young*. If one property in a correlation is compromised, other properties in its correlational structure can help retrieve it, or stand in for it.

The second factor that increases a category's robustness is its number of distinctive properties. Categories become lost when their distinctive properties no longer distinguish them from similar categories. The more distinctive properties a category has, the more damage is necessary to make it indistinguishable.

Finally, a category's robustness increases as its distinctive properties increasingly correlate with other properties, thereby protecting distinguishing information from damage. In artefacts, distinctive structural properties are often correlated with corresponding functions, thereby increasing their robustness. For example, *curved prong* for corkscrew distinguishes it from other artefacts and is uniquely correlated with its function—opening wine bottles. In natural kinds, distinctive properties are often not uniquely correlated. For example, *stripes* distinguish zebra from related categories but are not uniquely correlated with other zebra properties (e.g., *hooves*), which occur for other categories

---

[4] Correlated and distinctive properties are also central to McRae and Cree's (2002) account of category deficits (also see Cree & McRae, in press). However, they attempt to explain the importance of these factors within an elaborated version of sensory-functional theory, rather than within a single semantic space.

(e.g., horse, mule). Thus, correlated and distinctive properties work together to represent categories, with the weak points of statistical structure being most prone to damage.

Reconciling conflicting theories. Three different theories purport to account for the same data. Sensory-functional theory proposes that the conceptual system is organised around different types of sensory-motor information. Domain-specific theory argues that the conceptual system is organised into separate processing systems for evolutionarily significant categories. Conceptual structure theory holds that the conceptual system is organised by statistical relationships between categories and their properties.

A typical assumption in such debates is that one theory is correct and the others false. The spirited discussion on conceptual deficits often appears to proceed on this assumption. Significantly, however, the history of science is replete with rapprochements between competing theories. Theories viewed as conflicting and mutually exclusive often turn out to each capture important elements of truth and to even be compatible. In the early and mid-nineteenth century, British physicists argued heatedly over whether light was a particle or a wave. As we now know, both were correct (Chen, 1995). A major debate over the mechanisms of biological respiration similarly ended in rapprochement (Bechtel & Richardson, 1993). Closer to home, early research on the localisation of vision in the brain hotly debated whether vision resides in the occipital vs. parietal lobes. As became clear, both views were correct (Bechtel & McCauley, 1999; McCauley & Bechtel, 2001).

We believe that such a resolution may exist for theories of conceptual deficits. Although these theories may appear mutually exclusive at first blush, each captures important aspects of the conceptual system. Given this system's complexity, much room exists for different theories to capture different elements of truth about it.

Convergence zone (CZ) theory. An existing sensory-functional theory offers a preliminary basis for synthesising the three theories of conceptual deficits: Damasio's convergence zone (CZ) theory (A. R. Damasio, 1989; A. R. Damasio & Damasio, 1994). As argued in the next section, additional principles must be added to CZ theory to accomplish this synthesis. Here we review the theory's basic components: (1) modality-specific representations in sensory-motor areas, and (2) the binding of modality-specific information in convergence zones.

CZ theory begins with a widespread assumption about sensory-motor processing. When an entity is perceived, it activates feature detectors in the relevant sensory-motor areas (from hereon, we will refer to systems of these detectors as *feature maps*). During visual processing of a cat, for example, some neurons respond to line orientations, vertices, and planar surfaces. Others fire for colour, orientation, and direction of movement. The overall pattern of activation across this hierarchically organised distributed system represents the entity in visual perception (Palmer, 1999; Zeki, 1993). Similar patterns of activation in feature maps on other modalities represent how the entity might sound and feel, and also the actions performed on it. A similar account can be given for introspective states that arise while interacting with an entity. For example, patterns of activation in the amygdala represent emotional states produced in response to perceived entities. A tremendous amount of neuroscience research has addressed the structure of these feature maps across modalities and the states that arise in them.

The second core component of CZ theory concerns storage of the states that arise in feature maps. When a pattern in a feature map becomes active, neurons in an association area bind the pattern's features for later use (from hereon, we will refer to neurons in association areas as *conjunctive neurons*). Damasio refers to these association areas as "convergence zones" and assumes that they exist at multiple hierarchical levels, ranging from posterior to anterior in the brain. Most locally, CZs near a particular modality capture patterns of activation within it. Thus association areas near visual processing areas capture patterns of activation there, whereas association areas near motor processing areas capture patterns of activation there. Further downstream, higher-level association areas in more anterior areas, such as the temporal and frontal lobes, conjoin patterns of activation *across* modalities. We will say more about the localisation of these areas later.

In a given CZ, conjunctive neurons link active neurons in feature maps together. According to A. R. Damasio (1989, p. 129), neurons notice "combinatorial arrangements" of neurons active in a feature map, and establish associative relationships between them. On viewing a chair, for example, a subset of the conjunctive neurons in a visual CZ captures the active features across the hierarchically organised feature maps in vision. On sitting in the chair, a subset of the conjunctive neurons in a motor CZ captures the active features that executed the movement. At a higher level, a subset of the neurons in a cross-modal CZ correlates the visual form of the chair with the action taken on it. As will become clear, we assume that higher-level conjunctive neurons conjoin lower-level conjunctive neurons. Rather than conjoining feature map neurons directly, cross-modal conjunctions conjoin earlier conjunctions, thereby conjoining states in feature maps indirectly. Thus the conjunctive neurons that link visual and motor properties do so indirectly by linking modality-specific conjunctive neurons in the visual and motor systems.

This architecture of feature maps and convergence zones has the following functional implication: Once conjunctive neurons in a CZ capture an active pattern in a feature map, these neurons can later reactivate the pattern in the absence of bottom-up sensory stimulation (A. R. Damasio, 1989). In recollecting a perceived object, conjunctive neurons in CZs re-enact the sensory-motor states active while encoding it. Similarly, when representing a concept, conjunctive neurons in CZs reactivate the sensory-motor states characteristic of its instances. A given re-enactment is never complete, and biases may affect its reactivation, but at least some semblance of the original state occurs. This basic idea of re-enactment is essentially the same as neural accounts of mental imagery (e.g., Farah, 2000; Grezes & Decety, 2001; Jeannerod, 1995; Kosslyn, 1994). The one difference might be that imagery typically creates more complete and vivid re-enactments than those in memory, conceptualisation, and comprehension (Barsalou, 1999, in press-b).

Barsalou (1999, in press-a) takes the basic CZ architecture and shows how a fully functional conceptual system can be built upon it. Within this architecture, it is possible to implement categorical inference, the type-token distinction, argument binding, recursion, productivity, propositions, and abstract concepts. Because these phenomena are largely unrelated to the issue of how conceptual deficits originate, however, we do not pursue them here. Instead our interest is in how an elaborated CZ architecture synthesises the sensory-functional, domain-specific, and conceptual structure theories, thereby providing a unified account of conceptual deficits.

Distributed representations in CZ theory. As is evident, the representation of a concept in CZ theory is distributed across many brain areas that serve a variety of processes. On the one hand, a concept's representation resides in subsets of conjunctive neurons across multiple association areas. On the other hand, these conjunctive neurons re-enact sensory-motor states in feature maps, often in several modalities at once. Thus the representation of a concept may span the brain from frontal and temporal association areas to a variety of sensory-motor feature maps.

An important issue concerns the roles that parts of a distributed representation play in conceptual tasks. A. R. Damasio (1989, p. 46) proposes that conjunctive neurons in CZs play no representational roles. Instead they only constitute a means of reactivating previously active patterns in feature maps. Consider the representation of chairs. According to Damasio, the conjunctive neurons that conjoin the visual features of chairs cannot function as a stand-alone representation of this category. Instead these neurons only serve to reactivate chair features in visual feature maps, which then constitute a conceptual representation of chairs.

We agree that conjunctive neurons play the important role of reactivating patterns in feature maps during imagery, conceptual processing, and other cognitive tasks (Barsalou, 1999).

Reactivating features provides a powerful means of representing and processing categories. Unlike Damasio, however, we believe that patterns of conjunctive neurons in CZs can also function as stand-alone representations, in particular, during automatised feed-forward processes such as categorisation. During the categorisation of familiar objects (e.g., chairs), active feature detectors feed activation into the conjunctive neurons that integrate chair features. These conjunctive neurons then feed activation to response systems, such as the system that vocally produces a category name (e.g., "chair"). In this chain of feed-forward processing, the pattern of active conjunctive neurons functions as a representation sufficient to produce a correct response—reactivating a feature map pattern is not necessary. Nevertheless, research has shown that the visual system may be reactivated when categorisations are difficult (e.g., Humphreys & Forde, 2001; Kosslyn, Alpert, Thompson, Chabris, Rauch, & Anderson, 1994; Kosslyn, Thompson, & Alpert, 1997). Under demanding conditions, re-entrant activation produces visual representations that guide further feature extraction.

An open question is whether patterns of conjunctive neurons can function as stand-alone representations in other tasks that require the construction, manipulation, and evaluation of conceptual representations. For example, can people construct novel conceptual combinations using patterns of conjunctive neurons alone, or must they reactivate feature maps and then perform conceptual combination on these simulations? Similarly, can people perform language comprehension by simply manipulating patterns of conjunctive neurons, or are simulations in feature maps required (cf. Glenberg & Kaschak, 2002; Spivey, Tyler, Richardson, & Young, 2000; Stanfield & Zwaan, 2001; Zwaan, Stanfield, & Yaxley, 2002)?

We conjecture that conjunctive neurons only function as stand-alone representations when they feed forward automatically to responses under highly routinised conditions, such as categorisation and word association. When nonautomatised strategic processing must be performed to construct, manipulate, or evaluate a conceptual representation, simulations in feature maps are necessary. We increasingly believe that the componential symbolic operations that underlie complex conceptual processing can only be performed on simulations, not on stand-alone patterns of conjunctive neurons (cf. Kan, Barsalou, Solomon, Minor, & Thompson-Schill, 2003-this issue).

Distributed representations in other theories. Distributed representations are certainly not unique to CZ theory. Many theories, especially those grounded in connectionism, assume that knowledge is distributed across multiple systems. In sensory-functional theories, for example, a concept's representation resides in different knowledge stores for visual vs. functional properties.

In most other theories, however, conceptual content does not exist in feature maps—content exists only in other systems that *recode* patterns of feature map activation. In sensory-functional theories, for example, a sensory store is assumed to contain conceptual content that has presumably been recoded from the original sensory systems. Similarly, a functional store is assumed to contain conceptual content that has been recoded from the motor system. Typically, pools of hidden units in these theories recode sensory-motor properties into new representations, which then function as stand-alone representations.

Consider Humphreys and Forde's (2001) HIT theory. In their distributed architecture, knowledge stores exist for properties on different modalities (e.g., see their Figure 3 on p. 474). Importantly, however, "descriptions" of these properties represent them in conceptual knowledge—not simulations of them in sensory-motor systems. Furthermore, Forde and Humphreys' re-entrant activation is primarily a way to support difficult categorisation—it does not function as a simulation mechanism for representing imagery and knowledge in general. Thus a key difference between our approach and others is that simulations of content in feature maps represent categories, not re-descriptions of that content into amodal description languages.

Furthermore, nothing like our similarity-in-topography (SIT) principle (described later) exists in other current theories. Applying this principle to

A. R. Damasio's (1989) convergence zone theory produces a novel architecture that can be mapped to neural systems in ways that other theories have not. Additionally, this architecture makes novel predictions about conceptual deficits.[5]

Knowledge stores in the domain-specific and conceptual structure theories appear to have a similar status. In the domain-specific theory, stores outside sensory-motor systems recode the sensory-motor content of specific categories. In the conceptual structure theory, descriptions stored outside sensory-motor systems represent this content in a continuous semantic space. Furthermore, nothing like the SIT principle organises conceptual content in these theories.

Overview. In the first of the remaining sections, we supplement A. R. Damasio's (1989) CZ theory with the SIT principle, along with its corollary, the variable dispersion principle. These principles allow CZ theory to explain critical findings on conceptual deficits, and also to synthesise the sensory-functional, domain-specific, and conceptual structure theories. Within this framework, the second remaining section presents the conceptual topography theory (CTT). In the third remaining section, we bring neuroscience evidence to bear on CTT's neuroanatomical predictions. In the fourth section, we list critical findings in the lesion literature that any model of conceptual deficits must explain. In the final section, we show that CTT's architecture naturally predicts these findings.

TWO SUPPLEMENTARY PRINCIPLES FOR CONVERGENCE ZONE THEORY

On its own, A. R. Damasio's (1989) CZ theory is not sufficient to explain what is known about conceptual deficits; nor is it sufficient to synthesise the sensory-functional, domain-specific, and conceptual structure theories. At least, two additional principles are necessary: (1) *the similarity-in-topography (SIT) principle*, and its corollary, (2) *the variable dispersion principle*.

The similarity-in-topography (SIT) principle

We offer the SIT principle in the spirit of a conjecture that rests on modest empirical support but that clearly requires further examination and evidence. Perhaps the strongest current source of evidence is the SIT principle's ability to explain conceptual deficits, as the final section illustrates. Modest independent evidence, however, comes from the electrophysiological and neuroimaging literatures, as will be seen. Clearly much additional evidence is necessary before accepting this principle. Nevertheless we propose it here, first, because it seems like a conjecture worth pursuing, and second, because of its ability to explain conceptual deficits and to integrate theories.

The SIT principle concerns the organisation of the conjunctive neurons in CZs. Essentially the SIT principle claims that categorical structure in the world becomes instantiated in the topography of the brain's association areas. Specifically, the SIT principle states that:

> The spatial proximity of two neurons in a CZ reflects the similarity of the features they conjoin. As two sets of conjoined features become more similar, the conjunctive neurons that link them lie closer together in the CZ's spatial topography.

Consider the faces of a human, a monkey, and an elephant, and their representation in Figure 1. While viewing a given face, large numbers of neurons distributed throughout visual feature maps

---

[5] At the very end of their article, Humphreys and Forde (2001, p. 475) introduce something along the lines of our SIT principle, which they then dismiss as not being needed in the formulation of HIT that dominates their account of the literature. We originally and independently arrived at this principle earlier after reviewing the literature on conceptual deficits during a 1999 graduate seminar and seeing its potential for explaining the literature. This current article continues our application of A. R. Damasio's (1989) convergence zone theory to perceptual symbol systems (Barsalou, 1999), and shows how adding the SIT principle explains conceptual deficits within this framework.

