UNIVERSITY OF SOUTHAMPTON

The Contribution of Meaning in Forming Holistic and Segmented Based Visual Representations

Wendy Smith

Submitted for the degree of Doctor of Philosophy

Department of Psychology
University of Southampton

Submitted August, 2000

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL SCIENCE

PSYCHOLOGY

Doctor of Philosophy

THE CONTRIBUTION OF MEANING IN FORMING HOLISTIC AND SEGMENTED BASED VISUAL REPRESENTATIONS

By Wendy Smith

The visual cognitive system uses representations as a means of conveying information from one process to another. Understanding the representation and its format is important, therefore, in understanding visual cognition. One factor that could influence the representational format is the meaning associated with an item that is being represented in the cognitive system. The research question under investigation here is, specifically, whether the meaning associated with a stimulus encourages holistic representations or segmented representations.

Assuming that the underlying representation will constrain the processes performed on it, three paradigms were used to compare performance across meaningful and meaningless stimuli. Within each type of stimulus the relative performances with simple and complex objects were examined. The complexity was used as an analytical tool. In a segment-based representation, with relatively independent parts, there would be more parts in the representation of a complex stimulus that in that of a simple stimulus. In holistic representations the parts would be less independent; the number of parts would be less influential than the relationship among them.

The first study used a mental rotation task. A segmented representation would show an interactive effect of orientation and complexity, whereas a holistic representation would not. The findings suggested that the meaningful stimuli were rotated part by part, whereas the meaningless objects were rotated holistically. The second study used a part search task. It was predicted that the dependence among the parts in a holistic representation would result in a greater difference in performance across complexities than in a segmented representation with independent parts. The complexity of the stimulus showed a greater difference in the meaningful stimuli than in the meaningless stimuli. Assuming that holistic representations made more use of configural information than segment-based representations, the final study tested the contribution of configural information to the representation during a binary forced-choice probe task. The meaningful stimuli showed a greater advantage when configural information was present, especially the complex stimuli, relative to the meaningless stimuli.

The findings suggested that meaningful objects were represented more holistically than comparable meaningless objects; this difference is greater in the complex stimuli.

LIST OF CONTENTS

CHAPTER 1: OBJECT REPRESENTATIONS

CHAPTER 2: STIMULUS DESIGN AND ASSESSMENT

CHAPTER 3: THE EFFECT OF MEANING ON PERFORMANCE IN A ROTATION TASK

CHAPTER 4: THE EFFECT OF MEANING ON PERFORMANCE IN A PART SEARCH TASK

CHAPTER 5: THE EFFECT OF MEANING ON PERFORMANCE IN WHOLE AND PART PROBE TASKS

CHAPTER 6. OBJECT REPRESENTATION AND THE EFFECT OF MEANING

LIST OF TABLES AND FIGURES

CHAPTER 2: STIMULUS DESIGN AND ASSESSMENT

CHAPTER 3: THE EFFECT OF MEANING ON PERFORMANCE IN A ROTATION TASK

CHAPTER 4: THE EFFECT OF MEANING ON PERFORMANCE IN A PART SEARCH TASK

CHAPTER 5: THE EFFECT OF MEANING ON PERFORMANCE IN WHOLE AND PART PROBE TASKS

CHAPTER 6.OBJECT REPRESENTATION AND THE EFFECT OF MEANING

APPENDIX

ACKNOWLEDGEMENTS

First, I would like to thank Dr Itiel E. Dror for his contribution in supervising the research. The following members of the Department of Psychology at Southampton University offered useful advice and comments: Dr Nick Donnelly, Dr Richard Hastings, Professor Bob Remington, Dr Tony Roberts, and Dr Mark Van Vugt. The members of my Departmental Research Group provided valuable support, especially my fellow postgraduate students: Ms Wendy Kneller, Ms Liz Lee, Ms Vered Rafaely, and Ms Julie Winstone. I would also like to show appreciation for the unsung heroes of my research, the undergraduates who took part in my experiments, and provided the data.

Specific thanks are due to Ms Vered Rafaely and Ms Christine Weston for their assistance with the stimulus assessment reported in Chapter 2; and to Professor John Wixted, Professor Michael Corballis, Professor Apostolos Georgopoulos and an anonymous reviewer for their comments on a journal article based on Chapter Three.

Finally, I would like to make a special acknowledgement to Professor Stevan Harnad, for providing inspiration and support at the beginning of the research, and for being willing to take the first step of a long journey with me.

CHAPTER 1

OBJECT REPRESENTATIONS

ABSTRACT

Information processing theory posits that information entering the cognitive system via the senses is transformed to a state where a response can be decided through various processes acting upon mental representations. To understand cognition, therefore, we need to understand the nature of the processes and the nature of the representations. One such process (or, more likely, group of processes) is object recognition. The nature of the representations formed during object recognition is still underspecified. Different representational formats could be used; some representations appear more holistic, or unified, than do others. Several factors have been suggested as playing a role in deciding the format of mental representations. These include task demands, expertise, familiarity and stimulus properties. However, although considerable research has been carried out, the list is still far from exhaustive. This thesis examines one potential component in determining representations; that of meaning. The aim of this thesis is to address the question: Does the meaning of an object affect the representational format of that object, and, if so, how? Past research indicates that meaning may act to unify the representation, thus making it more holistic, although this remains an open issue. Furthermore, the term "holistic" has been used to describe three diverse conditions. The first use of the term is a single, undifferentiated unit; the second is a representation with the emphasis on configural processing; and the third is a number of parts processed in parallel. In common with other cognitive "black box" issues, this question is not one that can be examined directly. The research presented here takes three different approaches, with three different paradigms, with the purpose of assessing whether meaning encourages a more segment-based or holistic representation, and, if the representation is holistic, which of the three terms mentioned above best describe it.

CHAPTER 1

INTRODUCTION

Object representations provide a mental portrayal of the objects in the physical environment. This can serve at least two purposes. First, a mental representation can be used in the mental rehearsal, or planning, of potentially difficult, time consuming, or dangerous manipulations before the physical manipulation takes place. Second, a mental representation can be used as part of the input to the cognitive computational system. Information enters into human cognitive processes via the sensory mechanisms. The raw sensory information is converted into a format that is usable by the cognitive system. This information is processed. Mental representations are the means by which information can be transformed from one form into another, through these processes, until such time as a response can be made. Early work in the area of mental representations focused less on the second purpose, and concentrated more on the first. The resulting debate was based around whether representations as such actually existed or not. Mental representations and images were conceptualised as subjective experiences, leading to three main strictures. Representations were criticised for being homuncular, epiphenomenal, and the study of mental representations was criticised for the subjective nature of the topic. However, the understanding of mental representations moved away from the idea of a purely subjective experience, to the idea that such representations are the working tools of the cognitive system. As representations became envisaged more as information carriers during the computations of cognition, these criticisms became less relevant (see Markman & Dietrich, 2000).

Kosslyn (1994) provided a model of visual cognition in which mental imagery was a necessary part of visual perception. Several studies have demonstrated an equivalence between imagery abilities and perceptual abilities (e.g. Tlauka & McKenna, 1998). Imagery can interfere with perception (known as the Perky effect; Craver-Lemley & Reeves, 1992), and can also facilitate it (Ishai & Sagi, 1997). Imagery can induce perceptual priming

(McDermott & Roediger, 1994), visual illusions (Wallace, 1984; although see also Reisberg & Morris, 1985), and result in a memory for an imagined occurrence being interpreted as a real event (Roediger, 1996; Goff & Roediger, 1998). Imagery and perception appear to take place in the same structures within the brain; brain damage has been reported to cause similar deficits in image and perception (see Farah, 1988; Policardi, Perani, Zago, Grassi, & Làdavas, 1996).

However, representations are not just a form of perception. Differences have also been documented. For example, images appear to contain more than just an abstraction of the percept (Rouw, Kosslyn & Hamel, 1997). Preserved imagery with impaired perception has been described (e.g. Chatterjee & Southwood, 1995), although this may be due to the specialisation of the processes rather than the separation of the processes (Behrman, Moscovitch & Winocur, 1994). Recent research has shown a double dissociation between perception and mental imagery (Faw, 1997). Evidence has also supported a further dissociation between visual representation and spatial representation (e.g. Farah, Hammond, Levine, & Calvanio, 1988). The neuroscientific evidence reflects that the image shares characteristics with the percept, although the two do not overlap completely.

If representations are the "information-carriers" of the cognitive system then the type of information that is carried, and the way in which it is organised, could result in differences in the format of the representation. In particular, the type of information may influence whether that information is carried as a whole unit, or whether it is carried in segments. The work in this thesis assesses the role of meaning in mediating the type of representation that is found. One important process, or group of processes, making use of mental representations of meaningful objects is in the area of object recognition. This domain will be used as the framework for the research carried out. In the next section, we consider different representations within this framework.