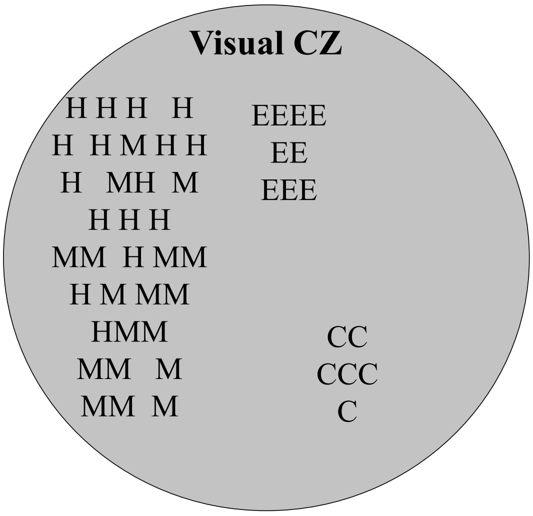

Figure 1. *Illustration of localised conjunctive neurons in a visual CZ for the features of a human (H), a monkey (M), an elephant (E), and a chair (C).*

become active to represent its features. Subsequently, neurons in a visual CZ conjoin these features by associating the respective feature map neurons. Of interest is the spatial proximity of the conjunctive neurons for the three faces. According to the SIT principle, the populations of conjunctive neurons for the human and monkey faces lie closer together on average in the visual CZ than does the population for the elephant face. Furthermore, the conjunctive neurons that represent all three faces lie closer together than the conjunctive neurons that represent some completely different type of object, such as a chair. In general, the topographic proximity of two conjunctive neurons reflects the similarity of the features they link.

The SIT principle has an important connection to the conceptual structure theory. As described earlier, this theory proposes that semantic space is "lumpy." Where clusters or "lumps" of correlated properties arise in the space, categories typically exist. Thus, one region of semantic space contains the correlated properties for mammals, whereas another contains those for plants. The SIT principle reproduces these lumps in the spatial topography of CZs. Whereas one population of neurons in the visual CZ tends to capture conjunctions of visual properties for animals, another population—perhaps partially overlapping—tends to capture conjunctions of visual properties typical of plants. Via the SIT principle, lumpiness at the conceptual level manifests itself physically at the topographical level of the CZ.[6]

Although statistical structure is central for both conceptual structure theory and the SIT principle, the actual implementation of this structure varies in the two approaches. Whereas the SIT principle realises statistical structure topographically in the brain's association areas, conceptual structure theories do not. Consequently, the two theories make somewhat different predictions about how statistical structure produces conceptual deficits, as we will see later.

Problems in specifying conjunctive similarity. To what feature combinations are conjunctive neurons tuned? How do conjunctive neurons acquire their tunings? We do not have space to address these issues in detail but briefly note their potential implications for our account.

In principle, any given feature could co-occur with any other feature. In the daunting face of this possibility, theorists often argue that biases must constrain the feature combinations that organisms anticipate (e.g., Murphy, 2002; Murphy & Medin, 1985). Perhaps some tunings evolve to anticipate feature combinations that are evolutionarily significant, such as those for animals. Similarly, other conjunctive neurons may be predisposed to link the feature correlations that pair words with their meanings, emotional expressions with emotional feelings, motor actions with visual outcomes, etc. The well-known fact that some associations are learned more easily than others is consistent with this proposal (e.g., Garcia & Koelling, 1966). By no

---

[6] As Barsalou (1983, 1985, 1991) shows, many categories do not share correlated properties (e.g., ad hoc categories). Because the lesion literature has not addressed these categories, we do not address them here. We assume, however, that other mechanisms beyond those in conceptual topography theory are necessary to account for ad hoc categories and other abstract concepts (see Barsalou, 1999, for accounts within the framework of perceptual symbol systems).

means do we imply that this categorical knowledge is “innate” or genetically encoded! To the contrary, we merely suggest that some tunings evolve that *anticipate* useful knowledge—category knowledge is not represented genetically. Only after actual experience utilises these conjunctive neurons does category knowledge develop epigenetically (cf. Elman, Bates, Johnson, Karmiloff-Smith, & Plunkett, 1996).

Clearly, though, people learn many categories that have little or no evolutionarily significance. Thus additional principles must determine the tunings of the conjunctive neurons that link the features of these categories. One possibility is that some tunings develop through learning. As new categories are encountered, free contiguous conjunctive neurons are recruited and tuned in the process. Another possibility is that constraints from higher-level knowledge establish tunings (cf. Murphy & Medin, 1985). For example, the belief that an organism’s internal organs control its external behaviour might tune conjunctive neurons to these particular correlations.

Topography as a general organising principle. Substantial evidence exists for topographic organisation in feature maps, ranging from the visual system to the motor, somatosensory, and auditory systems. In these areas, feature maps are often laid out according to the physical structure of the world. Such organisation in feature maps does not concern us here—our interest is in the topographic organisation of association areas. Nevertheless, it is intriguing to consider the possibility that topographic organisation underlies both. If so, this would suggest that topographical mapping constitutes a fundamental principle of brain organisation at multiple levels.

The variable dispersion principle

We do *not* assume that the conjunctive neurons for a category reside in a neatly circumscribed topographical region that only contains conjunctive neurons for that category and no other. To the contrary, we assume that the conjunctive neurons for a category are dispersed in clumps, with clumps for other categories falling between. Furthermore, a given clump may contain conjunctive neurons used by more than one category. As described later, both electrophysiological and fMRI evidence support these conclusions. Figure 2 summarises them. Unlike Panel A, the conjunctive neurons for a category are typically *not* contiguous in a CZ. Panel B illustrates instead how noncontiguous clusters of conjunctive neurons represent a category.

The variable dispersion principle applies to these noncontiguous clusters of conjunctive neurons, stating that:

> In a CZ, the proximity of the noncontiguous clusters for a category reflects the similarity of its instances. As the instances of a category decrease in similarity, its noncontiguous clusters of conjunctive neurons become increasingly dispersed in the CZ’s spatial topography.

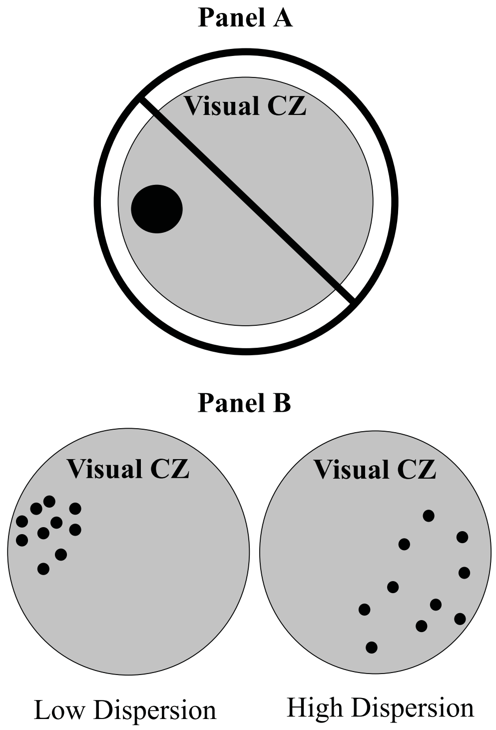

Figure 2. *Illustration of the variable dispersion principle. Panel A illustrates that the conjunctive neurons for a category are typically not contiguous in a CZ. Panel B illustrates how noncontiguous clusters of conjunctive neurons represent a category, with low dispersion for a category on the left, and high dispersion for a category on the right.*

Panel B in Figure 2 illustrates the variable dispersion principle, showing both low and high dispersion profiles for categories having similar vs. dissimilar instances respectively (e.g., mammals vs. artefacts).

The variable dispersion principle follows naturally from the SIT principle: Because similarity is instantiated in topography, the clusters of conjunctive neurons that represent a category lie closer together as within-category similarity increases. As will become clear later, this has significant implications for conceptual deficits. To the extent that the clusters for a category are tightly localised, a lesion can more easily disrupt the category. When a category's clusters are distributed more broadly, a single lesion is less likely to compromise them all.

CONCEPTUAL TOPOGRAPHY THEORY (CTT)

CTT develops a specific formulation of A. R. Damasio's (1989) CZ theory, coupling it with the SIT and variable dispersion principles. We begin with a general overview of CTT, and then turn to its specific components.

Overview

As Figure 3 illustrates, CTT contains the same configuration of four sub-systems on each of the six sensory-motor modalities, and also for emotion. Specifically, each modality contains *feature maps analytic CZs, holistic CZs*, and *modality CZs*. Although these four subsystems are only shown for the visual system, an analogous set is assumed to exist for each other modality. As Figure 3 further illustrates, cross-modal CZs integrate the modality-specific CZs. Together, all of these subsystems implement a core component of the human conceptual system.[7]

CTT further assumes that the lexicon is closely coupled with the conceptual system. Figure 3 does not illustrate a separate lexicon, however, because CTT assumes that the same modality-specific systems and cross-modal CZs that implement the conceptual system also implement the lexicon (as described later). The remainder of this section describes CTT's specific sub systems in greater detail.

Feature maps

Each modality is assumed to contain feature maps that code the content of modality-specific states. Thus feature maps code colour, line orientation, pitch, physical pressure at bodily locations, and so forth. It almost goes without saying that the perception of an object produces a tremendous amount of activity across multiple feature maps. Furthermore, because a category's instances vary, their feature map representations vary as well. As will become clear later, variability across a category's instances has implications for how the CTT model operates, and for how it explains conceptual deficits.

Analytic CZs

Thus far we have failed to address an important issue regarding CZs: Exactly what properties does a CZ capture for a category instance? A simple-minded possibility is that a CZ stores a complete image of them all. As Barsalou (1999) discusses, however, there is little evidence for complete images of this sort, and there are substantial theoretical problems as well (also see Hochberg, 1998; Wolfe, 2000).

Another much more likely possibility is that a CZ selectively stores information about *properties* of an entity, guided by selective attention. Barsalou (1999, in press-a) suggests that selective attention identifies informative subregions of a perceived entity and stores their content as perceptual symbols. Once selective attention focuses on a configuration of features in a particular subregion, the features become bound together in an analytic CZ (cf. Treisman & Gelade, 1980). As a result of being

---

[7] As mentioned earlier, a variety of additional components are necessary to handle intuitive theories, ad hoc categories, abstract concepts, and so forth.

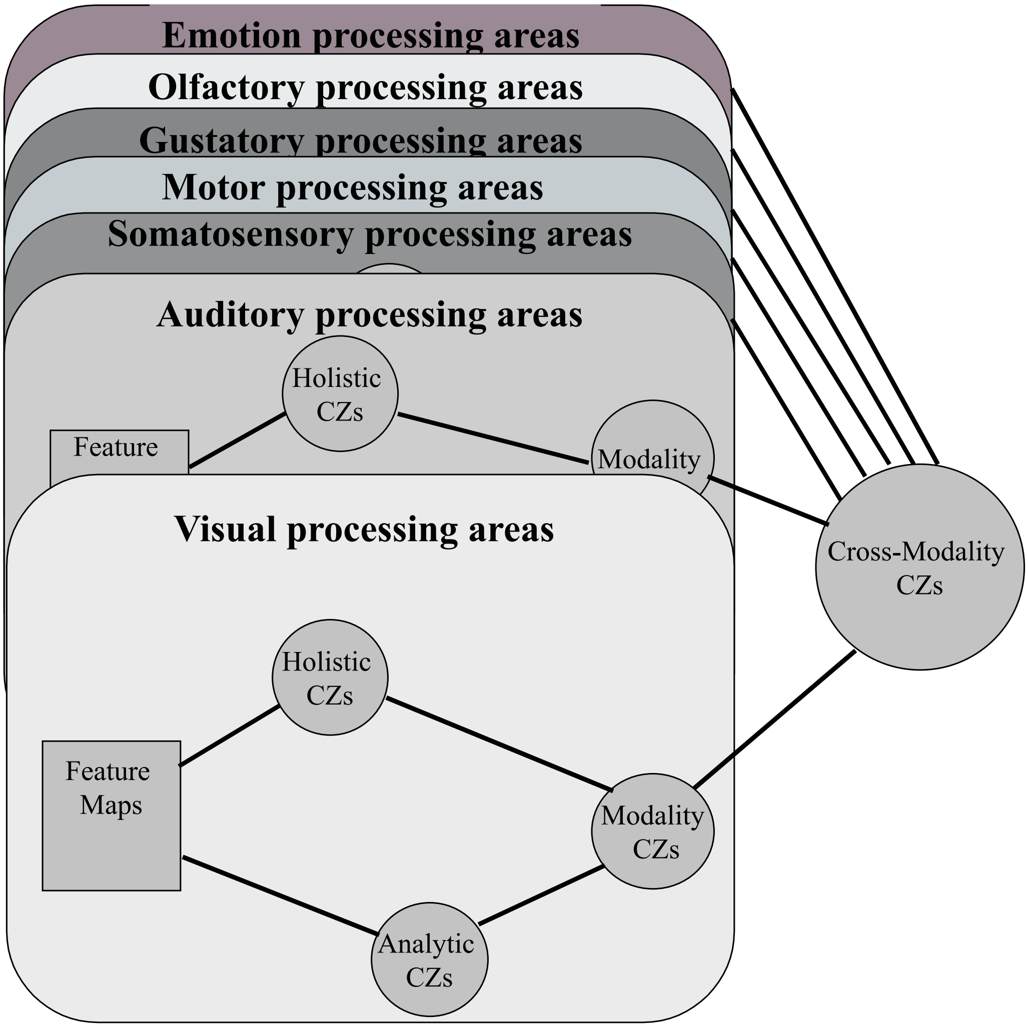

Figure 3. *The architecture of conceptual topography theory. The same configuration of four subsystems occurs in each of the six sensory-motor modalities and also for emotion, with cross-modality CZs integrating information across modalities.*

bound together, the features form an analytic property, which can then be used for a variety of conceptual tasks, such as categorisation and inference.