CHAPTER 1

OBJECT RECOGNITION

Overview

Object recognition allows sense to be made of the environment. It is a factor in linking present experience of the environment with past experiences, allowing a response to be made based on the knowledge gained through prior encounters as well as the immediate one. Several theories of object recognition exist; a summary of the common points of those theories is presented here.

Object recognition is a combination of two types of information. One type of information relates to the perceptual information that is extracted from the shape of the to-be-recognised object. The other is the representation of the structure of a potential candidate object, generated from memory. The object recognition process has to compare these two images for a match. For this mechanism to be successful, the format of the perceptual representation has to be compatible (or at least comparable) with the memory generated representation.

The comparison may reveal a match; if not, one of the representations may have to be transformed (e.g. rotated) so that it is compatible with the other representation. If a match is still not achieved, then the next step is to find another potential candidate object. A directed search to seek more information from the new object may also be needed (Kosslyn, 1994). These two strategies, or indeed any others, can be repeated until a satisfactory match is found. When a match has been found (or, in the absence of a perfect match, agreement reached on sufficient criteria), then other information associated with the object (e.g. knowledge about the item) will become available, through semantic memory, and can be linked to the newly recognised object. The information may have originally come from several different sources, through several different modalities. Once this point has been reached, a name is usually generated, though not always immediately. If the process has been successful, the environmental object now has a name and meaning attached to it. Other processes in the cognitive system can use the information to decide an appropriate response. Even if the recognition system was unsuccessful in generating an acceptable match, sufficient processing could have occurred to inform the decision, or a complete lack of information could also guide a suitable response.

The relationship between the name and the meaning, in terms of the semantic information associated with the object, is still equivocal. In some models a name has to be generated before semantic information can be accessed (e.g. Brennan, David, Fluchaine, Pellat, 1996); in others it is not necessary (e.g. Riddoch & Humphreys, 1987); and in yet others, the semantic information aids the comparison process (e.g. Boucart & Humphreys, 1997). In face recognition, although people can often fail to access a name while retaining a great deal of knowledge of the person, it is very rare to remember a name with no other knowledge at all. This finding has been questioned (e.g. Brennan, David, Fluchaine, Pellat, 1996), but mainly with case studies of individuals. Hodges and Green (1998) tested patients with Alzeimer's disease across 1,200 trials, and found no clear cases where naming occurred with a total absence of semantic information. However, some of the semantic information generated by the patients was meagre (e.g. the correct name, Ronald Reagan, generated the semantic information "a politician", p129). Again, the emphasis in these observations was on face recognition, rather than object recognition in general. Words and names do not appear to share the same manner of processing (Valentine, Hollis, & Moore, 1998). Burton and Bruce (1992) explain that names are unique to the individual, in a way that semantic knowledge is unlikely to be. The name of an object is not likely to be unique to the object in the usual level of use, for example, "cup", and could even be considered part of the semantic information.

To summarise, object recognition encompasses both perceptual and memory representations, and has to render them comparable to be able to generate a match or mismatch decision. One constraint operating on the representations used in object recognition is that the perceptual representation is likely to be of a specific two-dimensional view of a three-dimensional object. This has to match the stored representation, and previous representations may not be of the same viewpoint. Different viewpoints of an object can present considerably different shapes. One problem for theories of object recognition is how these two representations can be reconciled to allow some degree of object constancy and successful matching of representations.

Representations and Object Constancy

There are two basic solutions to the problem of shape constancy in object recognition. First, the stored representation can be of properties that do not change among the viewpoints (viewpoint independent, or VI theories). Second, the representation can be of properties that do change with viewpoint, and some allowance for this is made during the recognition process (viewpoint specific or VS theories).

Viewpoint-Independent Theories

Viewpoint independent (VI) theories consider that object recognition occurs through a process involving a representation containing information that is non-variant over multiple viewpoints. The non-variant information is based in some sort of “primitive”; that is, basic components from which all shapes can be derived. There are a finite number of such primitives; Biederman (1987), for example, suggested that there were 36.

Marr and Nishihara (1978) used cylinders, or generalised cones, as the basic component. One of the best-known theories, however, is probably that of Biederman (1987). Biederman called the primitive shapes "geons" and claimed each geon was defined by a set of features, which were a function of the main axis and a cross section of the shape. The features consisted of "non-accidental properties"; that is, the properties were invariant no matter what view was taken. Non-variant information may use, for example, parallel lines. If parallel lines are present in the representation, then there is a very high probability that they are present in the object, because the chance of them arising as an accident of the viewpoint is very low. So, a cube could be defined by the following properties: three parallel lines, one internal Y shape, and three external arrow shapes. It can be distinguished from a brick shape by the length of the parallel lines. These properties remain invariant for the majority of views of the object.

Geon theory was tested in an experiment in which participants had to name an object from a drawing (Biederman, 1987). There were three conditions. In the first, the drawing was complete. In the other two, only a proportion of the drawing was present. In one case, the intersections (containing the invariant features) remained, and in the other condition, the intersections were removed. Performance was better in the condition in which the invariant properties remained. However, although this supports the existence of non-accidental properties being useful in the recognition of objects, it does not provide evidence for the necessity of geons themselves (see also Cave & Kosslyn, 1993). Identification of the geons may be a function of the regularity of the objects, rather than invariant features (Leewenberg, van der Helm, & van Lier, 1994).

One criticism of the earlier version of Biederman's theory, recognition by components (RBC), was that it was not specific enough to allow recognition of similar objects. The main problem was that if only the geons were specified it would be impossible to tell apart two objects with similar geons in different positions; for example, a bucket and a cup (assuming similar sizes and shapes of geons). This difficulty, predicted by the RBC theory, is not reflected in human performance. This was resolved through the addition of a component that specified the spatial arrangement of the geons. The new representation was named a geon structural description (GSD; Hummel & Biederman, 1992).

The crucial test for VI approaches is whether the object can, indeed, be identified from any viewing position. Many studies have demonstrated that when an object is presented in a non-canonical, or change of, position, it can affect recognition performance. For example, presenting an object that has been rotated in the picture plane increases the time it takes to correctly name the object (e.g. Murray, 1995). Similar results are obtained when pictures and their mirror images have to be matched (e.g. Joliceour, Corballis, & Lawson, 1998). The response time in these tasks approximates a linear relationship between time and angle of orientation, from 0° up to 120°. This suggests that, at least between these angles, a mental rotation process is aligning one of the representations with the other.

The situation between 120° and 180° is not so clear cut (see Chapter Three for further consideration). The response times flatten out over these angles, suggesting that rotation is not taking place. Several explanations have been offered. One explanation is that a "flipping" strategy is used, whereby the representation is rotated in the depth plane, and this is quicker than rotating in the picture plane (e.g. Murray, 1997). An alternative is that VI processes are used at greater angles, where the task is more difficult (Lawson & Joliceour, 1998). In addition, the orientation effects can lessen with repeated exposure, although this reduction will not transfer to other objects (e.g. Joliceour, 1985). This would be predicted if a GSD is being formed over repeated exposures (Murray, 1995). For the purposes of the argument presented here, rotation would not be necessary at any stage in identification if the representation was viewpoint invariant. The findings of orientation effects are supported by the results from priming studies. Priming does not transfer well from one view to another (Lawson & Humphreys, 1998).

These findings weaken the position for VI processes. Biederman and Gerhardstein (1993) explain the findings by suggesting that performance will depend on whether the GSD is preserved within the rotation, or other transformation. They suggest that for the GSD to be a viable tool, three conditions need to be satisfied, and these will determine the extent to which there can be generalisation from one view to a novel view. If the three conditions do not hold, then one can expect to see viewpoint dependency. The conditions are as follows. First, the objects must be capable of being reduced to parts. Second, different objects must have different GSDs for them to be classified as different. Third, different views of the same object must have the same GSDs for them to be classified as the same.

Although these extra findings do provide an account for the previous findings, they also change the position of VI processes considerably. For example, if the GSD can only be seen from certain orientations, it suggests that it is not viewpoint invariant at all. It also implies that several encounters with the object may be necessary for the GSD to be established. If the object is seen from the "wrong" position, then a useful GSD will not be produced. It calls into question what information would need to be incorporated into the GSD to give adequate performance. For example, how well are the positions of the geons specified in the GSD? If the positions are coarse (e.g. above, below), then objects from the same category may be very difficult to tell apart using this approach. If the positions are precisely defined, then the GSD loses its invariance across views. For accurate recognition, the information in the representation has to be unique for any given object; if not, errors will occur. Given the low number of errors in normal recognition, a high degree of distinctiveness is indicated. The theory has to account for how this distinctiveness in achieved, while at the same time maintaining accurate constancy over several viewpoints with the same representation.