Consider *wheel* for car. On perceiving a car, the edges of a wheel are represented in visual feature maps. If selective attention focuses on this region of the perceived car, conjunctive neurons in an analytic CZ capture the features in this region. Later, on reactivating these conjunctive neurons, the visual representation of this particular wheel is partially re-enacted. As the perceived wheels of subsequent cars similarly receive attentional processing, they activate overlapping conjunctive neurons in the analytic CZ, thereby linking the visual features across different wheels to each other. The result is what Barsalou (1999) and Solomon and Barsalou (2001) refer to as a *property simulator*, namely, a system that can simulate the various forms a property takes in different categories. As the simulator for wheel develops, it produces simulations of different wheels in different objects, such as cars, bicycles, and skates. Furthermore, on perceiving a new car, processing its wheels might activate the simulator for wheel, which in turn might contribute to categorising the object as a car, given that the analytic property and the category are positively correlated.

As a result of such learning, knowledge of analytic properties develops. For car, such properties might include *wheels*, *doors*, *windshields*, *engines*, *exhaust fumes*, *steer*, and so forth. It is essential to note that analytic properties are not the same as features in feature maps—indeed, we use *feature* for one, and *property* for the other. Whereas the features in feature maps are sensory, supporting the construction of perceptions, analytic properties are conceptual, supporting the representation of cognitive-level categories. Whereas the features in feature maps are largely processed independently of attention, analytic properties typically result from applying attention to perception, thereby allowing conjunctive neurons to integrate the attended features. The results of this process are the "high-level" properties that constitute conceptual knowledge.

We assume that analytic properties arise on every modality, such as *loud* in audition, *pungent* in olfaction, and *sad* in emotion. Analytic properties need not have names. In many cases, feature configurations may become fused into unnamed analytic properties, such as the diagnostic properties that determined the categorisation of Martian rocks in Schyns and Murphy (1994). Many tastes and smells similarly appear unnamed. In general, any configuration of features that attention can select is a potential candidate for an analytic property (Schyns, Goldstone, & Thibaut, 1998). Sometimes these configurations have names and some times they do not. What is common across all analytic properties is that they play roles in the processing of concepts.

Property organisation according to the SIT principle. In CTT, the conjunctive neurons in analytic CZs are organised according to the SIT principle. Notably, the SIT principle does not generally organise analytic properties by category. Instead, it generally organises them by *property type*. As the SIT principle states, two conjunctive neurons are spatially close to the extent that they conjoin similar information. At the level of analytic properties, this similarity manifests itself in how similar two sets of conjoined features are. Because two different shapes are more similar than a shape and a colour, the conjunctive neurons for the two shapes should lie closer together than the conjunctive neurons for the shape and the colour. Figure 4 illustrates this state of affairs. As can be seen, regions of conjunctive neurons develop that represent shape, colour, movement, and other property types.

Notably, strong organisation by category should *not* be apparent in analytic CZs. For example, the properties of animals should not be tightly clustered. Because animals have diverse types of visual properties, these properties should be distributed widely throughout analytic CZs in vision. Following the SIT principle, the large differences between property types should cause them to be dispersed broadly, rather than being adjacent. Furthermore the visual properties for animals should be heavily mixed with those for other categories that have similar property types. To the extent that multiple categories share shape, colour, texture, and movement, their properties should reside in highly similar regions, such that the categories are not distinguished topographically.

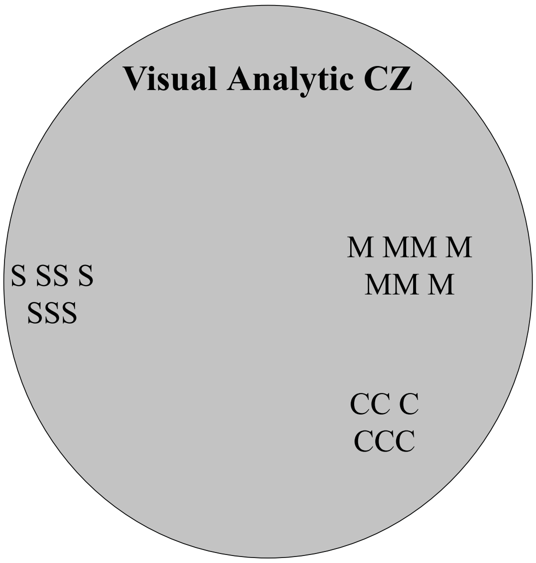

Figure 4. *Illustration of the conjunctive neurons in a visual analytic CZ that capture properties for colour (C), shape (S), and movement (M).*

Nevertheless, limited potential for category-specific topographic organisation exists in analytic CZs. Imagine that a category has a relatively unique type of property that no other category has, and that this property type is central to the category's identity. Following the SIT principle, the conjunctive neurons that capture these properties should be set off topographically from the conjunctive neurons that capture other property types. Furthermore, if the area holding these conjunctive neurons were lesioned, this category, and only this category, might show a deficit. In general, though, deficits at this level are unlikely to eliminate specific categories selectively, given that so many categories usually depend on a given property type. We pursue this issue later when we describe CTT's account of classic conceptual deficits.

Holistic CZs

Much research indicates that global and configural properties are extracted during the holistic processing of perceived entities. For example, the extraction of low spatial frequency information establishes the rough shape of an object and the location of its major and minor axes, in contrast to the extraction of high spatial frequency information that yields classic analytic properties (e.g., De Valois & De Valois, 1988; Morrison & Schyns, 2001). Similarly, holistic processing captures the configural relations between multiple analytic properties, such as the configuration of the eyes, nose, and mouth in a face (e.g., Macrae & Lewis, 2002). In general, holistic properties are relatively large, being distributed across much or even most of a perceived entity.

Research suggests that holistic processing often finishes before analytic processing (e.g., Breitmeyer & Ganz, 1976, 1977; Navon, 1977). Other research suggests that holistic processing dominates early category learning, with analytic processing refining subsequent learning (Kemler Neslon, 1989). Such differences between holistic and analytic processing suggest that different neural systems underlie them.

Analogous to analytic CZs, we assume that holistic CZs contain conjunctive neurons that capture activation patterns in feature maps for holistic properties. Following the SIT principle, we assume that two conjunctive neurons are topographically close to the extent that the holistic properties they link are similar. Thus one region of a holistic CZ in vision might capture overall shape information, another might capture information about major axes, and another might capture configural relations between analytic parts. As a result, holistic information about a category may be distributed across holistic CZs, rather than being localised spatially, analogous to the distribution of analytic properties across analytic CZs (Figure 4).

Holistic CZs may also play important roles in expertise. Recent evidence suggests that the acquisition of configural information underlies various forms of visual expertise, ranging from face perception, to greeble perception, to bird watching (e.g., Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Johnson & Mervis, 1997; Rackover, 2002). In some cases, this configural information allows experts to shift the basic level of categorisation down to the subordinate level. We assume that holistic CZs capture the relevant configural properties. Of interest, though, is whether the configural properties that underlie expertise are processed before analytic properties, like nonexpert holistic properties, or whether they are processed afterwards. Current theories of skill learning suggest that the storage of exemplars underlies expert performance, and that these exemplars are learned *after* the initial formation of rules during analytic processing (Logan, 1988). This suggests that different kinds of holistic properties may be extracted early vs. late during perception, with those extracted later being relevant to expertise. Regardless, we assume that these properties are organised around the SIT principle in holistic CZs.

Modality CZs

As the conceptual structure theory emphasises, the members of taxonomic categories tend to share correlated properties. As a result, the distributed pattern of activity for a category in analytic CZs provides a rich brew of statistical information just waiting to be extracted by higher-level systems. This is the task of modality CZs.

During perception, conjunctive neurons in modality CZs capture statistical regularities across the distributed activity in analytic and holistic CZs. Modality CZs neither process features (the job of feature maps), nor do they conjoin features into conceptual properties (the job of analytic and holistic CZs). Instead modality CZs capture correlations between various analytic and holistic properties. Because such property correlations are central to category structure, modality CZs are essentially category learners—they form category representations. At this point, however, modality CZs are only learning a category in *one* modality. Whereas modality CZs for vision capture the visual properties of a category, modality CZs for motor movement capture its action properties. As Figure 3 shows, modality CZs exist for each modality. As discussed in the next section, cross-modal CZs integrate the category's properties across modalities.

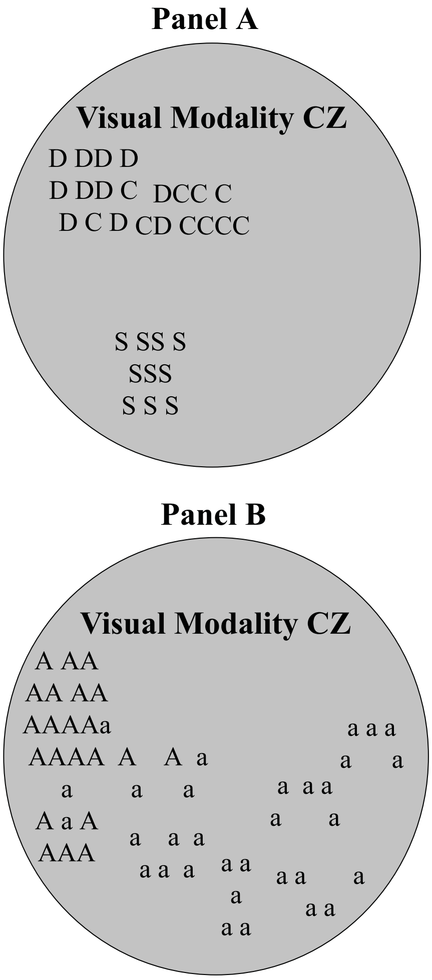

Figure 5. *Panel A illustrates the conjunctive neurons in a visual modality CZ that capture property correlations for dogs (D), cats (C), and shoes (S). Panel B illustrates the differential dispersion of the conjunctive neurons in a visual modality CZ for animals (A) vs. artefacts (a).*

Within a modality CZ, conjunctive neurons are tuned to particular conjunctions of properties within the respective modality. Thus, a conjunctive neuron in the modality CZ for vision might capture the co-occurring presence of *red*, *round*, and *smooth*. Following the SIT principle, conjunctive neurons should be near each other to the extent that they capture similar conjunctions of properties. Conversely, conjunctive neurons that capture completely different sets of conjunctive properties should be far apart. To make this clearer, consider the categories of dogs, cats, and shoes. As different dogs are perceived, similar but somewhat different properties will be linked for them in analytic and holistic CZs. Because the property conjunctions for different dogs tend to be similar, conjunctive neurons in the modality CZ that are topographically close will capture them, as Panel A of Figure 5 illustrates. Next consider cats. Because different cats have similar analytic and holistic properties, conjunctive neurons that are spatially close in modality CZs will capture them. As Figure 5a illustrates, however, the greater similarity within dogs and within cats than between them results in some separation between their conjunctive neurons in modality CZs. Finally, consider shoes. Again, conjunctive neurons will cluster for the category in modality CZs. Because shoes differ more from dogs and cats than dogs and cats differ from each other, the conjunctive neurons for shoes lie relatively far from those for dogs and cats. As this example illustrates, modality CZs topographically capture the statistical "lumps" of conceptual information important in conceptual structure theory.

Another important property of modality CZs follows from the variable dispersion principle, which states that the dispersion of conjunctive neurons for a category reflects the similarity of the category's instances. Panel B of Figure 5 illustrates this principle in modality CZs. First consider conjunctive neurons in visual modality CZs for animals. Because different animals share many visual properties, the conjunctive neurons that capture their property conjunctions should lie near each other. Conversely, consider the conjunctive neurons for artefacts. Because different artefacts have more diverse visual properties, the conjunctive neurons that capture their property conjunctions should be more dispersed. As discussed later, these differences in dispersion play important roles in explaining various conceptual deficits.

Cross-modal CZs

Whereas modality CZs conjoin properties within modalities, cross-modal CZs conjoin properties between modalities. Within a modality, the modality CZ creates a category representation but only on that modality. In contrast, cross-modal CZs create the complete representation of a category across modalities. Cross-modal CZs operate according to the same principles as modality CZs, except that the conjunctive neurons in cross-modal CZs link the conjunctive neurons in modality CZs, which in turn link conjunctive neurons in analytic and holistic CZs.

One might think that cross-modal integration of this sort is impossible: Because information on different modalities is represented in different formats (e.g., visual, auditory, olfactory, motor), they could not possibly be recoded into a common format. This is not a problem for CTT's cross-modal CZs, because conjunctive neurons in cross-modal CZs simply note the simultaneous firings of conjunctive neurons in modality CZs—they do not care about the format of the information ultimately represented in feature maps. Recoding into a common format is not necessary.

Complex property profiles. As we saw earlier, the sensory-functional theory divides conceptual properties into two types—sensory and functional—and proposes that conceptual deficits result from the loss of one property type or the other. In recent work, McRae and Cree (2002) showed that this simple division of property information is not adequate to explain the patterns of conceptual deficits in the literature (also see Cree & McRae, in press). However, a more detailed analysis of property types, based on a coding scheme developed by Wu and Barsalou (2003), does explain these deficits. Specifically, McRae and Cree coded categories for the extent to which they included external components, internal components, external surface features, internal surface features, materials, functions, and entity behaviours. McRae and Cree then computed how similar different categories are in possessing these different property types. For example, artefacts and foods are similar in both having *functions*. Most importantly, how categories cluster based on sharing these property types predicts the clusters of categories lost in category deficits.

This result bears on the organisation of conjunctive neurons in cross-modal CZs. To the extent that two categories are similar on one or more of McRae and Cree's seven property types, the closer their conjunctive neurons lie in cross-modal CZs, following the SIT principle. The closer two categories are, the easier it is for a localised brain lesion to compromise both.

One further implication of McRae and Cree's analysis concerns the relation of cross-modal CZs to modality CZs: It would be incorrect to presume that the topographic organisation of cross-modal CZs mirrors the topographic organisation of modality CZs. When two categories have similar properties on a specific modality, but different properties across modalities, their conjunctive neurons should be closer in modality CZs than in cross-modal CZs. For example, live animals and stuffed animals look similar visually, thereby causing their conjunctive neurons in visual modality CZs to lie close. However, because these categories differ substantially in auditory and olfactory properties, their conjunctive neurons in cross-modal CZs lie farther apart. As this example illustrates, the SIT principle interacts with complex property profiles to determine the proximity of conjunctive neurons at each level of analysis.