Viewpoint-Specific Theories

Viewpoint specific (VS) theories consider that object recognition occurs through a process involving several representations that differ over multiple viewpoints. One version of viewpoint-specific theories posits that many representations are stored - one for every viewpoint, or an aggregate of every encounter with an object. The ultimate VS of this type is an (infinite) number of templates for every object encountered - totally view-specific representations. This would seem unnecessary; for example, humans can recognise objects without having to examine them from every possible angle. There may be problems, too, in linking together very different views of the same object. The theory assumes that the viewpoints are independent of each other, and were linked through either semantic knowledge or the name of the object (see Minsky, 1975). This does not appear to be the case, however, because, following brain injury, patients show a double dissociation between the ability to recognise objects and the ability to name or relate any knowledge about them (Warrington & Taylor, 1978). At some point, all the views must be either combined or reduced (Edelman, 1995). The exact information that would be left by the combination or reduction, however, is not clear.

One possibility involves a descriptive representation, based on common properties derived from a combination of views (Ullman, 1998). Although these are not invariant properties, and Ullman considers the representations to be viewpoint dependent (Ullman, 1989), this still has striking similarities to the later GSD theories (Hummel & Biederman, 1992). The descriptive representation can be compared with one or more stored representations. Edelman (1995) suggests a prototype is produced from the different views, and stored for comparison. An alternative to a prototype is the exemplar-based approach, in which a selection of views is represented in memory, with both shape and orientation stored (Gauthier & Tarr, 1997).

Novel viewpoints would be unrecognisable unless a familiar viewpoint could also be accessed; performance does indeed appear to be affected by novelty (see Tarr & Bülthoff, 1995), although to a lesser extent than it is affected by orientation. Solutions involving viewpoint specific mechanisms have been presented. For example, combining multiple views can lead to the recognition of a novel view, especially if the object is symmetrical (Srinivas &

Schwoebel, 1998). However, this is moving away from an approach in which there is a template for every viewpoint. It is also possible that novel representations could be extrapolated from those gained through experience. This implies a move to more VI processes, as not all the information will be extrapolated easily.

A second version of VS theories posits that, although more than one representation may be stored, generally there is a canonical viewpoint. The canonical viewpoint is the "preferred" viewpoint when observers are asked to choose, and displays certain advantages in processing performance compared with non-canonical viewpoints. This viewpoint is often, but not always, a three-quarter view of the object (Blanz, Tarr & Bülthoff, 1999). When a non-canonical percept is compared to the canonical stored representation, one of the representations is transformed to match the other (Tarr & Pinker, 1989). This can explain the reduced effects of orientation with repeated exposure. As more representations are added to the set, the need for transformations is reduced (Tarr & Pinker, 1989). Unfortunately for this explanation, although it provides a neat account for the findings, it is not supported empirically (Murray, 1995).

It is still not clear, however, why sometimes transformations are needed, and sometimes they are not. This may relate to the information needed to complete the task (e.g. Corballis & Corballis, 1988; Takano, 1989). Different tasks will need different information for successful completion, and Takano (1989) suggested that the nature of this information determined whether mental rotation was necessary. He described two types of information, each with two variants that, when combined together, result in four types of information in total. Information can be orientation-free or orientation-bound, and each of these can operate at an elementary level, or a conjunctional level. Orientation-free information refers to the distinction between, for example, a straight line and a curved line. This information can be at this level, or it can refer to how two such pieces are related to each other; for example, attached or unattached. Orientation-bound information holds orientation information, such as horizontal line versus vertical line. Again, this can also be presented at another level, which specifies how two pieces of information relate, but this time in terms of "above" or "below" and so forth.

In summary, some VS theories posit numerous representations from many different viewpoints; this appears unnecessary, and does not fit in with the empirical evidence. Other VS theories suggest a canonical viewpoint, and other viewpoints being transformed to match. However, transformation is not always needed; this might depend on the task demands, and the information needed from the representation to complete the task.

Evaluation

To summarise, the representation in VI approaches consists of the parts of an object and how these parts are arranged. The parts are defined by the non-accidental properties that are invariant over most viewpoints. Incoming and stored representations share the same representation if they are from the same object. VS approaches posit the collection of information from several different viewpoints per object, which are combined together. The incoming and stored representations may be different, in terms of having been derived from different viewpoints. If this is the case, one representation can be transformed to match the orientation or position of the other representation.

One problem with the VI theories is that, although objects can be recognised before they are transformed to canon (Hamm & McMullen, 1998), there are advantages for the canonical viewpoint, both from the point of easier recognition and better priming for the object (Lawson & Humphreys, 1998). This can still be accommodated within the viewpoint independent theories by the canonical viewpoints revealing the non-accidental properties more clearly than other viewpoints, and VI developing over repeated exposures to the object. This suggests there may be limits to when VI approaches can be used. Furthermore, empirical evidence shows canonical superiority even when the same features are shown (Liu, 1996); this finding generalises to situations when real objects, with geon-like properties, are used (Humphreys & Khan, 1992).

A more pressing problem is when the non-accidental properties are derived. If they are readily available from the retinal image at the first encounter, no matter what the viewpoint, then this does not explain the advantage of canonical viewpoints. If they are derived from several encounters with the object (Murray, 1995), then this suggests that viewpoint specific processing may be needed until or unless invariant features can be extracted. In either case there is a role for VS representations in addition to VI representations. It also casts doubt on how invariant the features actually are if several exposures are needed before they can be incorporated into the representation. Another related question concerns the variant information sampled in the retinal array. If only VI information is processed, then what becomes of the other information? If it is not incorporated in some sort of representation, then it will not be available should the VI information prove inadequate for recognition.

The mechanism involved in defining the geons limits the number of objects which can be represented and recognised (Kurbat, 1994). The objects that cannot be represented and recognised using VI processes apparently include geons themselves (Tarr, Williams, Hayward & Gauthier, 1998). Following from this, another issue relates to the manner in which the invariant features are integrated into geons. Just as different geons can be arranged in different configurations to form different objects, so, presumably, can invariant features be arranged in different configurations to form different geons. This implies a process of successive elaboration of representation, or a hierarchy, beginning with very basic features, and building up to an object, or even a scene.

In light of several clear restrictions on VI processes, they were muted to view-restricted rather than invariant, and limitations were put on their use.

Geons and their configurations have to be visible, and several GSDs may represent a single object (Hummel & Biederman, 1992; see also Ullman, 1998 for a similar account from a different perspective).

These criticisms suggest that VI approaches cannot account for object recognition on their own. If this is so, does this mean there is no role for them to play at all, or are they used along with VS representations? In other words, can VS approaches account for object recognition without recourse to VI representations? Tarr and Bülthoff (1995) suggest that object recognition is based on multiple views of objects, with sets of views of an object forming a complete representation. Percepts formed outside this set will be transformed to one of the views. Once the percept has been recognised as an example of the object, the new view can be added to the set. This process may give the appearance of VI if the stored representations are complete and reduce the need for transformations. In this way, VS representations can account for the findings of both VI and VS representations. However, VS approaches have also been criticised.

One problem with VS theories is the heavy storage and processing requirements, although this does not preclude them if this is the best method of achieving successful recognition. If VI representations can achieve the same purpose, however, then it could restrict excess resource use.

The canonical viewpoint that is found in many studies proposing a viewpoint specific explanation also needs more specification. The preferred view for an object tends to be consistent over participants (e.g. see Perret, Harries, & Looker, 1992). This is not always the case, however. In contrast to Perret *et al*, who used wooden models, Cutzu and Edelman (1994) found that canonical views of paperclips were not consistent.

Edelman and Bülthoff (1992) considered that a property in the shape of the object was encouraging a canonical view. Blanz, Tarr and Bülthoff (1999) tried to assess what property this was. Viewer-oriented properties included several contextual effects. They found that task affected the preferred view. When participants chose a view for a photograph, then physical properties such as stability came into play, along with informational properties such as seeing as much of the object as possible. Familiarity affected objects with a clear context; for example, views of an aircraft from below were acceptable, but views of cars from below were not. Functionality was found to be important where the view affected how the object would be easily used. For example, right-handed and left-handed people imaged a teapot with the handle on the appropriate side for use. Object-oriented properties were less well specified. One possible explanation for canonical views is that it allows a GSD to be created. However, Blanz, Tarr and Bülthoff dismiss this because the factors determining the canonical viewpoint tend towards viewer-oriented properties rather than object-oriented properties.

One of the major criticisms of VS accounts is the lack of specificity for the format of the representation, compared with viewpoint independent theories (Biederman & Gerhardstein, 1995). In fact, Biederman and Gerhardstein question whether VS representations as such exist, or whether these approaches are better considered in terms of VS processes, or even methodological artifacts. The latter point is not generally accepted (see Tarr & Bülthoff, 1995).

The two types of theory may be better described as extremes on a continuum rather than as independent processes (Bülthoff, Edelman & Tarr, 1995), with a combination of the two necessary to produce a functional model of object recognition (Tarr & Bülthoff, 1998). VI representations might be of use in certain conditions; for example, without some sort of viewpoint independence operating, different representations could arise because of the ambient lighting. An object first met in bright sunlight would be unrecognisable under fluorescent lighting. Although recognition does appear affected by the lighting (Tarr, Kersten & Bülthoff, 1998), this is a small effect compared to that of the viewpoint. The information derived from the lighting could also be useful;

shading may provide data about the three-dimensional shape. In this instance, the two types of representation could act in tandem.

If a choice of representations based on task demands is proposed, then this implies that more than one type of representation may be available at any one time. VS representations appear related to the level of similarity of a set of objects that need to be discriminated or recognised. When the similarity is high, then VS representations are predominant over VI representations. This would indicate VI representations for between category processes, and VS representations for within category processes (Tarr & Gauthier, 1998). Liter (1998) showed that there was a difference in whether VI or VS processes were used depending upon whether the features used in the task were based in the parts (VI processes) or the connections between parts (VS processes).

In summary, one problem for comparing the two representations during object recognition is that they have to be comparable. Two approaches have been described to achieve this, involving VI or VS representations. The strength of the VI representations lies in the invariant information contained in the representation. However, their weakness is that this information is insufficient on its own to allow successful recognition, particularly when similar objects need to be distinguished. To achieve this, the position of the invariant information also needs to be specified, and the addition of this information promotes a VS representation (see Lawson, 1999, for a review). Both types of representation, therefore, appear to have a role in object recognition.

However, this does not address the format of those representations in any depth. Next, we will review what information is contained in a representation, and how the information is organised.

CHAPTER 1

REPRESENTATIONAL FORMATS

Information in visual representations can come from two sources: the array of light arriving via the retina, and information already stored in memory. The perceptual representation is formed as the array of light reflected from objects in the environment strike the retina. The pattern of light is then transformed into a percept. Low level perceptual mechanisms focus on producing the original percepts from edges, giving information about non-accidental properties, or invariant features. Details on these processes are not discussed here. For the purposes of the research in this thesis, the main focus of interest at this point is how this information is represented in the percept and memory.

Early Models: Features and Templates

Traditionally, there have been two main approaches to the representational format. One approach posited that the representation was a whole, unified template, and the other posited that the representation was formed of smaller units, or features.

Feature Theories

Feature-based theories assume that the incoming representation consists of small segments, forming a list of features. The feature lists of each object can be stored, and then used for comparison. The exact information that constitutes a feature can vary among the specific theories. The information that is used for recognition in feature theories, however, appears an abstraction of the total information available (Eley, 1983).

The features can be broadly classified into two types. One type of feature consists of edge-based information and building up the whole can be thought of as a type of jigsaw puzzle. The other category consists of information such as different spatial frequencies, or principle component analyses. In these cases, building up the whole can be thought of as putting one transparent sheet over the top of another. Some of the information may be more relevant than other information for particular tasks. For example, low spatial frequencies give a view of the overall shape whereas higher spatial frequencies fill in the contrast.

It may be that the edges detected in the first category of features are derived by the processes specified in the second category; an edge occurring, for example, at a coincidence of several spatial frequencies. For the purposes of this thesis, the lower level visual processes will not be addressed in any detail.

One problem with these theories is how recognition takes place with only this level of information. For example, this information will allow a face to be recognised as a face, as opposed to something completely different such as a cat. It does not provide such a good explanation of how one face can be recognised specifically, or even distinguished from another.

Even if the basic features are processed into larger segments (similar to the non-variant features forming a geon) three more factors need to be taken into account. First, the segments themselves can differ among similar objects. For example, one cup handle may be a different size or shape from another, but the two objects are both still cups. The features themselves may need to be decomposed back into a subset of features - perhaps based on lines and angles. If so, the featural information may become very basic and require a unifying process to combine the basic features into a coherent whole; Gelstalt properties have been suggested for this role (see Saariluoma, 1992). If this does not happen, the problem of distinguishing very similar objects will return. For example, the cups may have identically shaped handles, but in different positions. This leads back to the problems in discriminating similar stimuli described above. The exact definition of a feature may vary depending upon the needs of the processing.

This leads to the second factor, which is the relationship among the features; their spatial arrangement, or configuration, can be important in recognition (e.g. Hole, 1994; Leder, 1996). For example, the difference between a cup and a bowl depends more upon the relationship between the handle and the container than the properties of the handle and container themselves. This information is lost from a representation containing merely a list of features, and must be incorporated alongside the featural information to allow full recognition. Such configurational properties may well prevail over featural properties, the latter being resorted to when the configurational properties are unhelpful or unavailable (Kimchi & Bloch, 1998).

Finally, the features have to be combined together into the correct object. For example, if a cup is placed in a saucer, the system has to correctly combine those features pertaining to the cup separately from combining those features that belong to the saucer. Failure to do this will not lead to successful recognition (Ashby, Prinzmetal, Ivry & Maddox, 1996).

Feature theories account for the level of abstraction that is often present in information stored in memory, but do not account well for the configural information which is essential for some cases of recognition.

Template Theories

The "whole" approach involves pictorial representations, an example of which is template theory. Within these theories the incoming stimulus can be represented in virtually the same format as the retinal image, and can also be stored in much the same format. This suggests that all the information is represented, and used, rather than a proportion of it. The implication in template theories is that one part of the stimulus is no more or less salient than any other part; it is the whole that allows comparison. A template can also be based on spatial locations rather than features.

The problem for which template theories are usually criticised is their lack of parsimony. Several templates will have to be stored for any single pattern. For example, to recognise the letter H, the templates would have to be stored in several fonts, for each font, a template in lower case and in upper case would be needed, for each of these a template in italics, and in bold are needed, and so forth. However, this is not necessarily a problem. The mechanism may well rely on storage processes rather than computational processes.

A bigger problem is that novel views can also be recognised. If this is taken beyond letter recognition to more general shape recognition, the problems become more apparent as the possible variations in shape for the same item or type of item become apparent. Furthermore, a shape can be recognised when it is partially occluded, or changes apparent size. Completing an occluded shape appears constrained by the need to maintain the elements and the regularity already present, suggesting feature-based rather than template representations (see Siddiqi, Tresness, & Kimia, 1996).

A further problem is that certain parts of the shape appear more salient to recognition than other parts (e.g. Hoffman & Singh, 1997) to the extent that caricature effects can occur (e.g. Rhodes & McLean, 1990). A template model does not readily explain this. However, Farah, Wilson, Drain and Tanaka (1998) explain that different parts of a template can carry different weights. The template formation can occur through some sort of parsing procedure. However, this suggests that the definition of the template is moving towards a feature-based definition. All the suitable templates stored in memory could be compared in a one to one matching of each point with the corresponding points on the incoming stimulus. This can then generate a ratio of matched points to unmatched points. The highest ratio signifies the closest match. Farah, *et al* (1998) claim that the template refers to the best overall fit, rather than a sum of the parts.

Templates are able to describe the perceptual image in a way compatible with the cognitive processes. For example, templates are compatible with the finding of global precedence. Global precedence is demonstrated by an advantage for global information, or information based on the whole shape, over localised information, in either accuracy or speed, during the early stages of object recognition. In addition, global information interferes with local information, but the reverse does not occur (Navon, 1977). However, the information stored in memory does not appear to share a template format. For example, memory can be affected by featural information alone (Reinitz, Lammers, & Cochran, 1992), suggesting that the information in memory is not stored as pure templates.

Structural Descriptions

Neither a feature-based model nor a template model is able to provide a complete account of the representational format. Feature-based models do not contain the configural information necessary for recognition. A template model does, but the featural information cannot be accessed individually. It would seem that some combination of the two is necessary for the practical application of the representations as cognitive tools.

One possibility is that the production of a prototype can take place. The prototype can be produced either by storing the common elements across a group of templates, or by producing an average, or norm, across examples. Storing the common elements would mean the representation would apply generally to all the items and specifically to none of them. This may be too limited for full recognition purposes, leading to the problem of differentiation between similar items. Norm based accounts of recognition have not received much empirical success (see Rhodes, Carey, Byatt, & Proffitt, 1998).