The lexicon

Thus far we have discussed the organisation of the human conceptual system. Clearly, though, words are tied closely to the use of concepts in humans. Although we believe that the lexicon plays a central role in conceptual processing, we do not show it as a separate system in Figure 3. This reflects our belief that the lexicon is distributed throughout the same brain systems as the conceptual system. Perhaps the main difference is that the lexicon's primary presence is in the left hemisphere, whereas the conceptual system is bilateral (e.g., Gainotti, Silveri, Daniele, & Giustolisi, 1995; Lambon Ralph, McClelland, Patterson, Galton, & Hodges, 2001).

At the level of feature maps, the perception of words in vision, audition, and touch is represented in the respective modalities. Maps that represent the acoustic features of phonemes in the auditory system detect spoken words; line and vertex feature maps in the visual system detect written words; analogous feature maps in the somatosensory system detect Braille. Clearly these feature maps may play roles in the processing of nonlinguistic stimuli as well. For spoken language, however, certain auditory feature maps may be relatively specialised for linguistic processing.

At the level of analytic and holistic CZs, localised areas may develop for linguistic stimuli in the respective modalities. Following the SIT principle, the conjunctive neurons that link analytic and holistic features for words may be relatively localised, because the conjunctions of features that constitute phonemes and letters are relatively unique. For example, the conjunction of a bilabial stop and voicing that specifies /b/ may be relatively unlike most other auditory properties.

At the level of modality CZs, the conjunctive neurons that integrate phonemes and letters into words may be relatively localised. Again, following the SIT principle, the uniqueness of these conjunctions may distinguish them. In auditory CZs, spoken words may be relatively distinct from other auditory stimuli, with analogous distinctions existing for written language and Braille. As a result, the conjunctive areas representing words may be separate topographically from the conjunctive areas representing concepts, creating the appearance that lexical and conceptual knowledge are processed in separate systems.

Finally, at the level of cross-modal CZs, subsets of conjunctive neurons link different forms of the same word across modalities. Thus, conjunctive neurons link the spoken and written forms of "chair" in the auditory and visual modalities, respectively, thereby providing a higher-order lexical entry of sorts. On this view, a "lexical item" is simply a population of conjunctive neurons that point to other conjunctive neurons that ultimately organise the respective features in feature maps.

As this account illustrates, the lexicon is distributed throughout the same brain systems that represent conceptual knowledge (except not as bilaterally). Because of the SIT principle, however, the representations of words are somewhat localised within these systems. The unique perceptual features of words isolate them, especially in the topographies of higher-order CZs.

Finally additional cross-modal CZs integrate conjunctive neurons for words with conjunctive neurons for concepts. Following the SIT principle, conjunctive neurons in these CZs are topographically close to the extent that the words and concepts they conjoin are similar. Specifically, the spatial proximity of two conjunctive neurons should be a function of two factors: (1) the similarity of the words they link in terms of their phonological and orthographic properties, and (2) the similarity of the concepts they link in terms of their analytic and holistic properties. An intriguing issue is whether word organisation or conceptual organisation dominates topographical clustering in these CZs, or whether clustering is an equal function of both.

NEURAL LOCALISATION OF CTT COMPONENTS

The previous section outlined a functional architecture for representing knowledge. As Table 1 illustrates, various anatomical predictions follow from this architecture. The current section reviews neuroscience evidence for these predictions. A later section analogously presents predictions for

Table 1. *Neuroanatomical predictions from conceptual topography theory (CTT)*

CTT component	*Neuroanatomical predictions*
Feature maps	1. Early topographic regions in sensory-motor processing streams detect and represent features that are utilised, not only in perception and action, but also in conceptual processing.
Analytic CZs	2. Just downstream from feature maps, modality-specific association areas conjoin conjunctions of features into analytic conceptual properties.
	3. Analytic CZs are organised topographically by property, not by category.
Holistic CZs	4. Also just downstream from feature maps, modality-specific association areas conjoin conjunctions of features into holistic conceptual properties.
	5. Holistic CZs are organised topographically by property, not by category.
Modality CZs	6. Modality CZs near the end of the processing stream on a sensory-motor modality conjoin analytic and holistic properties only on that modality.
	7. Modality CZs are organised topographically by category. The organisation of conjunctive neurons by similarity produces a rough categorical topography implicitly (there are no category-specific CZs).
Cross-modal CZs	8. Cross-modal CZs reside in brain regions having projections to multiple modality CZs, thereby supporting the integration of category properties across modalities.
	9. Modality CZs are organised topographically by category, again implicitly.
	10. Some cross-modal CZs may be specialised for conjoining properties across modalities for evolutionarily important categories, such as emotions, words, etc.

conceptual deficits and reviews the clinical evidence for them.

Feature maps

Feature maps are some of the most well-understood systems in the brain. Each modality contains hierarchically organised regions of neurons that capture specific features. Perhaps the most studied system is the macaque visual cortex. Starting in primary visual cortex, V1, and moving into areas V2, V3, V4, MT, and TEO, columns of neurons tuned to specific features process the visual stream (Tanaka, 1997a). Neurons in area MT, for example, extract movement direction, whereas neurons in V4 extract colour. Analysis in these regions occurs largely in parallel, with each area typically having projections to the others (Kandel, 1991).

Similar feature systems reside in other modalities besides vision. In somatosensory cortex, columns of neurons along the post-central gyrus are arranged somatotopically to detect physical stimulation at specific bodily locations (Dodd & Kelly, 1991). In the motor system, neurons along the pre-central gyrus are analogously arranged to generate movements in specific muscle groups (Ghez, 1991). In the auditory system, columns of neurons along the superior temporal gyrus are ordered tonotopically, from the lowest frequencies at the anterior end of primary auditory cortex, to the highest frequencies at the posterior end (Schreiner, Read, & Sutter, 2000).

The distribution of feature maps. An open question is how far feature maps are distributed down the processing stream. Some evidence suggest that the borders between visual feature maps and adjacent CZs are graded, not discrete. For example, Tanaka (1997a) reports that the percentage of neurons responsive to moderately complex feature conjunctions increases gradually as one moves from V4 to TEO to TE. Thus feature detectors and the conjunctive neurons that integrate them appear mixed together in transitional regions. Findings from the visual imagery literature can be interpreted as telling a similar story. In some studies, imagery activates primary sensory-motor areas, whereas in other studies it just activates secondary ones (e.g., Thompson & Kosslyn, 2000). In both cases, human subjects appear to experience imagery, which suggests that feature maps are distributed from primary to secondary visual areas.

Regardless of how far feature maps are distributed down processing streams, their task is to analyse complex sensory signals and extract the information to which they are tuned. Activation of specific neurons constitutes a signal that a particular feature is present in the input stream.

Analytic CZs

As described earlier, conjunctive neurons in analytic CZs link patterns of active feature detectors to represent analytic conceptual properties. According to the SIT principle, conjunctive neurons that link similar features lie near each other topographically in analytic CZs.

If this account is correct, two neuroanatomical predictions follow. First, neurons that conjoin features into higher-level conjunctions should lie just downstream from feature maps in the processing stream for a modality. Second, the proximity of conjunctive neurons in analytic CZs should follow the SIT principle. The evidence reviewed in this section offers preliminary evidence for both predictions.

Evidence from area TE in the macaque. Much evidence for the existence and placement of analytic CZs comes from work on the macaque visual system. According to Tanaka (1997a, 1997c), neurons along the occipital-temporal juncture, and also in the ventral and inferior temporal cortex, capture relatively complex feature conjunctions in earlier areas that contain feature maps (e.g., V1, V2, V4, TEO). Using a combination of single-unit recording and optical imaging techniques, Tanaka and his colleagues have focused on area TE in the inferior temporal cortex, mapping its functional and topographic organisation (e.g., Fujita, Tanaka, Ito, & Cheng, 1992; Kobatake & Tanaka, 1994; Tanaka, 1996, 1997b, 1997c; Wang, Tanaka, & Tanifuji, 1996). Consistent with CTT, these researchers show that TE has significant projections back to virtually all earlier visual processing areas (TEO, V4, V2, V1), with TE's cells tuned selectively to the presence of specific feature arrangements (Tanaka, 1997a).[8]

Conjunctive neurons in TE are organised into columns that stretch nearly the depth of the cortex, from layers two to six. Consistent with the SIT principle, these columns appear to be organised around the similarity of feature conjunctions. Fujita et al. (1992) inserted electrodes vertically into TE columns, isolated single neurons, and then recorded spike activity to three-dimensional objects. When a cell was particularly responsive to an object, the researchers produced a digital image of the object and used a computer to systematically pare down its features. Images were simplified until the minimum feature combination that maximally activated the neuron was identified, thereby isolating the feature combination for which the cell was tuned selectively. For example, a neuron that initially responded to an upside-down water bottle with a straw hanging out of the lid was eventually found to respond selectively to a vertically aligned ellipse with a small, thin line protruding from the bottom (for more on this technique see Tanaka, 1997a). Within a TE column, Fujita et al. found that all of its neurons responded to similar feature conjunctions. In further research using optical imaging, Wang et al. (1996) found that adjacent columns responded to related, but not identical, feature conjunctions.

Based on these findings, Tanaka (1997c, p. 524) speculated that "one function of the inferotemporal columns may be to augment feature variations around selected center points in the feature space; they may also facilitate computations among these variations. One or two principle components of the variations are mapped along the cortical surface." Such organisation follows directly from the SIT principle: Because two different orientations of a feature configuration are so similar, the conjunctive neurons that capture their respective features lie close to one another.

Although Tanaka and his colleagues report evidence consistent with the SIT principle, they note an important limitation. The sheer size of the feature space that must be laid out across a two-dimensional cortical surface means that the global organisation of TE cannot be a smooth function of similarity. Instead discontinuities should arise

---

[8] TE neurons are not the earliest cells that respond to co-occurring features—some cells in V4 respond in this way—but the proportion of such cells is significantly greater in TE (Kobatake & Tanaka, 1994).

where adjacent columns conjoin feature conjunctions that vary considerably. Nevertheless, two findings clearly support the SIT principle, first, the finding that neurons within a column represent a common conjunction, and second, that adjacent columns sometimes code similar conjunctions. Although it is physically impossible for an entire analytic CZ to reflect similarity, similarity still has considerable influence on the CZ's topography. We assume that such discontinuities are likely to arise in CZs at all levels of abstraction.

Distributed representation in analytic CZs. As we just saw, similar properties are often located near one another topographically. As described earlier, an implication of this organisation is that the representation of a particular category should be distributed widely over analytic CZs. The category's representation should not be localised at this level, because the SIT principle organises conjunctive neurons by property type—not by category. Because a given category typically has many types of properties, its representation in analytic CZs is distributed broadly. Because the members of a category tend to share properties, however, a shared set of distributed regions develops to represent these properties. Findings from two recent studies support this prediction.

Thomas, Van Hulle, and Vogels (2001) recorded from neurons distributed across macaque inferior temporal cortex (IT) during a simple categorisation task (e.g., whether a pictured object was a tree or something else). Using a Kohonen self-organising map to model the information-bearing properties of IT neurons, Thomas et al. found that, as a whole, these neurons contained sufficient information to distinguish the categories from one another.

Thomas et al. report two further findings in support of the conclusion that IT neurons code the analytic properties of categories. First, Thomas et al. observed some neurons that only responded to a single category, versus other neurons that responded to both categories, but to one more than to the other. These latter neurons, which were distributed throughout the recording region, played the most important role in categorisation. The authors concluded that these "broadly tuned feature detecting (almost certainly complex features) neurons" were central to the visual categorisation task (p. 198).

Thomas et al. further found that none of the recorded neurons responded to all of a category's exemplars. This finding is consistent with the idea that different exemplars of a category have somewhat different properties, linked by somewhat different sets of conjunctive neurons. Although all of these neurons may be useful in predicting category membership, none represents a property true of every exemplar, given that no such properties exist. Furthermore, the finding that some neurons responded to multiple categories is consistent with the fact that some properties occur in many categories, not just one.

Together, these findings support the conclusion that IT neurons capture the informational structure of visual categories at the level of conceptual properties. Consistent with CTT, a widely distributed set of conjunctive neurons in IT appears to represent the properties of a category, with some properties being unique for particular categories, and other properties being shared.

A second line of work further supports these conclusions. Using fMRI, Haxby et al. (2001) imaged the ventral temporal lobe in humans as they viewed pictures of cats, faces, houses, chairs, scissors, shoes, bottles, and nonsense patterns. Similar to Thomas et al., Haxby et al. found that a unique pattern of widely distributed activity predicted membership in each category. By observing a given pattern of brain activation, it was possible to predict the category viewed with 67% to 100% accuracy.

The structure of these activation patterns conforms to both Thomas et al.'s findings and to CTTs proposals about analytic CZs. First, the activation pattern that predicted each category contained topographically distributed regions of active neurons. This is again consistent with CTT's proposal that the multiple properties representing a category are distributed widely throughout analytic CZs. Within each activation pattern, one region was always more active than any of the others, namely, the *local maximum*. Interestingly, when the local maxima were excluded from analyses, the remaining active areas still predicted the category viewed. Thus multiple distributed regions represent a category's informational structure, not just one (i.e., the local maximum). Enough distinguishing information exists in the remaining properties to perform categorisation.

Haxby et al. further found that the activation patterns for different categories often overlapped, suggesting that a shared subregion represented a property common to multiple categories. Similar to Thomas et al.'s findings, these brain areas appear to capture a category's informational structure. Because categories often overlap in their properties, the brain regions that represent properties overlap as well.

Haxby et al. make an additional claim that is relevant to later sections on conceptual deficits: A lesion to the ventral temporal lobe could not produce a deficit in one and only one category. Because any given subregion responded to multiple categories, no well-placed lesion could knock out a single one. Haxby and colleagues conclude that "ventral temporal cortex has a topographically organized representation of attributes of form that underlie face and object recognition" (Haxby, Gobbini, Furey, Ishai, Schowten, & Pietriui, 2001, p. 2425). At this time, it is not known which particular properties these regions represent. Nevertheless, these regions do appear to represent category properties, given that they differentially predict category membership (if these regions were used to process all categories, they wouldn't be diagnostic). Regardless, these results are consistent with CTT's predictions about analytic CZs: Property information is distributed widely throughout these CZs, and multiple categories typically share properties.