An alternative is the structural description. The structural description is another list form, with two lists present for each object. One list consists of the features, and, independent but associated, is a list specifying the spatial arrangements of the features. This combines both the completeness of the template format, and the succinctness of the features format. All the information is present, should it be needed, but in a format which is readily dissociable, allowing cognitive processing to occur on a sub-set of information. This not only appears to account for the two types of information used in object recognition processes, but they are in formats that allow the information to be transcribed into other cognitive processes too.

The main problem with the model in this format is that the two lists are posited as being dissociable. If they are not, then the overall tendency is towards a template situation. The dissociation of the material would be compatible with a VI approach to object recognition. However, findings during object recognition suggest that these two sets of information are not completely independent. For example, it takes longer to recognise a rotated object than it does to recognise the same object in an upright position (Shepard & Mezler, 1971). If spatial arrangement and featural information could be easily separated, this would not be predicted. However, not all researchers agree with this finding. For example, recognition of alphanumeric characters appears immune to orientation (Corballis, Macadie, Crotty & Beale, 1985). These findings could depend upon the level of recognition needed to complete the task, and the relative familiarity the observer has with the object (Hamm & McMullen, 1998).

Although structural descriptions appear a good basis for representations, the relationship between the two lists (parts and arrangements) needs more exploration.

Parts and the Arrangement of the Parts

Trying to categorise the representations into any particular type, such as featural or template, may not be the best method for explaining representations. An understanding of the representation might be better served by a specification of the parts, the arrangement of the parts, how these two pieces of information are related together, and how they are used in whatever task is under consideration.

Parts of Representation

One problem with representing the parts and the spatial relationship among the parts is that, as we saw with features, it can be difficult to discern exactly what constitutes a part. The definition of the part can depend upon the relationship the part has with the whole, and on the relationship one part has with any other part. Another problem is that the spatial list is expensive to maintain in terms of processing resources. This is particularly so if the spatial relationship is not crucial to recognition, because recognition is in any sense based on the parts. For example, although a steering wheel aids recognition of a car, the position of the steering wheel (e.g. on the left or the right of the car) is not important. Saiki and Hummel (1998a) also point out that even if the spatial relationship is important, even simple objects have several permutations of relations, and complex objects can have many relations to record. In addition, objects are rarely observed in isolation, and if relations are also computed across objects, then the resource use becomes extremely high. Furthermore, if the parts vary from being basic level features to larger segments, then the spatial relations will vary among the levels, too.

Saiki and Hummel (1998a) offer two solutions. One is to use the least number of parts that is practical. So, rather than operating at a low level of feature detection, using something like a geon or other volumetric primitive is more economical in resource use. The second solution is to employ the property of connectedness to avoid between object relation computations. This property could be used by weighting the relationships of parts that are connected relative to those parts that are not connected (Saiki & Hummel, 1998b). Other research suggests that uniform connectedness is as efficient as proximity and more efficient than similarity for grouping local components (Han, Humphrey

& Chen, 1999b). Connectedness also appears to play a role in some categorisation tasks, especially if the whole/part relationship is diagnostic for the category (Saiki & Hummel, 1998b). The parts, therefore, are probably better thought of as operating at the level of a geon (or similar level primitive) rather than at the level of an invariant feature.

Spatial Arrangement of the Parts

The spatial arrangement can include several items of information. First, it may define the direction that a part is from that of other parts. Second, it may define the distance of the part from other parts. Third, it may define the Gestalt relationship among the parts, for example, symmetry, closure, or connectedness. One problem with positing the representation of the spatial arrangement or configuration is that, as was the case with VI and VS representations, two forms may be needed (e.g. see Rhodes, Brake & Atkinson, 1993). One form is adequate for basic level recognition between different classes of objects, such as a car and a donkey, but will not suffice for recognising and discriminating between homogeneous classes, such as two similar cars. An alternative to needing two different levels of configural information is that homogeneous classes could be compared to a prototype, or norm of that class of objects (e.g. Rhodes & McLean, 1990). So, rather than the configuration of an object being defined by the relationship among the parts, the configuration is defined in terms of a comparison to a prototype. The norm is produced through expertise and experience with the class of objects. The prototype has not been a successful account when applied to faces (Rhodes, Carey, Byatt, & Proffitt, 1998).

First- and Second-Order Relations

These and similar terms are not used consistently by different researchers in the area. In particular, a distinction has to be made between first- and second-order features, and first- and second-order relational features. First order features refer to the parts of an object, and second order features refer to the configuration (e.g. Farah, Wilson, Drain & Tanaka, 1998). First- and second-order relational features refer to different levels (e.g. Cooper & Wojan, 2000) or different types (e.g. Tanaka & Sengco, 1997) of configural information.

Diamond and Carey (1986) introduced the terms first-order relations and second-order relations to explain the association of parts in an object; specifically, they were considering faces. Although Diamond and Carey did not make the exact terms explicit, the gist was that first-order relations are described in terms of the positioning of the parts, and second-order relations are a more precise definition of the spatial relationship among the parts. First-order relations applied when there were similar parts, but different configurations. Second-order relations described the situation where items shared the same parts, and also had the same configuration; that is, have a common first-order relationship. Diamond and Carey defined the same configuration as being when there were corresponding points on two or more items, and that the average of these points would also produce a recognisable example of the class of items. To allow discrimination among such items meant that the relationship had to be described in very specific terms.

It would appear the two types of representation are fundamentally similar, but vary quantitatively along a continuum. However, a continuum leads to uncertainty in clearly defining what type of relations are present, unless the representation is at one or other end of the continuum. Rhodes, Brake and Atkinson (1993) point out the ambiguity when trying to distinguish among them empirically. They explain Diamond and Carey's account in terms of three levels. The first level consists of the isolated parts themselves. Some objects may be distinguished by the completely different parts they have. First order relational features refer to the arrangement of the features where the position of the parts are not constant across the examples, such as might be found in heterogeneous classes. The usefulness of this is open to debate; items from different or varied classes often have different features, too. This would appear better, perhaps, as a coarse version of the second level information. Second order relational features refer to the precise arrangement of features where the gross position is constant across examples, such as would occur in a homogeneous class. The three levels occur on a continuum, and access to the second-order level is dependent upon expertise (Tanaka & Farah, 1993).

It is not clear what information is included in the first-order relations. In particular, it is not clear if distance information is included. Tanaka and Sengco (1997) interpret the first order as not having a distance measurement at all. For example, they suppose first order information to be "letter-box to right of door handle". Second order information supplies a distance measurement in one of two quantitatively different forms. For example, the same relationship as above could be expressed as "letter-box to right of door handle, wide gap" or "letter-box to right of door handle, 40cm".

Recent work by Cooper and Wojan (2000) found that a shared configuration is better defined by two objects having the same basic parts, and the same positions, defined by such terms as "above", "to the left of", and so forth. They claim that first- and second-order relations are not adequate to explain the empirical findings, and they offer an alternative account. In their account, the relationship among the parts falls into one of two types: categorical or co-ordinate descriptions.

Categorical Descriptions

The categorical description conforms to the classic idea of a structural description as defined in viewpoint-independent theories of object recognition. This description is among all the parts in a representation. The relationship of each part with every other part is specified in categorical terms, based on the direction, with no reference to distance. The categories will include such terms as “above”, “below”, “to the right of”, “to the left of”. It is not clear whether properties such as “parallel to” or “symmetrical with” would also be included. If they were not included here, it is not certain where they would occur. Although they have no reference to distance, they are not really directional, either. Several researchers consider such Gelstalt relationships to be important in encoding the configural information (see Kimchi, 1994; Kimchi & Bloch, 1998).

The categorical description can explain several findings in object recognition. First, objects with slightly different positions of features can still be grouped together, because the categorical description is constant across the examples. For example, a car can be recognised as a car, even when the exact positions of the parts vary across makes (consider an Alhambra, a Corsa and a Ferrari). This is understandable if categorical descriptions are used. The doors, headlights, bonnet and so forth are still in the same positions if they are considered in purely directional terms. This leads on to the second point. Categorical description will be relatively tolerant of noise. Providing one part remains above the other, for example, the exact relationship can vary and still retain the same description. However, these strengths also give a weakness to their use, because a categorical description will not differentiate between the three makes of car, or indeed, between any objects of the same class of items.

The categorical description would also appear to have an empirical base in viewpoint-independent theories of object recognition. Although the descriptions are not completely dependent of orientation, their generalisability and tolerance of noise allow them to be flexible across orientations, up to the point where the direction of the relationship changes. The references are also centred with the object. The description is segment-to-segment, and hence the references for each part are relative to other parts within the object.

Co-ordinate Descriptions

The co-ordinate description not only gives the direction of the relationship for a part, but also specifies the distance of that part from a point that is separate from all the parts of the object. The object can be imagined as having a grid overlaid upon it, as if it had been converted into an Ordinance Survey map, with appropriate co-ordinates for each part. For example, the "Q" key on a standard keyboard could be described as "four units right, eight units up" if the common reference point was just below the bottom left corner; the "G" key could be described as "nine units right, six units up". If the grid is outside the frame of reference of the object, then these units will have to be absolute. If any relativity was present (e.g. for the "Q" key the height is twice the width of the space), then there must be some object-centred properties because each relative measurement would need at least one other measurement.

The resulting description of all the parts will give a very precise arrangement, although this precision is likely to be at the cost of processing resources. The co-ordinate description will not allow the generalisability of the categorical description, nor will it be as tolerant of noise. It would, however, allow the Alhambra, the Corsa and the Ferrari to be readily distinguished from each other. The reference point arises outside any particular feature, although it may be confined to a particular position relative to the viewer (e.g. bottom left corner), making it viewer-centred. The resultant representation will, therefore, be viewpoint-specific.

One unresolved point is how the co-ordinates are produced in the first place. The configuration of the parts appears to be encoded alongside the parts, rather than being a product of later learning (Ceraso, Kourtzi, & Ray, 1998). However, the co-ordinate description, as opposed to the categorical description, may arise as the products of experience or expertise, as with the second-order relations. If so, the question then arises of whether the introduction of the distance parameter is an all or none procedure. In the categorical level, does the distance remain completely unspecified, or can it be introduced if warranted? Forming this representation would be costly in terms of processing resources, but a localised or relative (possible in terms of "wide" versus "narrow") distance may be feasible.

Rhodes, Carey, Byatti, and Proffitt (1998) have supplied some support for the view presented by Cooper and Wojan (2000). They found that configurations of homogeneous classes were best explained by an absolute coordinate system. Furthermore, they found that the use of this system did not depend on expertise. Whether the two coordinate systems described share enough similarities to be comparable has yet to be established.

In summary, the representation will contain segments, arranged in a certain way. The arrangement of the segments is known as the configuration. There appear to be several levels of configural information. One area to be examined is how the different levels relate to each other, and how this will affect the processing of the representations.

Hierarchy of Representations

Holistic and Segment-Based Representations

Early work considered representations as a dichotomy. Representations were either holistic, and subject to holistic processing, or part-based, and subject to piecemeal or analytical processing. Later work suggested this was too simplistic, and a continuum of representation was proposed, whereby a representation is located somewhere from one extreme as a single, holistic unit, to the other extreme of individual features (e.g. see Farah, Wilson, Drain & Tanaka, 1998). The position on the continuum could be determined by the nature of the object, by past experience with the object, by the task to be performed, or, indeed, by any additive or interactive combination of these factors. Distinguishing between a type of representation and a corresponding type of processing may, however, not be so straightforward.

The representation contains parts, and the configuration of the parts. In a segment-based representation, the parts are relatively independent. In a segmented or part-based representation (the two terms are used interchangeably throughout this thesis), the parts and configuration of the parts will be more easily separated; the parts themselves can also be processed independently of each other, in a piecemeal process.

Holistic processing involves interaction among the parts, and is identified when one of the parts cannot be isolated from the other parts. The definition of a holistic representation is not as clear as that of a segment-based representation. It is generally agreed that the configural information is predominant in such a representation, but the fate of the parts is not as clear. Some theorists suggest that the parts are present, and still relatively well defined. Other researchers are less specific on the fate of the parts, merely emphasising the importance of the relations relative to the parts (Farah, Wilson, Drain & Tanaka, 1998). The configural information could actually be considered as another form of feature (Macho & Leder, 1998); possibly an "emergent" feature arising from the relationship among the parts (Pomerantz & Pristach, 1989). The idea of emergent features has been criticised for its lack of clarity, and some researchers prefer to specify the relationships, usually in terms of Gestalt principles (Kimchi, 1994). The parts may maintain some measure of independence, but will influence the processing of each other (Sergent, 1984).

The alternative view is that the parts lose their structure in a holistic representation. For example, Palmer and Rock (1994) claim that when a representation displays uniform connectedness across all the parts, there is no internal structure. This view is echoed in that of Tanaka and Farah (1993), who claim that in a holistic representation the between part relations are specified to the same extent as the within part relations, resulting in no internal structure for the parts, and a template for the representation. This extreme view is not without its critics. Donnelly and Davidoff (1999) point out that if the representation is a template, then the parts would not be segregable at all, and yet parts can be detected in such a representation. Pomerantz and Sager (1975; p461) capture the argument very well: "An element can exist in isolation without a configuration, but a configuration must by definition be constructed from elements".

Whether an object has to be parsed into parts at an early stage of processing is not clear (Farell, 1984). Palmer (1977) considered that the first stage of processing was a representation of an undifferentiated whole, followed by parsing the representation into suitable parts. When this was accomplished, the representation was re-formed back into the whole, but the whole was now clearly structured, rather than undifferentiated. Even if the object is parsed, it is not always clear when a relation is purely a relation and when it is a feature (e.g. see Rhodes, Brake, & Atkinson, 1993). For example, thinking of a belt, is the area of material around the holes better described as a part of the belt, or the spatial arrangement of the holes in the belt?

This confusion has led some researchers to use an operationalised definition of holistic representations (e.g. Tanaka & Sengco, 1997). The "whole" shows some degree of interaction between the parts and the configural information. They claim that a feature (part) is the "parts of an object or a face that can be identified on the basis of its natural discontinuities" (p592). The configural information is "the difference in recognition [performance] when a feature is tested in isolation versus when tested in the context of a whole object" (p591). The part and configural information is encoded together for holistic representations, and the parts are not independent of each other.

Processing the Representations

The format of the representation has the potential to affect how it is processed. For example, segment-based representations lend themselves readily to piecemeal processes, in which each part is processed one by one.

They also allow one part to be processed independently of the rest. Holistic representations lend themselves to processes where the configural information is more important than the parts for successful completion of the task.

Holistic processing will take the representation as a single unit, and transform it as a single unit. In a holistic representation, all the information is represented as a single unit, or at least integrated together. The parts and configurational information will be maintained throughout the process. In piecemeal processing each part will be transformed, then the other parts re-aligned. One way this could be achieved is by the first part being transformed to the full extent of the process, then the other features being re-aligned. Alternatively, the first part could be partially transformed, the other parts re-aligned, then the process repeated until the full transformation is completed. This means that the configural information is disrupted and re-constructed.

However, the segment-based representation is also open to other forms of processing. For example, the parts, rather than being processed one by one, could be processed in parallel, giving the appearance of a holistic representation. The entire image does not have to be processed in a part-based representation; a sub unit may be processed. This may also give the appearance of a holistic representation if the sub unit can effectively be transformed as a single unit.

Piecemeal processing may be more flexible than holistic processing. Piecemeal processing appears to allow a wider variety of transformations (e.g. folding or bending, where the configural information is broken and reformed). Holistic processing may impair such transformations, because all the parts are constrained by the configural information. Piecemeal processes may facilitate transformations, through allowing individual parts to separate and re-combine in a variety of ways. However, although piecemeal processing may be more flexible, holistic processing may be "safer". It is safer in the sense that the spatial information is more easily preserved and remains relatively intact throughout the process, rather than being re-constructed. Therefore, even though piecemeal processing may produce the more flexible processing, if there are no representations or task demands to re-reinforce its use, holistic processing may used.

The availability of both piecemeal and holistic processing enables flexibility; this maximizes the accurate processing of information, in face of constraints imposed by the circumstances (see also Cohen & Kubovy, 1993). Different representations appear more or less suitable for different processes. The next area we consider is how these different representations relate to each other.

Levels of Representation

A dichotomy of representation is too simplistic; a continuum of representation is a better option (e.g. see Farah, Wilson, Drain & Tanaka, 1998). However, there is a third option; a hierarchy of representation, in which one representation is nested in a representation at another level (e.g. Navon, 1977). Each level of the hierarchy is a complete representation, but in addition is an essential contributor to the levels above and below. The latter two models appear similar, but the continuum model suggests a single type of representation at any one time per object, whereas the hierarchy model suggests that access to multiple representations at any one time is possible. There is some debate in the literature in the field about whether the two models can be alternatives to each other, or whether the information they represent is actually orthogonal (e.g. see Kimchi, 1994).

Palmer (1977) found a simple geometrical line figure had three levels of representation, and he posits a basic unit at each level. The first level was the whole, or global unit; the second level was the component parts; and the third level was the lines from which the other levels were formed. At each level the basic unit is a representation in its own right.