Localised property representations in occipital-temporal cortex. In neuroimaging studies, Martin and his colleagues have extensively assessed the neural localisation of property representations (for reviews, see Martin, 2001; Martin & Chao, 2001; Martin, Ungerleider, & Haxby, 2000). In particular, these researchers have identified the brain regions that process object colour, form, and motion, for both animals and tools. Their findings indicate that brain regions just downstream from early vision areas capture these object properties. Furthermore, these properties appear organized according to the SIT principle, with similar properties lying nearer to each other than dissimilar properties.

For object colour, areas just anterior to V4 represent colour properties. For object form, areas in ventral occipito-temporal cortex represent *form* properties. For object motion, areas in the lateral temporal cortex represent *motion* properties. The localised representations of these three property types conform to the SIT principle's prediction that property types should cluster separately in analytic CZs. Different property types reside in different neural areas.

The representations of these property types for animals vs. tools offers even stronger support for the SIT principle. Form properties for animals reside right next to form properties for tools (lateral fusiform vs. medial fusiform gyrus, respectively). As the SIT principle predicts, form properties for animals and tools lie closer to each other than to colour and motion properties. The same relationship holds for motion properties. Motion properties for animals reside in superior temporal sulcus, whereas motion properties for tools reside just adjacently in middle temporal gyrus. Both of these areas are just anterior to area MT that underlies visual motion perception.

The overall layout of properties in occipital-temporal cortex provides evidence that the SIT principle organises analytic CZs. Not only are property types localised separately, the spatial proximity of property types reflects their conceptual similarity. The same property type for animals and tools (e.g., form) resides in different but adjacent areas, which lie closer together than two different property types for the same category (e.g., form and motion for animals).

Holistic CZs

We only address the possible location of holistic CZs in the visual system, but assume that similar CZs exist for other modalities. It is not clear where the feature maps for holistic processing in vision end, and where the CZs for capturing holistic properties begin. We suspect that holistic feature maps reside in early visual processing areas. For example, Zhou, Friedman, and Von der Heydt (2000) report that cells in V2 and V4 (and to a lesser extent in V1) code invariant boundaries of objects. In contrast, holistic CZs probably lie further along the visual processing stream. For example, Kanwisher and her colleagues report that the lateral occipital complex (LOC) and fusiform gyrus are important for processing configurations of object parts (Kanwisher, Woods, Iacoboni, & Mazziotta, 1997; Kourtzi & Kanwisher, 2001). Other research similarly localises configural processing in both the LOC (Hasson, Hendler, Ben Bashat, & Malach, 2001) and the fusiform gyrus (Gauthier et al., 1999). Currently we know of no evidence that the SIT principle organises these areas, although we predict that it does, just as we saw for the Martin studies in analytic CZs. Analogous to property CZs, conjunctive neurons in holistic CZs should be organised by the type of holistic property they encode. In addition, distributed regions of conjunctive neurons should capture the holistic property structures of different categories, with the same conjunctive neurons often representing holistic information across multiple categories.

Modality CZs

As we have seen, analytic and holistic CZs extract property information about categories. Furthermore, the topography of these CZs appears to reflect informational structure. Because properties vary across exemplars, no single conjunctive area becomes essential for a given category. Because multiple categories often share properties, a given conjunctive area is active for multiple categories. Thus analytic and holistic CZs contain the requisite information for a later processing stage that captures categories, not just their properties. In this spirit, Thomas et al. (2001, p. 190) state, “only one additional layer of processing is required to extract the categories from a population of IT neurons.” In CTT, modality CZs and cross-modal CZs extract these category-level representations. We discuss the neural localisation of modality CZs in this section, and the localisation of cross-modal CZs in the next.

As described earlier, CTT proposes that conjunctive neurons in modality CZs link the conjunctions of analytic and holistic properties for a category exemplar in a given modality (where these properties are represented by earlier conjunctive neurons that link features in feature maps). Furthermore, the SIT principle specifies that the topographical proximity of conjunctive neurons in modality CZs generally increases with the similarity of the property conjunctions they link.

Importantly, these neurons are *not* tuned to categories. Instead these neurons are tuned to specific conjunctions of properties that may happen to occur for a category’s exemplars, but not necessarily so. *There are no category-specific CZs*. Nevertheless, to the extent that a category’s exemplars share many properties, the conjunctive neurons in modality CZs that link them should cluster together topographically. As a result, an imprecise category-specific region may emerge for the category within the CZ. Although individual conjunctive neurons do not represent categories, larger groups of them may, although conjunctive neurons for other categories may be somewhat interspersed, following the variable dispersion principle. In the terms of conceptual structure theory, lumpy conceptual space becomes represented in lumpy neural space.

In the visual system, the modality CZ should be located at the end of the ventral pathway, most likely in the anterior inferior temporal cortex, or in parts of the perirhinal and/or entorhinal cortices. Staining and lesion studies indicate that these locations receive significant projections from the IT cortex (Tanaka, 1997a, 1997c), and that they also have projections to and from other polymodal brain areas, which will be important in the next section for cross-modal CZs (e.g., Murray & Richmond, 2001).

Perhaps the most compelling evidence for modality CZs in these visual areas comes from Kreiman, Koch, and Fried (2000a, 2000b). In these studies, single neuron responses to visual stimuli were recorded from electrodes implanted in the medial temporal lobes of human epilepsy patients. In Kreiman et al. (2000a), the stimuli were complex images of objects drawn from various categories: animals, foods, manipulable objects, cars, famous faces, emotional faces, spatial layouts, and complex patterns. Neurons in the entorhinal cortex, amygdala, and hippocampus exhibited category-specific changes in firing rates. Some responded selectively to all animal pictures in the stimulus set. Others responded selectively to all face pictures. Category-selective neurons were observed for every category presented.[9]

Kreiman et al. (2000b) replicated these findings and observed firing rate changes in category-specific neurons, not only during picture presentation, but also for mental imagery of previously viewed pictures. Interestingly, many of the same neurons were active during both the perception and imagery conditions, a finding that coincides nicely with CTT's assumption that perceptual and conceptual processing share neural mechanisms (Barsalou, 1999).

These findings support CTT's prediction that conjunctive neurons in modality CZs capture regularities in the properties for specific categories. At this time, we know of no direct evidence that the SIT principle organises these conjunctive neurons, although we predict that the more similar two categories are, the closer their conjunctive neurons should lie in modality CZs. The later section on conceptual deficits provides indirect support for such organisation.

Cross-modal CZs

As we just saw, CTT assumes that modality CZs integrate the properties of a category on a single modality. At that level, multiple representations of the category exist on all the modalities for which it has properties. What remains is to integrate information for the category across modalities. Cross-modal CZs play this role.

Perirhinal cortex. To accomplish cross-modal integration, cross-modal CZs must be located in brain areas that are highly connected to lower-order convergence zones for the various modalities. One region with such connections is the hippocampus. However, because patients with near total loss of the hippocampus often maintain conceptual knowledge (e.g., Mishkin, Vargha-Khadem, & Gadian, 1998), this structure is an unlikely candidate for the respective cross-modal CZs. The perirhinal cortex is a more likely possibility. Accumulating evidence indicates that perirhinal cortex is important for (1) integrating information about objects across sensory modalities, and (2) establishing associative relations between different concepts (e.g., Murray & Richmond, 2001). For example, lesion studies with monkeys demonstrate that perirhinal cortex is required for learning arbitrary associations between visual objects and the taste/smell of particular foods (Parker & Gaffan, 1998). Furthermore, monkeys with perirhinal lesions are impaired at visually identifying foods that taste good. Together these results suggest that perirhinal cortex integrates visual and gustatory information, implicating this area as a cross-modal CZ.

Perirhinal lesions similarly impair the ability to link an object's auditory and tactile properties with its visual and gustatory properties (e.g., Goulet & Murray, 2001; Higuchi & Miyashita, 1996; Holdstock, Gutnikov, Gaffan, & Mayes, 2000; Murray, Gaffan, & Mishkin, 1993; Parker & Gaffan, 1998). Based on these accumulating findings, Murray and Richmond (2001, p. 191) conclude that the perirhinal cortex "is an essential part of a system for storing fact-like information about objects . . . [and that] this 'object knowledge' system appears to be analogous to a semantic memory system in humans."

In this spirit, many models of conceptual deficits in humans propose "encyclopaedic knowledge" stores that are tantamount to associative, fact-like representation systems. In CTT, cross-modal CZs such as the perirhinal cortex implement this type of knowledge via associations between multiple concepts across modalities. Thus the knowledge

---

[9] Such neurons should not be confused with the infamous grandmother cell. As should be clear in the text, we assume that *populations* of conjunctive neurons integrate category features—not just a single conjunctive neuron. Also a given conjunctive neuron is likely to participate in the representation of *multiple* categories—not just one.

that a beaver builds a dam may be mediated by conjunctive neurons in the perirhinal cortex that link concepts for beavers, build, dam, trees, river, etc. Notably, however, cross-modal CZs do not literally contain this conceptual content, as is typically the case in current theories that employ knowledge stores. Instead conjunctive neurons in these CZs point to a hierarchy of lower-order conjunctive neurons that eventually activate features in feature maps. Ultimately the representation of a fact is a complex simulation of its content involving multiple concepts on multiple modalities.

Cross-modal CZs in emotion. Recent research suggests that specific areas organise emotion representations across modalities. For example, Anderson and Phelps (2000) suggested that somatosensory areas in the right hemisphere integrate "functionally distinct affective maps at different levels of organization" (also see Adolphs, Damasio, Tranel, Cooper, & Damasio, 2000). In particular, these CZs appear to integrate facial and bodily information for particular emotions.

Following A. R. Damasio's (1994) somatic marker theory, various researchers have similarly found that orbito-frontal cortex integrates emotional and conceptual information (e.g., Davidson, 2000; Davidson, Jackson, & Kalin, 2000). Specifically, conjunctive neurons in orbitofrontal cortex link a category's properties with associated emotional states (somatic markers). When a new category instance is encountered, the category's properties become active, which then activate its somatic markers, via conjunctive neurons in orbitofrontal cortex. The average of the activated somatic markers then provides a plausible emotional response to the new instance.

Besides providing evidence for cross-modal CZs, this work on emotion suggests an important constraint on them: Cross-modal CZs may often be tailored to conjoin particular sets of modalities that serve an important role in the cognitive system. For example, because different sources of information about emotion must be integrated, cross-modal CZs may develop in right somatosensory cortex for this purpose. Similarly, because conceptual and emotional information must be integrated, cross-modal CZs may develop in orbito-frontal cortex. Such cases suggest that some cross-modal CZs develop to integrate specific sources of modality-specific information that serve some purpose.

Cross-modal CZs in language. Similar sorts of cross-modal CZs appear tailored to language. Based on lesion studies, Tranel and his colleagues found that left prefrontal and premotor regions underlie the retrieval of words for actions, whereas left anterior and inferior temporal regions underlie the retrieval of words for concrete objects (Tranel, Adolphs, Damasio, & Damasio, 2001; Tranel, Damasio, & Damasio, 1998). In these studies, lesion patients had to name actions or concrete objects, thereby requiring them to link conceptual representations of categories with the respective words. When patients had lesions in left prefrontal and premotor regions, they could not make these mappings for action words; when they had lesions in left anterior and inferior temporal regions, they could not make these mappings for object words. These findings suggest that the respective brain areas constitute cross-modal CZs that link concepts and words.

Summary

Findings in the neuroscience literature support predictions of the CTT architecture. We hasten to add, however, that this evidence is open to alternative interpretations. Furthermore, we are the first to agree that much more evidence is necessary before embracing CTT with high confidence. Although we find the existing evidence encouraging, much further research is clearly needed. The evidence that initially motivated CTT—conceptual deficits—is the topic of the next two sections.

CONCEPTUAL DEFICITS

We next describe six findings for conceptual deficits that we believe any account must explain. In the subsequent section, we systematically step through CTT, addressing how damage to its various components produces these findings.

Finding #1: Multiple-category deficits that reflect property profiles

In many cases, patients who suffer a conceptual deficit lose more than one category. In their review, McRae and Cree (2002) list seven common forms that multiple-category deficits take: (1) Multiple categories that constitute creatures, (2) multiple categories that constitute nonliving things, (3) fruits and vegetables, (4) fruits and vegetables with either creatures or nonliving things, (5) foods with living things, (6) musical instruments with living things, (7) gemstones with living things.

The standard goal of sensory-functional theories has been to explain these patterns (e.g., Farah & McClelland, 1991; Humphreys & Forde, 2001; Warrington & Shallice, 1984). As McRae and Cree note, however, the standard forms of sensory-functional theory cannot explain all seven (also see Caramazza & Shelton, 1998). Standard theories typically distinguish only between sensory and functional properties, which provide too few parameters for explaining all seven patterns. McRae and Cree's contribution is to extend the number of critical property types significantly, with each pattern resulting from loss to one or more property types.

This set of findings offers two important constraints on theories of conceptual deficits. First, a conceptual deficit can include multiple categories, not just one. Second, when multiple categories are lost, they reflect damage to specific property types that the lost categories share.

Finding #2: Single-category deficits

At the opposite end of the deficit continuum lie single-category deficits. Rather than losing multiple categories, some individuals lose just one (Capitani, Laiacona, Mahon, & Caramazza, 2003-this issue; Caramazza & Shelton, 1998; Hart & Gordon, 1992). Furthermore, some of these patients lose the ability to process the category across all the sensory modalities tested—the deficit is not limited to a single modality. For example, EW demonstrated impaired recognition, naming, and knowledge comprehension for animals, whether tested in the visual or auditory/verbal modalities (Caramazza & Shelton, 1998). Her deficit cannot be attributed to low familiarity and frequency of animals, given careful control of these factors. In addition, perceptual deficits did not accompany her conceptual deficits, thereby implicating higher-order conceptual structures.

Other patients exhibit different single-category deficits. One patient reportedly lost knowledge of number (Cipolotti, Butterworth, & Denes, 1991); another lost knowledge of medical concepts (Crosson, Moberg, Boone, Rothi, & Raymer, 1997); yet another lost knowledge of musical notation (Cappelletti, Waley-Cohen, Butterworth, & Kopelman, 2000). Because of such cases, theories must account for single-category deficits as well as multiple ones.

Finding #3: Deficits for nonliving things are infrequent

Relative to deficits for living things, deficits for nonliving things are relatively infrequent. Reports of the former are much more common than the latter, as reviewers of this literature often note (Caramazza & Shelton, 1998; Humphreys & Forde, 2001; Tyler et al., 2000).