The units can be described in terms of their relationships with other basic units on the same level, and also with superordinate and subordinate units; however, they retain an independent status. Every level is functional as a stimulus in its own right (Kimchi, 1992), with one of the levels as an "entry" level (Kimchi, 1998). There may be different types of hierarchies. Some of the levels that are closer together are found to be difficult to separate, suggesting that they may not represent independent hierarchical levels (Kimchi & Palmer, 1985). If the entry level consists of many small parts, the entry level and the global level may be better viewed as separate representations, rather than one nested in the other. Alternatively, if the entry level consists of fewer and larger parts, this level may interact with the global level. One level of representation is automatically associated with the other, and to some extent, cannot exist without it; the object and its parts (Kimchi, 1998). The idea is that parsing and grouping are taking place together, with the purpose of structuring the representation (Saiki & Hummel, 1998a).

There is some debate concerning the order of processing of the different levels. Navon (1977) suggested that processing went from the highest, global level down to the lowest, atomic level. This arose, originally, from the global advantage. Navon found that the higher levels showed an advantage in terms of more accurate and faster responses to the information than was found for information at the lower, or local, levels (see Kimchi, 1992, for review). Global information also interfered with local information more than local information interfered with global information. Other evidence has supported middle out progression (Kinchla & Wolfe, 1979), or a beginning at the most salient level (Hoffman, 1980).

Both high-level down and low-level up accounts are subject to some problems. The low-level up account cannot explain the finding of advantages for global processing (see Navon, 1977). However, there are also problems for a strictly high-level down account. Many factors have been found to interfere with and remove the global advantage (see Kimchi, 1992). Kimchi includes visual angle (e.g. Lamb & Robertson, 1990), acuity (e.g. Navon & Norman, 1983), location on the retina (e.g. Pomerantz, 1983), goodness of form (e.g. Hoffman, 1980) and exposure duration (e.g. Paquet & Merickle, 1984) in the list of factors. One way out of this paradox is to consider whether the two processes could be interactive (e.g. see Vecera & Farah, 1997, for such a process in figure-ground segregation). Instead of full processing at one stage preceding the processing of the next stage, partly completed outcomes can be offered as input. In this way, processing of a "later" stage can begin before an "earlier" stage is finished, and can even feed in outcomes to the earlier stages. Such a model has not been supported in the global/local domain, however. Sanocki (1993) states that global information is processed first, and local processing uses this as its input, hence local processing cannot begin until the global processing is complete.

Love, Rouder and Wisniewski (1999) suggest that to identify the whole stimulus requires local components to be grouped together, probably using Gestalt principles. From this, the global shape will emerge. The question that arises is whether the local component has to be identified or processed as an entity in itself before grouping can occur, or whether the global shape can emerge before local processing takes place. Love, Rouder and Wisniewski found evidence that local components did not need to be identified before being grouped. However, in other circumstances, for example, if the task focused on local information, or grouping can not proceed easily, then a local to global progression could not be ruled out.

One questions that arises concerns the relationship between hierarchies and configurations. Are more global levels in any way equivalent to more holistic representations? The global advantage has been taken as an indication of holistic representation, or processing (e.g. Robertson & Lamb, 1991; see also Donnelly & Davidoff, 1999; Tanaka & Farah, 1993). However, the relationship may not be as simple as this suggests. Kimchi (1994) pointed out some problems with equating the two systems too closely. One point raised is whether the global and local properties can devolve into an issue of the size of the stimulus, which is definitely not the case with the holistic/part-based distinction. Another point is whether global and local hierarchies are a processing hierarchy rather than a representational hierarchy. Configural, or holistic, properties are a function of their relationship among the parts; that is, they may be processed first, last, or at any point in the middle of a series of processes. Global properties are a function of their position in the hierarchy; they are at the highest level, and therefore they are processed first. However, this then begs the question of why they are processed first. For example, in cases where a local advantage is found, and small letters appear to get precedence over the large letters, are the small letters elevated to a higher level? Do they become global? If not, then this difference needs further clarification.

Kimchi (1994) provides evidence for a dissociation between configural and global properties. Configural properties aided a classification task regardless of whether the configural properties were at a global or local level. A global advantage was only found when the non-configural property was used for the task. Further evidence, however, suggests that Gestalt properties may be associated with hierarchical processing (Love, Rouder and Wisniewski, 1999), although not all the Gestalt properties act in the same direction in the hierarchy (Han, Humphreys, & Chen, 1999a). Gestalt properties have also been associated with configural information, in that they can define the relationship among the parts in a representation.

The nature of the representation and the relationships between and within the representational formats are still open to debate. We now consider factors that may influence the format. The most important of these within the context of the research presented here is the meaning associated with the object being represented.

CHAPTER 1

THE ROLE OF MEANING IN VISUAL COGNITION

Object recognition is a process that provides information about the objects in the environment. For object recognition to be useful in interacting with a physical object, information about the physical object is required. For example, knowing whether an object is heavy or light can determine how the object should be moved. In addition, the semantic information could aid the shape information in determining what the object is (Boucart & Humphreys, 1997). The semantic information is closely tied in with the structural representation; however, it is not certain whether it influences the format of that representation.

The Meaning of Meaning

The first problem in looking at meaning within the cognitive system is to find some sort of agreement on what meaning is. The researchers within the realm of object recognition attach importance, as far as meaning is concerned, to the stage at which a name is attached to an object. As we have seen, many theorists posit a cascade mechanism for object recognition; when the process has progressed sufficiently that a name can be generated, the associated information in semantic memory has been "released". The association of meaning and name suggests that perhaps we should be looking within the realm of language and verbal behaviour for the root of meaning in the cognitive system. However, the tradition within this area is that the meaning of a word is closely linked to the recognition of the object it represents. This can be through one of three routes. First, it can replace the object with a verbal symbol (e.g. Watson, 1920). Second, it can acquire a response to an object (e.g. Dewey, 1925). Finally, it can recognise an underlying concept or category (e.g. Fodor, 1987).

Behaviourists think that a word becomes associated with an object until, eventually, the word can stand in place of the object. This idea was similar to the more cognitive idea of categorisation. Plato claimed that the correct name maps directly onto the underlying concept; this view is echoed with more recent philosophers who consider meaning to arise from the images and concepts associated with objects. Cognition was the manipulation of mental representations; these representations constituted an innate language of thought. Language was equated with pre-existing concepts. There may a distinction between forming an abstract concept and performing the practical task of assigning category names. It has been suggested that concept formation is based on knowledge, whereas categorisation is based on explanation (Komatsu, 1992).

However, the problem with using these theories as an explanation of meaning is that we have already described them as part of, or at least associated with, the object recognition system. It would appear that the researchers in object recognition look to the language researchers for an explanation, while the language researchers are looking towards the object recognition researchers for an explanation.

Learning to name objects, which presumably involves acquiring the meaning of the objects and the names, would appear, therefore, to be a two way process. Both appear able to access the same semantic representations, but also to retain specific representations (Bajo, 1988). However, the two types of representation appear additive, which suggests a level of independence (Paivio, Walsh, & Bons, 1994). A name can invoke an object; and an object can invoke a name. The problem with this is that it is circular. At some point the learning needs to be set against a context in which meaning develops for both the object and the name.

This still does not give an exact definition of meaning in terms of cognitive processing. Looking for the meaning of meaning in absolute terms appears fruitless. Perhaps the best way forward is to develop a more operational definition of meaning, by looking at the processing of meaningless and meaningful items. However, cognising meaningless stimuli is also problematical. We did not evolve our cognitive systems in a meaningless environment - the environment is made up of meaningful objects. Furthermore, the whole purpose of our cognitive systems is to help us to make sense of the environment; to pose meaning on the patterns of light, or sound, or smells which arrive through our senses. Therefore, there may be a significant difference between the way we cognise an object which is meaningful, and one which is meaningless.

Studying cognition through the use of meaningless stimuli has been a long tradition; for example, learning nonsense syllables. There were good reasons for using this type of variable; meaning could influence the memory for a word. Miller (1956) found that the capacity of short-term memory was seven, plus or minus two, units. However, “chunking” the material together, or combining smaller units into larger units could increase this capacity. The definition of a unit became related to the meaning, because through meaning, small chunks could be grouped together. For example, six numbers could be associated together through the meaning of a birthdate.

Recognition memory is large for meaningful, visual stimuli (Nickerson, 1965), although studies of the memory of scenes suggests that the memory for details is not as good as reputed (e.g. Pezdek, Whetstone, Reynolds, Askari, & Dougherty, 1989). Incongruent items in a scene may be better remembered than items that are more congruent, and therefore perhaps more meaningful for that scene. Memory would appear to be better for meaningful items, producing better performance for items which are more "object-like" than those which clearly do not resemble an object (Mou, Anderson, Vaughan, & Rouse, 1989). This difference in memory could be due to changes in the structure of memory, or changes in the representation of the information.