Finding #4: Deficits occur at the superordinate level, not at the basic level

Conceptual deficits are typically characterised by the loss of relatively large categories, such as superordinates, that reside toward the top of taxonomic hierarchies (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). To our knowledge, there are no reported cases of conceptual deficits for basic or subordinate categories. For example, deficits occur for categories like musical instruments, jewels, artefacts, or animals, but not for individual categories like violins, diamonds, chairs, or dogs. The relatively broad loss of categorical knowledge may offer an important constraint on theories of conceptual deficits.

Finding #5: Equal loss of functional and visual properties in some living things deficits

The sensory-functional theory asserts that knowledge of living things depends more on visual/sensory information than on functional information, whereas knowledge of nonliving things depends more on functional information. As a consequence, the sensory-functional theory predicts that living things deficits should be more associated with impairments in visual/sensory properties than with impairments in functional properties. The literature does not always support this prediction. Many researchers have reported patients with living things deficits whose knowledge of functional properties was just as impaired as their knowledge of visual/sensory properties (Capitani et al., 2003-this issue; Caramazza & Shelton, 1998; Funnell & De Mornay Davies, 1996; Laiacona, Barbarotto, & Capitani, 1993; Lambon Ralph, Howard, Nightingale, & Ellis, 1998; Samson, Pillon, & De Wilde, 1998). We believe that this finding offers another important constraint on theories of conceptual deficits.

Finding #6: Loss of visual properties, but a greater deficit for artefacts than for living things

A related problem for sensory-functional theory is a patient with an artefacts deficit whose knowledge of visual properties was more impaired than her knowledge of functional properties (Lambon Ralph et al., 1998). If living things depend more on visual properties than nonliving things, this patient should have exhibited a deficit for living things, not for artefacts. Coltheart, Inglis, Cuples, Michie, Bates, and Budd (1998) report another patient who offers a similar challenge, namely, significant loss of visual property knowledge, but equivalent knowledge of living things and non-living things. Similar to Finding #5, the relationship of living things and nonliving things to visual and functional information is far from systematic, offering a further challenge for theories.

EXPLAINING CONCEPTUAL DEFICITS WITH CTT

The CTT architecture makes natural predictions about the types of conceptual deficits that should be observed in the clinical literature. Three mechanisms underlie CTT's predictions:

1. The presence of modality-specific processing systems that begin with feature maps and proceed through holistic, analytic, and modality CZs.
2. The presence of cross-modal CZs that integrate properties for a category across all the relevant modalities.
3. The organisation of CZs according to the SIT principle, thereby organising analytic and holistic CZs topographically by property type, but organising modality and cross-modal CZs topographically by category.

Given this hierarchical, topographically organised architecture, diverse possibilities for localised lesions exist, along with correlated possibilities for conceptual deficits. Table 2 summarises CTT's predictions. The potential for diverse deficits is realised in the literature. Indeed, such diversity is one of CTT's primary predictions. No other current theory predicts the full range of deficits. In this next section, we step through CTT's systems, discuss the implications of lesions in each, and show how such lesions produce the six findings reviewed in the previous section.

Damage to feature maps

Damage to feature maps can disrupt both perceptual and conceptual processing. When a feature map is damaged, perception in the modality is likely to be impaired. Such perceptual deficits are not uncommon. For example, damage to V4 in the occipital lobe often produces impaired colour perception (Kandel, 1991). Similarly, bilateral lesions of the superior temporal gyrus can produce cortical deafness (Coslett, Brashear, & Heilman, 1984).

According to CTT, one function of feature maps is to help represent knowledge during conceptual processing. Thus damage to feature maps should produce conceptual deficits, not just

Table 2. *Deficit predictions from conceptual topography theory (CTT)*

CTT component	*Deficit prediction*
General prediction	1. A wide variety of conceptual deficits should be observed, each reflecting the part of the CTT architecture damaged. Damage to feature maps, analytic CZs, holistic CZs, modality CZs, or cross-modal CZs should each produce a different deficit. All possible deficits could, in principle, be observed.
Feature maps	2. A feature map lesion should not only compromise perception on the respective modality, it should also produce conceptual deficits across those categories that utilise these features conceptually (Finding #1).
	3. Feature map damage should primarily impair performance on strategic conceptual tasks, not on automatic ones.
Analytic CZs	4. Following a lesion to an analytic CZ, perception on the respective modality should remain relatively intact, while a conceptual deficit should arise for categories that rely on the analytic property damaged.
	5. A lesion to an analytic CZ should typically disrupt multiple categories that utilise the damaged analytic property (Finding #1). Deficits for single categories should not be observed, unless the damaged analytic property is unique to a category (Finding #2).
Holistic CZs	6. Following a lesion to a holistic CZ, low-level perception on the respective modality should remain relatively intact, while a conceptual deficit should arise for categories that rely on the holistic property damaged.
	7. A lesion to a holistic CZ should typically disrupt multiple categories that utilise the damaged holistic property (Finding #1). Deficits for single categories should not be observed, unless the damaged holistic property is unique to a category (Finding #2).
Modality CZs	8. When a lesion falls largely on the conjunctive neurons for one category in a modality CZ, a category-specific deficit may result (Finding #2), but only on that modality.
	9. When a lesion falls on the conjunctive neurons for multiple categories in a modality CZ, multiple categories should be compromised (Finding #1).
	10. A localised lesion should be more likely to eliminate the conjunctive neurons for a narrowly dispersed category on a modality than a broadly dispersed one (Finding #3).
	11. Because basic level categories tend to share many properties with each other, damage to a modality CZ will usually compromise multiple basic level categories. As a result, conceptual deficits will occur mostly for superordinate categories, not for basic ones (Finding #4).
	12. Contrary to the sensory-functional theory, a category can exhibit a deficit following damage to a secondary—not a primary—modality CZ (Finding #6) (e.g., an artefacts deficit following damage to a visual CZ, not to a motor CZ).
Cross-modal CZs	13. When a lesion falls largely on the conjunctive neurons for one category in a cross-modal CZ, a category-specific deficit may result across all modalities for which the category has properties (Finding #2).
	14. When a lesion falls on conjunctive neurons for multiple categories in a cross-modal CZ, multiple categories should be compromised (Finding #1).
	15. A localised lesion should be more likely to eliminate the conjunctive neurons for a narrowly dispersed category on a modality than for a broadly dispersed one (Finding #3).
	16. Because basic level categories tend to share many properties with each other, damage to a cross-modal CZ will usually compromise multiple basic level categories. As a result, conceptual deficits will occur mostly for superordinate categories, not for basic ones (Finding #4).
	17. Contrary to the sensory-functional theory, a category can exhibit deficits for properties on both primary and secondary modalities, following a lesion to cross-modal CZs (Finding #5) (e.g., equal losses of functional and visual properties for nonliving things).

perceptual ones. Indeed, classic agnosias offer such examples, where damage to a sensory system produces impairments in conceptual knowledge: For example, damage to colour processing areas produces deficits in colour knowledge (e.g., A. R. Damasio & Damasio, 1994; De Renzi & Spinnler, 1967).

If a particular modality is important for representing the properties of a concept, then damage to its feature maps should produce conceptual deficits on strategic tasks. Because animals depend heavily on visual properties, damage to visual feature maps should compromise knowledge of animals. Conversely, because tools depend heavily on motor properties, damage to motor maps should compromise knowledge of tools. As described earlier in the section on Distributed Representation in CZ Theory, the top-down reactivation of feature maps may only occur in conceptual tasks that require the strategic construction, manipulation, and evaluation of conceptual representations. Thus we only predict conceptual deficits on such tasks following feature map lesions. In routinised tasks, where automatised feed-forward pathways are sufficient for satisfactory task performance, we don't predict conceptual deficits in these patients.[10]

In general, damage to feature maps has the potential to produce multiple-category deficits, as described for Finding #1 in the previous section. Consistent with the sensory-functional theory, damage to a particular property type should have repercussions for any category whose instances rely on it. Because multiple categories typically use a given property type, multiple categories may often show deficits when it is damaged.

A potential problem for this account arises when patients lose perception on a modality but can still conceptualise properties on it. CTT explains such cases in several ways. First, loss of bottom-up pathways into feature maps may disable perception, while intact feature maps, along with intact top-down projections, still implement conceptualisation. Similarly, damage to feature maps may be incomplete, such that a partially intact feature map retains the capacity to support conceptual processing but not perception. Finally, it is unclear to what extent the intact conceptual performance of patients with damaged feature maps reflects their reliance on word associations to support seemingly conceptual processing (Kan et al., 2003-this issue; Solomon & Barsalou, 2003). Most, if not all, standard neuropsychological tests of conceptual knowledge are not designed to limit the degree to which patients can use word association strategies to perform tasks, thereby precluding assessments of knowledge.

Damage to analytic CZs

Lesions in analytic CZs should impair conceptual knowledge, and possibly late perceptual processes, while leaving early sensation largely intact. In sensation, the bottom-up channels that activate feature maps remain unaffected, as do the feature maps themselves, while conjunctive neurons in an analytic CZ that combine features into conceptual properties are lost. As a result, activating the corresponding feature conjunctions is no longer possible. The content of the respective properties can no longer be retrieved, and the concepts that utilise them suffer.

As we saw earlier, the SIT principle organises analytic CZs by property type. Topographically, the conjunctive neurons that represent similar properties reside in the same region of an analytic CZ. In vision, conjunctive neurons for shape reside together, whereas conjunctive neurons for visual motion lie elsewhere. As a consequence of this organisation, localised lesions should tend to disrupt knowledge of specific property types, thereby producing multiple category deficits (i.e., Finding #1). Consistent with the sensory-functional theory, all categories that utilise a property type should be affected by its loss. Unlike damage to feature maps, however, perception should remain relatively unaffected. Many cases of conceptual deficits in the literature fit this pattern, namely, multiple category loss with no perceptual deficits.

The multiple category deficits that result from damage to an analytic CZ should exhibit certain patterns (Cree & McRae, in press; McRae & Cree, 2002). Given the wide variety of property types represented in analytic CZs, there should not just be one or two patterns, there should be many. From McRae and Cree's perspective, deficit patterns that initially appeared puzzling become straightforward. For example, fruits/vegetables and nonliving things may sometimes be lost together when the shared property type of function is lesioned.

---

[10] One problem in diagnosing such cases is that debilitating perceptual deficits often mask accompanying conceptual deficits, making the latter difficult to detect. Thus it would be useful to assess more thoroughly whether patients with feature map damage tend to have corresponding conceptual deficits.

Second, when multiple categories are lost, they should share properties on the damaged modality, but not necessarily on any other. Because damage is to a single modality, the categories lost should exhibit salient similarities only on it. As we will see later, when the categories lost are similar on multiple modalities, this implicates damage to cross-modal CZs instead.

Typically, loss to a particular property type should not produce a single-category deficit. However, if a property type is relatively unique for a category—no other categories share it—then damage to this area of an analytic CZ could compromise only that category (i.e., Finding #2). In general, though, because properties are generally shared across categories, single-category deficits should rarely arise this way. Instead they should typically arise after lesions to higher-order CZs that are organised by category, as we will see shortly.

Damage to holistic CZs

Analogous to damage in an analytic CZ, damage to a holistic CZ should typically affect multiple categories that share the holistic property which is lost (i.e., Finding #1). In audition, for example, failure to apprehend the temporal envelope of a sound could reflect damage to a holistic CZ in the auditory system (Lorenzi, Wable, Moroni, Derobert, Frachet, & Belin, 2000). In vision, some patients with apperceptive visual agnosia may have suffered damage to a visual holistic CZ. Although these patients can typically perceive and recognise object parts, they cannot organise these parts into integrated perceptual wholes (e.g., Shelton, Bowers, Duara, & Heilman, 1994). When asked to copy a complex figure, these patients can occasionally draw and remember image fragments, but cannot organise them properly (e.g., Lezak, 1995). Such patterns suggest that feature maps and analytic CZs are intact, but that holistic CZs are not. Furthermore, the categories that pattern on such deficits should be similar on the same modality, but not necessarily on others.

Damage to modality CZs

Whereas the SIT principle organises analytic and holistic CZs by property type, it organises modality CZs by category to a considerable extent. Because the proximity of conjunctive neurons reflects the similarity of the properties they conjoin, neurons that capture the correlated properties of a category tend to be localised topographically. Following the variable dispersion principle, however, the conjunctive neurons for a category are mixed with those for other categories that share similar property correlations. When high similarity exists between a category's exemplars (e.g., mammals), the dispersion of its conjunctive neurons is relatively low. Although these conjunctive neurons are somewhat dispersed, they nevertheless cluster within a relatively well-circumscribed region of the modality CZ. Conversely, when low similarity exists between a category's exemplars (e.g., artefacts), the dispersion of its conjunctive neurons is relatively high, covering a relatively broad region of the modality CZ (Figure 5b).

Modality-specific single-category deficits. When a lesion in a modality CZ falls squarely on the region containing the conjunctive neurons for a particular category, loss of category knowledge on that modality should result. Although other categories may have conjunctive neurons in the region, the bulk may lie elsewhere, thereby preserving them (e.g., Farah & McClelland, 1991). The result is a single-category deficit, but only on one modality, thereby making it a special case of Finding #2 (i.e., where a category is lost across all modalities).

Consider the patient Michelangelo, who exhibited impaired knowledge of visual properties for animals, relative to artefacts, but whose knowledge of functional/associative properties for animals appeared relatively intact (e.g., where they live, what they eat, how they sound, etc.; Sartori & Job, 1988). CTT explains this pattern as the result of localised damage to the visual modality CZ that integrates visual properties for animals. Because the conjunctive neurons that integrate functional/associative properties reside in different modality CZs, knowledge of these properties remained intact.

Modality-specific multiple-category deficits. Whenever a lesion fails to fall squarely on a categorical region in a modality CZ, multiple categories may be affected (i.e., Finding #1). Under such conditions, the categories affected should pattern according to the shared properties on the damaged modality (but again not on other modalities). Patterns of the sort addressed by McRae and Cree (2002) should arise. One prediction is that the lost categories may typically share multiple properties. Because conjunctive neurons at this level capture *correlations* of properties, a lesion should compromise the ability to represent complex *conjunctions* of properties, such that the multiple categories affected have multiple properties in common. In contrast, a lesion to an analytic or holistic CZ should be more likely just to disrupt a single property type, such that the multiple categories affected tend to be similar on one (or at least fewer) properties.

The relative infrequency of artefact deficits. The variable dispersion principle explains the relatively low probability that a localised lesion compromises artefacts (Finding #3). Consider Figure 5b. To the extent that the conjunctive neurons for a category are distributed broadly throughout a modality CZ, it should be difficult for a localised lesion to disrupt them completely. Because within-category similarity is relatively low for artefacts, its conjunctive neurons should be highly dispersed, thereby making it less susceptible to damage. Conversely, because animals has much higher within-category similarity, its conjunctive neurons should be more localised, and therefore more susceptible to damage.

A related implication is that certain concepts may be more susceptible to lesions in some modality CZs than in others. If a category has a narrow dispersion on one modality CZ but a broad dispersion on another, a deficit should be more likely after damage to the narrow dispersion than to the broad one. It follows that deficits for tools should be particularly associated with lesions in fronto-parietal areas that mediate action, whereas deficits for animals should be particularly associated with lesions in temporal areas that support vision (Gainotti et al., 1995; Tranel, Damasio, & Damasio, 1997). If the actions associated with tools are more similar than their visual properties, tool deficits should be more likely following lesions to motor areas than to visual areas. Conversely, if the visual properties for animals are more similar than their other properties, animal deficits should be most likely to follow lesions to visual areas.

Deficits at high taxonomic levels. The SIT and variable dispersion principles offer an explanation for why conceptual deficits cover superordinate categories, not basic level ones (i.e., Finding #4). Because the members of superordinate categories share many properties in a family resemblance structure (Rosch & Mervis, 1975), the conjunctive neurons that code them should all be mixed together topographically within a modality CZ. Thus a lesion that damages the conjunctive neurons for one basic level category is likely to damage the conjunctive neurons for other basic level categories that share its family resemblance features. The result is the loss of a superordinate category, or at least much of it. The relative sparing of a few exemplars that often occurs may reflect the sparing of conjunctive neurons for properties that represent a subset of the category's family resemblance structure.

Simultaneous deficits in visual properties and nonliving things. Damage to a modality CZ also explains Finding #6, in which patients with a deficit for visual properties also have a deficit for nonliving things. Again such findings constitute a problem for the sensory-functional theory, because it predicts that loss of visual properties should affect animals more than nonliving things. According to CTT, however, the large region in a visual modality CZ that integrates the visual properties of nonliving things has been damaged, while the relatively circumscribed region that integrates the visual properties of animals has been spared. Because the categories are topographically localised, such lesions are possible. Paradoxically, in this case, the tight clustering of animal properties in the visual modality CZ spares it.

Damage to cross-modal CZs

As for modality CZs, CTT predicts that the conjunctive neurons in cross-modal CZs exhibit considerable topographic organisation by category. The difference is that the category knowledge linked in cross-modal CZs is multi-modal, whereas it is unimodal in modality CZs. Because of their similar organisation, however, many of the phenomena that arise in modality CZs also arise in their cross-modal counterparts.

Single-category deficits. When a lesion lands largely on the region of a cross-model CZ that integrates the multi-modal properties of a category, only that category may be lost (i.e., Finding #2). Unlike analogous lesions in modality CZs, such lesions in cross-modal CZs produce complete single-category deficits across modalities—the entire category is lost.

Forde, Francis, Riddoch, Rumiati, and Humphreys (1997) report a patient, SRB, with such a deficit. Following damage to an area stretching from his left medial temporal lobe to his occipital lobe, SRB exhibited impaired naming of living things, whether tested in the visual, gustatory, somatosensory, or auditory modality. SRB's loss of the category was apparently complete, consistent with a localised lesion in cross-modal CZs. Although SRB could have suffered individual damage on each of these modalities, it is more parsimonious to assume that he suffered damage to a single system that integrated information across them. The primary location of SRB's lesion—the inferior medial temporal lobe—is consistent with this conclusion. As described earlier, this area is widely implicated in cross-modal integration.

Findings #1, #3, and #4. Analogous to lesions in modality CZs, lesions in cross-modal CZs can produce a variety of important deficits noted in the literature. If a lesion does not fall squarely on a region that links a category but falls on a region that links multiple categories, a multiple category deficit should occur (i.e., Finding #1). Again the categories lost should tend to be similar on multiple properties, not just on one, because the conjunctive neurons in these CZs link conjunctions of properties. Furthermore, the similarities on which multiple deficits pattern should typically be multi-modal, not unimodal. Because conjunctive neurons in cross-modal CZs capture statistical regularities across modalities, conjunctive neurons lying together should link similar conjunctions of properties on different modalities.

Notably, categories that pattern together in a modality CZ may not necessarily pattern together in a cross-modal CZ. As described earlier, two categories may have similar properties on one modality but very different properties on another. Thus the same category may be lost with different other categories depending on whether the lesion falls in a modality or cross-modal CZ.

Because the variable dispersion principle applies to cross-modal CZs, categories having low within-category similarity at this level should be distributed more broadly than categories having high within-category similarity. As a result, a localised lesion should be less likely to completely disrupt all the conjunctive neurons for highly dispersed categories, such as artefacts (i.e., Finding #3).

Finally, deficits in cross-modal CZs should again primarily disrupt high-level categories, such as superordinates, rather than low-level categories, such as those at the basic level (i.e., Finding #4). Because the conjunctive neurons for different basic level categories are highly intermixed topographically, a localised deficit never disrupts just one—instead the larger superordinate is typically affected.

Equal loss of visual and functional properties for living things. Patients with a lesion that falls on conjunctive neurons for living things in a cross-modal CZ could exhibit equal loss of visual and functional properties (i.e., Finding #5). Again this pattern is difficult to reconcile with the sensory-functional theory, given its prediction that visual properties should suffer more than functional properties in a living things deficit. Nevertheless, this deficit is consistent with localised damage in cross-modal CZs. According to CTT, the conjunctive neurons for animals in these CZs conjoin properties across modalities. Two types of properties likely to be conjoined are those for visual recognition and function. Most importantly, a lesion to the conjunctive neurons that conjoin these properties would not only produce a deficit for animals, it would also produce roughly equivalent deficits in the retrieval of visual and functional properties for this category.

CONCLUSIONS

Rather than being mutually exclusive, competing theories of conceptual deficits complement each other in important ways, all contributing important insights. The sensory-functional theory stresses the modality-specific representation of knowledge. The domain-specific theory stresses the circumscribed representation of individual categories. The conceptual structure theory stresses the statistical relationships that hold for properties within and between categories.

CTT integrates these insights into a single theory. Feature maps, analytic CZs, holistic CZs, and modality CZs capture the importance of modality-specific representations. The SIT principle maps the statistical structure of property information into the topographical organisation of CZs. In modality CZs and cross-modal CZs, this organisation produces representations of individual categories that are relatively circumscribed. Lesioning this distributed system produces a variety of deficits that correspond to many reported in the literature, ranging from single category deficits to multiple category deficits organised around various patterns of similarity.

A concern might be that such a theory is too complex and powerful. The conceptual system, however, is far from a simple structure, and conceptual deficits exhibit tremendous variability (Coltheart et al., 1998). CTT's ability to produce a diverse array of deficits does justice to the complexity of the phenomena. Indeed, CTT predicts that conceptual deficits should take a variety of specific forms, which they do.

Perhaps CTT's most unique contribution is the SIT principle. Although topographic organisation is well established for feature maps, its role in CZs is much less clear. Nevertheless, modest independent evidence exists for this principle in the neuroimaging and electrophysiological literatures, as we saw earlier. Perhaps the strongest evidence, though, is the SIT principle's ability to explain the wide variety of conceptual deficits. By assuming topographical organisation, we were able explain these deficits within a single framework, organised around one principle. Clearly additional direct evidence for the SIT principle is needed. Future assessments of the relation between conceptual similarity and topography within CZs are critical to assessing CTT. We would be most grateful to hear from readers who know of findings that bear on this issue.

Surprisingly little direct evidence for A. R. Damasio's (1989) convergence zone hypothesis exists as well. To what extent do conjunctive neurons in association areas reactivate feature configurations in feature maps? Again further research is essential to answering this critical question. Clearly much evidence across many literatures indicates that high-level conceptualisations re-enact sensory-motor processing. Nevertheless, little if any direct evidence exists that this re-enactment is the result of conjunctive neurons reactivating assemblies of feature detectors. Even if our assumptions about re-enactment and the SIT principle turn out to be false, the research required to rule them out may lead to the discovery of important new mechanisms. Electrophysiological studies of the relations between associative and feature neurons may be particularly informative on this issue.

Conversely, CTT fits well with a variety of new findings in the literature that implicate the modality-specific representation of knowledge. Such evidence has not just been reported in the lesion and neuroimaging literatures on concepts, but also in the neuroimaging literatures on imagery (Farah, 2001; Grezes & Decety, 2001; Jeannerod, 1995; Kosslyn, 1994) and episodic memory (Nyberg, Habib, McIntosh, & Tulving, 2000; Wheeler, Petersen, & Buckner, 2000). The behavioural literatures similarly report increasing evidence that modality-specific representations underlie concepts (Barsalou, 1999, in press-b), memory (Glenberg, 1997), and language comprehension (Glenberg, 1997; Glenberg & Kaschak, 2002; Spivey et al., 2000; Stanfield & Zwaan, 2001; Zwaan et al., 2002). Thus there is reason to believe that CTT's basic assumptions are on the right track, although much further assessment is required.

Finally, we are intrigued by the potential importance of topography throughout the brain. The topographical organisation of feature maps is well known, but the topographical organisation of CZs is not. On reflection, it is not so surprising that evolution capitalised on such a simple principle to organise the brain at multiple levels. McClelland (personal communication, 2002) suggests that related representations in the brain need to influence each other frequently, and that being topographically close minimises the use of long-distance connections. For example, the facilitory relations that produce thematic priming benefit from close connections, as do the inhibitory relations that allow one category to suppress close competitors during object recognition. In general, understanding why topographical organisation exists at multiple brain levels may lead to insights about the brain's evolution and function.

REFERENCES

Adolphs, R., Damasio, H., Tranel, D., Cooper, G., & Damasio, A. R. (2000). A role for somatosensory cortices in the visual recognition of emotion as revealed by three-dimensional lesion mapping. *Journal of Neuroscience*, *20*, 2683–2690.

Anderson, A. K., & Phelps, E. A. (2000). Perceiving emotion: There's more than meets the eye. *Current Biology*, *10*, R551–R554.

Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition*, *11*, 211–227.

Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 629–654.

Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*, *Vol. 27* (pp. 1–64). San Diego, CA: Academic Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–660.

Barsalou, L. W. (in press-a). Abstraction as dynamic construal in perceptual symbol systems. In L. Gershkoff-Stowe & D. Rakison (Eds.), *Building object categories. Carnegie Symposium Series*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Barsalou, L. W. (in press-b). Situated simulation in the human conceptual system. *Language and Cognitive Processes*.

Barsalou, L. W., Sloman, S. A., & Chaigneau, S. E. (in press). The HIPE theory of function. In L. Carlson & E. van der Zee (Eds.), *Representing functional features for language and space: Insights from perception, categorization and development*. Oxford: Oxford University Press.

Bechtel, W., & McCauley, R. N. (1999). *Heuristic identity theory (or back to the future): The mind-body problem against the background of research strategies in cognitive neuroscience*. Paper presented at the Twenty-First Meeting of the Cognitive Science Society. Mahwah, New Jersey, USA.

Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press.

Breitmeyer, B. G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual-pattern masking, saccadic suppression, and information-processing. *Psychological Review*, *83*, 1–36.

Breitmeyer, B. G., & Ganz, L. (1977). Temporal studies with flashed gratings: Inferences about human transient and sustained channels. *Vision Research*, *17*, 861–865.

Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology*, *20*, 213–261.

Cappelletti, M., Waley-Cohen, H., Butterworth, B., & Kopelman, M. (2000). A selective loss of the ability to read and to write music. *Neurocase*, *6*, 321–331.

Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, *7*, 161–189.

Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate–inanimate distinction. *Journal of Cognitive Neuroscience*, *10*, 1–34.

Chen, X. (1995). Taxonomic changes and the particle-wave debate in early 19th-century Britain. *Studies in History and Philosophy of Science*, *26*, 251–271.

Cipolotti, L., Butterworth, B., & Denes, G. (1991). A specific deficit for numbers in a case of dense acalculia. *Brain*, *114*, 2619–2637.

Coltheart, M., Inglis, L., Cupples, L., Michie, P., Bates, A., & Budd, B. (1998). A semantic subsystem of visual attributes. *Neurocase*, *4*, 353–370.

Coslett, H. B., Brashear, H. R., & Heilman, K. M. (1984). Pure word deafness after bilateral primary auditory-cortex infarcts. *Neurology*, *34*, 347–352.

Cree, G. S., & McRae, K. (in press). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*.

Crosson, B., Moberg, P. J., Boone, J. R., Rothi, L. J., & Raymer, A. (1997). Category-specific naming deficit for medical terms after dominant thalamic/capsular hemorrhage. *Brain and Language*, *60*, 407–442.

Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, *33*, 25–62.

Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Grosset/Putnam.

Damasio, A. R., & Damasio, H. (1994). Cortical systems for retrieval of concrete knowledge: The convergence zone framework. In C. Koch & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain. Computational neuroscience* (pp. 61–74). Cambridge, MA: The MIT Press.

Damasio, H., Grabowski, T. J., Tranel, D., Hichwa, R. D., & Damasio, A. R. (1996). A neural basis for lexical retrieval. *Nature*, *380*, 499–505.

Davidson, R. J. (2000). Cognitive neuroscience needs affective neuroscience (and vice versa). *Brain and Cognition*, *42*, 89–92.

Davidson, R. J., Jackson, D. C., & Kalin, N. H. (2000). Emotion, plasticity, context, and regulation: Perspectives from affective neuroscience. *Psychological Bulletin*, *126*, 890–906.

De Valois, R., & De Valois, K. (1988). *Spatial vision*. New York: Oxford University Press.

De Renzi, E., & Spinnler, H. (1967). Impaired performance on color tasks in patients with hemispheric damage. *Cortex*, *3*, 194–217.

Dodd, J., & Kelly, J. P. (1991). Trigeminal system. In E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of neural science* (3rd ed., pp. 701–710). New York: Elsevier.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development.* Cambridge, MA: MIT Press.

Farah, M. J. (2000). The neural bases of mental imagery. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 965–974). Cambridge, MA: MIT Press.

Farah, M. J. (2001). Consciousness. In B. Rapp (Ed.), *The handbook of cognitive neuropsychology: What deficits reveal about the human mind* (pp. 159–182). Philadelphia, PA: Psychology Press.

Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, *120*, 339–357.

Forde, E. M. E., Francis, D., Riddoch, M. J., Rumiati, R. I., & Humphreys, G. W. (1997). On the links between visual knowledge and naming: A single case study of a patient with a category-specific impairment for living things. *Cognitive Neuropsychology*, *14*, 403–458.

Fujita, I., Tanaka, K., Ito, M., & Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature*, *360*, 343–346.

Funnell, E., & De Mornay Davies, P. (1996). JBR: A reassessment of concept familiarity and a category-specific disorder for living things. *Neurocase*, *2*, 461–474.

Gainotti, G., Silveri, M. C., Daniele, A., & Giustolisi, L. (1995). Neuroanatomical correlates of category-specific semantic disorders: A critical survey. *Memory*, *3*, 247–264.

Garcia, J., & Koelling, R. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, *4*, 123–124.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform "face area" increases with expertise in recognizing novel objects. *Nature Neuroscience*, *2*, 568–573.

Ghez, C. (1991). Voluntary movement. In E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of neural science* (3rd ed., pp. 609–625). New York: Elsevier.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences*, *20*, 1–18.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*, 558–565.

Gonnerman, L. M., Andersen, E. S., Devlin, J. T., Kempler, D., & Seidenberg, M. S. (1997). Double dissociation of semantic categories in Alzheimer's disease. *Brain and Language*, *57*, 254–279.

Goulet, S., & Murray, E. A. (2001). Neural substrates of crossmodal association memory in monkeys: The amygdala versus the anterior rhinal cortex. *Behavioral Neuroscience*, *115*, 271–284.

Grezes, J., & Decety, J. (2001). Functional anatomy of execution, mental simulation, observation, and verb generation of actions: A meta-analysis. *Human Brain Mapping*, *12*, 1–19.

Hart, J., & Gordon, B. (1992). Neural subsystems for object knowledge. *Nature*, *359*, 60–64.

Hasson, U., Hendler, T., Ben Bashat, D., & Malach, R. (2001). Vase or face? A neural correlate of shape-selective grouping processes in the human brain. *Journal of Cognitive Neuroscience*, *13*, 744–753.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 2425–2430.

Higuchi, S., & Miyashita, Y. (1996). Formation of mnemonic neuronal responses to visual paired associates in inferotemporal cortex is impaired by perirhinal and entorhinal lesions. *Proceedings of the National Academy of Sciences of the United States of America*, *93*, 739–743.

Hillis, A. E., & Caramazza, A. (1991). Category-specific naming and comprehension impairment: A double dissociation. *Brain*, *114*, 2081–2094.

Hochberg, J. (1998) Gestalt theory and its legacy: Organization in eye and brain, in attention and mental representation. In J. Hochberg (Ed.), *Perception and cognition at century's end: Handbook of perception and cognition* (2nd ed., pp. 253–306). San Diego, CA: Academic Press.

Holdstock, J. S., Gutnikov, S. A., Gaffan, D., & Mayes, A. R. (2000). Perceptual and mnemonic matching-to-sample in humans: Contributions of the hippocampus, perirhinal and other medial temporal lobe cortices. *Cortex*, *36*, 301–322.

Humphreys, G. W., & Forde, E. M. E. (2001). Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral and Brain Sciences*, *24*, 453–509.

Jeannerod, M. (1995). Mental-imagery in the motor context. *Neuropsychologia*, *33*, 1419–1432.

Johnson, K. E., & Mervis, C. B. (1997). Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General*, *126*, 248–277.

Kan, I. P., Barsalou, L. W., Solomon, K. O., Minor, J. K., & Thompson-Schill, S. L. (2003). Role of mental imagery in a property verification task: fMRI evidence for perceptual representations of conceptual knowledge. *Cognitive Neuropsychology*, *20*, 525–540.

Kandel, E. R. (1991). Perception of motion, depth, and form. In E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of neural science* (3rd ed., pp. 440–466). New York: Elsevier.

Kanwisher, N., Woods, R. P., Iacoboni, M., & Mazziotta, J. C. (1997). A locus in human extrastriate cortex for visual shape analysis. *Journal of Cognitive Neuroscience*, *9*, 133–142.

Kemler Nelson, D. G. (1989). The nature and occurrence of holistic processing. In B. E. Shepp & S. Ballesteros (Eds.), *Object perception: Structure and process* (pp. 357–386). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the Macaque cerebral cortex. *Journal of Neurophysiology*, *71*, 856–867.

Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: The MIT Press.

Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Chabris, C. F., Rauch, S. L., & Anderson, A. K. (1994). Identifying objects seen from different viewpoints: A PET investigation. *Brain*, *117*, 1055–1071.

Kosslyn, S. M., Thompson, W. L., & Alpert, N. M. (1997). Neural systems shared by visual imagery and visual perception: A positron emission tomography study. *NeuroImage*, *6*, 320–334.

Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science*, *293*, 1506–1509.

Kreiman, G., Koch, C., & Fried, I. (2000a). Category-specific visual responses of single neurons in the human medial temporal lobe. *Nature Neuroscience*, *3*, 946–953.

Kreiman, G., Koch, C., & Fried, I. (2000b). Imagery neurons in the human brain. *Nature*, *408*, 357–361.

Laiacona, M., Barbarotto, R., & Capitani, E. (1993). Perceptual and associative knowledge in category-specific impairment of semantic memory: A study of two cases. *Cortex*, *29*, 727–740.

Lambon Ralph, M., Howard, D., Nightingale, G., & Ellis, A. (1998). Are living and non-living category-specific deficits causally linked to impaired perceptual or associative knowledge? Evidence from a category-specific double dissociation. *Neurocase*, *4*, 311–338.

Lambon Ralph, M. A., McClelland, J. L., Patterson, K., Galton, C. J., & Hodges, J. R. (2001). No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a connectionist model. *Journal of Cognitive Neuroscience*, *13*, 341–356.

Lezak, M. (1995). *Neuropsychological assessment* (4th ed.). Oxford: Oxford University Press.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.

Lorenzi, C., Wable, J., Moroni, C., Derobert, C., Frachet, B., & Belin, C. (2000). Auditory temporal envelope processing in a patient with left- hemisphere damage. *Neurocase*, *6*, 231–244.

Martin, A. (2001). Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone (Eds.), *Handbook of functional neuroimaging of cognition* (pp. 153–186). Cambridge, MA: MIT Press.

Martin, A., & Chao, L. (2001). Semantic memory and the brain: structure and process. *Current Opinion in Neurobiology*, *11*, 194–201.

Martin, A., Ungerleider, L. G., & Haxby, J. V. (2000). Category-specificity and the brain: The sensory-motor model of semantic representations of objects. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 1023–1036). Cambridge, MA: MIT Press.

McCauley, R. N., & Bechtel, W. (2001). Explanatory pluralism and heuristic identity theory. *Theory and Psychology*, *11*, 736–760.

Macrae, C. N., & Lewis, H. L. (2002). Do I know you? Processing orientation and face recognition. *Psychological Science*, *13*, 194–196.

McRae, K., & Cree, G. S. (2002). Factors underlying category-specific semantic deficits. In E. M. E. Forde & G. W. Humphreys (Eds.), *Category-specificity in mind and brain* (pp. 211–249). Hove, UK: Psychology Press.

Mishkin, M., Vargha-Khadem, F., & Gadian, D. G. (1998). Amnesia and the organization of the hippocampal system. *Hippocampus*, *8*, 212–216.

Morrison, D. J., & Schyns, P. G. (2001). Usage of spatial scales for the categorization of faces, objects, and scenes. *Psychonomic Bulletin and Review*, *8*, 454–469.

Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.

Murray, E. A., & Bussey, T. J. (1999). Perceptual-mnemonic functions of the perirhinal cortex. *Trends in Cognitive Sciences*, *3*, 142–151.

Murray, E. A., Gaffan, D., & Mishkin, M. (1993). Neural substrates of visual stimulus association in Rhesus monkeys. *Journal of Neuroscience*, *13*, 4549–4561.

Murray, E. A., & Richmond, B. J. (2001). Role of perirhinal cortex in object perception, memory, and associations. *Current Opinion in Neurobiology*, *11*, 188–193.

Navon, D. (1977). Forest before trees: The precedence of global features in visual processing. *Cognitive Psychology*, *9*, 353–383.

Nyberg, L., Habib, R., McIntosh, A. R., & Tulving, E. (2000). Reactivation of encoding-related brain activity during memory retrieval. *Proceedings of the National Academy of Sciences of the United States of America*, *97*, 11120–11124.

Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.

Parker, A., & Gaffan, D. (1998). Lesions of the primate rhinal cortex cause deficits in flavour: Visual associative memory. *Behavioural Brain Research*, *93*, 99–105.

Rackover, S. S. (2002). Featural vs. configurational information in faces: A conceptual and empirical analysis. *British Journal of Psychology*, *93*, 1–30.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.

Samson, D., Pillon, A., & De Wilde, V. (1998). Impaired knowledge of visual and non-visual attributes in a patient with a semantic impairment for living entities: A case of a true category-specific deficit. *Neurocase*, *4*, 273–290.

Sartori, G., & Job, R. (1988). The oyster with four legs: A neuropsychological study on the interaction of visual and semantic information. *Cognitive Neuropsychology*, *5*, 105–132.

Schreiner, C. E., Read, H. L., & Sutter, M. L. (2000). Modular organisation of frequency integration in primary auditory cortex. *Annual Review of Neuroscience*, *23*, 501–529.

Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1–54.

Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 31, pp. 305–349). San Diego, CA: Academic Press.

Shelton, P. A., Bowers, D., Duara, R., & Heilman, K. M. (1994). Apperceptive visual agnosia: A case study. *Brain and Cognition*, *25*, 1–23.

Silveri, M. C., Gainotti, G., Perani, D., Cappelletti, J. Y., Carbone, G., & Fazio, F. (1997). Naming deficit for non-living items: Neuropsychological and PET study. *Neuropsychologia*, *35*, 359–367.

Solomon, K. O., & Barsalou, L. W. (2001). Representing properties locally. *Cognitive Psychology*, *43*, 129–169.

Solomon, K. O., & Barsalou, L. W. (2003). *Perceptual simulation in property verification*. Manuscript under review.

Spivey, M., Tyler, M., Richardson, D., & Young, E. (2000). Eye movements during comprehension of spoken scene descriptions. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 487–492). Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Stanfield, R. A., & Zwaan, R. A. (2001). The effect of implied orientation derived from verbal context on picture recognition. *Psychological Science*, *12*, 153–156.

Tanaka, K. (1996). Representation of visual features of objects in the inferotemporal cortex. *Neural Networks*, *9*, 1459–1475.

Tanaka, K. (1997a). Columnar organization in the inferotemporal cortex. In K. Rockland, J. Kaas, & A. Peters (Eds.), *Cerebral cortex* (pp. 469–498). New York: Plenum Press.

Tanaka, K. (1997b). Inferotemporal cortex and object recognition. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-network models of cognition: Biobehavioral foundations* (Vol. 121, pp. 160–188). Amsterdam, The Netherlands: North-Holland/Elsevier Science Publishers.

Tanaka, K. (1997c). Mechanisms of visual object recognition: Monkey and human studies. *Current Opinion in Neurobiology*, *7*, 523–429.

Thomas, E., Van Hulle, M. M., & Vogels, R. (2001). Encoding of categories by noncategory-specific neurons in the inferior temporal cortex. *Journal of Cognitive Neuroscience*, *13*, 190–200.

Thompson, W. L., & Kosslyn, S. M. (2000) Neural systems activated during visual mental imagery. In A. W. Toga & J. Mazziotta (Eds.), *Brain mapping: The systems* (pp. 535–560). San Diego, CA: Academic Press.

Tranel, D., Adolphs, R., Damasio, H., & Damasio, A. R. (2001). A neural basis for the retrieval of words for actions. *Cognitive Neuropsychology*, *18*, 655–674.

Tranel, D., Damasio, H., & Damasio, A. R. (1997). A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia*, *35*, 1319–1327.

Tranel, D., Damasio, H., & Damasio, A. R. (1998). The neural basis of lexical retrieval. In R. W. Parks & D. S. Levine (Eds.), *Fundamentals of neural network modeling: Neuropsychology and cognitive neuroscience* (pp. 271–296). Cambridge, MA: MIT Press.

Treisman, A. M., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, *12*, 97–136.

Tyler, L. K., & Moss, H. E. (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences*, *5*, 244–252.

Tyler, L. K., Moss, H. E., Durrant-Peatfield, M. R., & Levy, J. P. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. *Brain and Language*, *75*, 195–231.

Wang, G., Tanaka, K., & Tanifuji, M. (1996). Optical imaging of functional organization in the monkey inferotemporal cortex. *Science*, *272*, 1665–1668.

Warrington, E. K., & McCarthy, R. A. (1987). Categories of knowledge: Further fractionations and an attempted integration. *Brain*, *110*, 1273–1296.

Warrington, E. K., & McCarthy, R. A. (1994). Multiple meaning systems in the brain: A case for visual semantics. *Neuropsychologia*, *32*, 1465–1473.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829–854.

Wheeler, M. E., Petersen, S. E., & Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory-specific cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *97*, 11125–11129.

Wolfe, J. (2000). Visual attention. In K. K. De Valois (Ed.), *Seeing* (2nd ed., pp. 335–386). San Diego, CA: Academic Press.

Wu, L., & Barsalou, L. W. (2003). *Perceptual simulation in property generation*. Manuscript submitted for publication.

Zeki, S. (1993). *A vision of the brain*. Cambridge, MA: Blackwell Scientific Publications.

Zhou, H., Friedman, H. S., & Von der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. *Journal of Neuroscience*, *20*, 6594–6611.

Zwaan, R. A., Stanfield, R. A., & Yaxley, R. H. (2002). Do language comprehenders routinely represent the shapes of objects? *Psychological Science*, *13*, 168–171.