This argument can be taken a stage further. Do objects have meaning from the very first time we meet them? Or do we have to learn the meaning of objects? If we do, then how does the representation and processing change as an object acquires meaning? And what stages does it have to pass through to get meaning? If we want to examine meaning, we need to compare the processing of meaningless and meaningful items to see how the one becomes the other. If we can not name an object until we have semantic information, how can people associate nonsense names with non-real shapes? What happens the first time an object is met; can the name not be learned until some information is also learned? This suggests that with regard to objects, if semantic knowledge has to be triggered before the name can be triggered, the amount of semantic knowledge needed may be extremely small. It may even be that familiarity is all that is needed, in the sense that the semantic knowledge amounts to "I saw a similar object yesterday, and that object was called a glock".

The Role of Expertise in Representation

The information provided about objects can exist at one or more of several levels. First, an object may be recognised as having been encountered before, and therefore familiar to the observer. The observer may have no other information than this; or may incorporate it with information derived specifically from the previous encounter, for example, whether the object was moving or still. The object is therefore familiar to the observer, but related information is limited in scope. The observer has met the object, may even be able to give the object a name, but does not know the purpose of the object, or how best to interact with it. In some cases this information may be known and the object familiar, but some information is not present in the representation. For example, despite handling coins on a daily basis, most people are poor at being able to report the details on a coin (Nickerson & Adams, 1979; Jones, 1990; Jones and Martin, 1992). The information selected for the representation could depend on the meaning put on the stimulus by the observer (Chambers & Reisberg, 1992).

Second, an object may be recognised as belonging to a certain category. The object itself may not have been encountered before, but other objects within the category may have been. Again, this can be accompanied by other information. Common properties for objects within the category can be linked with the new member of the category. In this way, an unfamiliar object can become linked with a substantial amount of information through the similarities to other familiar objects. Tasks located at this level have been linked with VI representations, and the use of part information (Tarr & Gauthier, 1998).

Third, the object may be identified as being a very specific object. This will generally involve the object having been encountered previously, and with sufficient information that the object is nameable. The ends of such processes are often seen as providing the name of the object. However, once the name has been generated, all information pertaining to that object, and that class of objects, is also available. Providing the semantic information would, to a large extent, be a more useful purpose when interacting with the object than merely providing an arbitrary symbol to stand in place of the object. The process links three types of information: perceptual information, semantic information, and verbal information. At this level of recognition, very similar objects need to be discriminated. Such tasks rely more on VS representations and configural information (Tarr & Gauthier, 1998). One important question is whether the same representations are available in all cases, and task demands determine which are used; or whether the information available can determine the type of representation that is formed.

An area worth considering in this light is that of expertise. Expertise is generally considered in terms of the familiarity that a person has with a topic or items, but alongside that familiarity is also an issue of meaning. Neither familiarity nor meaning *per se*, however, captures the whole of expertise. What is expertise, in cognitive terms, and how can it relate to meaning? Gauthier, Williams, Tarr and Tanaka (1997) asked a very similar question with regard to familiarity. When does an expert become an expert? When does a person change from novice to competent to expert? Also, what, exactly, does this change entail?

Experts differ from non-experts in that they have acquired skills or knowledge within their area of expertise. This can manifest as a difference in memory performance. For example, Hassebrock, Johnson, Bullemer, Fox, and Moller (1993) looked at different levels of expertise in physicians, and their abilities to solve diagnostic problems. There were no immediate differences in performance. A week later, although the non-experts could recall more general information, the experts could recall more relevant information. Saariluoma & Kalakoski (1997) found that, in blindfold chess, experts were quicker at taking in the necessary information, and more accurate in their responses than the non-experts. Again, the difference between the experts and the non-experts may not just be in the use of memory, but in the type of representations in terms of the information used (Diamond & Carey, 1986).

Research has found a difference in recognition or discrimination performance between experts and novices in diverse topics (e.g. chicken sexing, Biederman & Shiffran, 1987; x-ray interpretation, Norman, Brooks, Coblentz, & Babcock, 1992; blindfold chess, Saariluoma & Kalakoski, 1998). Modigliani, Loverock, and Kirsen,(1998) suggest that experience may modify the salience of a particular piece of information. When the task involves distinguishing between very similar stimuli, one source of information used by experts to a greater extent than novices is the configural information. For example, Diamond & Carey (1986) found that sporting dog experts show more configural processing than do novices. Rhodes (1995) summarises the research by saying that experts use relational features as opposed to isolated features.

Expertise is often defined through operationalisations (see Stevenage, 1995, for a review). For example, an expert is deemed an expert if they display the "differential inversion" effect (Yin 1969). Performance in some task (usually a form of recognition) with an upright object is disrupted when the object is inverted. To display this effect, the object needs to be monooriented (to give a perceptible effect of inversion); and the observer needs to be an expert with the object (Diamond & Carey, 1986). It is thought that the inversion of the object disrupts configural processing.

There are criticisms of operational definitions of expertise. First, they are indicators of the use of configural information with highly homogeneous sets of objects. Whether they are indicators of expertise beyond this is not clear. It limits the definition of an expert to someone who can tell apart very similar objects that they are familiar with. Although this is true, it does not give a comprehensive definition of the term. In some cases the value of the experts is that they are able to combine very different objects together and recognise them as the same class of objects. For example, physicians may be able to look at symptoms that are not immediately similar, but still be able to group them into the same diagnosis (Hassebrock, Johnson, Bullemer, Fox & Moller, 1993). Second, any operationalised definition is at risk of becoming circular, and this is no exception. An expert is someone who displays an inversion effect, and they display the inversion effect because they are expert. The use of this definition in research is saved from circularity by relating it to configural processing. However, this automatically limits the scope for examining expertise.

Gauthier, Williams, Tarr and Tanaka (1997) provided an alternative definition. They claimed that the criterion for being an expert was that the candidate was as fast to recognise examples of a category as they were to recognise the category. This gives the idea that novices can function at the category level, but their performance falls at subordinate levels. However, it may be more useful in general terms to consider expertise as being no difference in performance between category and subordinate levels of the category.

Gauthier *et al* (1997) looked at this issue when they trained volunteers to become "greeble" experts. Greebles were nonsense figures that came in two genders, and four/five families. Gauthier *et al* found that there were multiple facets to expertise at any one time, and that expertise may not be a unitary skill over a period of time. The first facet to be demonstrated was the ability to perform tasks using the overall shape. After this, the participant was able to use individual parts. Next, a combination of these parts could be produced, and finally, the configuration was available. This also implies that being an expert is not an all or none event. There may be various stages involved in travelling the road from novice to expert.

The change from novice to expert is accompanied by a change from the use of part-based information to the use of configurally-based information. Although familiarity has been clearly shown as a factor responsible for this change, it may not be the only factor. Evidence from comparative studies suggests that global precedence is a result of a top-down factor (Fagot & Deruelle, 1997). Baboons show a local precedence. The global precedence found under the same test conditions in humans was not due to familiarity (that is, previous experience with the stimuli) because this was comparable in the baboons, and the global precedence was not removed when unfamiliar stimuli were used. Although this is not evidence for meaning *per se* having an effect, it does suggest that familiarity may not be the only factor involved.

The Role of Meaning in Representation

Saiki & Hummel (1998a) claim that representations are broken down, then presented as parts and arrangements for a memory match. Palmer (1977) adds in an extra stage. The undifferentiated input is broken down, then re-built in an ordered, structured fashion. However, what drives the "building up" is not clear. One property could be connectedness of the object and its parts (Han, Humphreys, & Chen, 1999; Saiki & Hummel, 1998a). If expertise is associated with configural processing, what happens when we meet an object with which we have no expertise at all? Three options appear available. First, it is treated in the same way as if there were expertise. Second, it is broken down, but as the lack of configural information will prevent full use of the whole representation, a part-based representation will have to be used. Finally, it remains as an undifferentiated input.

Palmer (1977) suggested that representation went through three stages: an undifferentiated whole, parsing into parts, re-assembling into a structured whole. Alternatives to this view include the global precedence view (e.g. Navon, 1977) where a visual input begins as a whole unit, and is broken down into smaller and smaller units. Conversely, the visual input could begin as basic level features and be built up into the whole (e.g. Biederman, 1987). If each level is a stimulus in its own right, as Kimchi (1994) proposes, then meaningful and meaningless stimuli may vary. Meaningful stimuli can be seen as having a clear hierarchy. For example, a human being can be seen as a whole. The next level could contain the head, torso, legs and arms. The next level again will represent the thigh, knee, calf and foot. The hierarchy goes on with the heel, ankle, and toe, and so forth. Each level is as meaningful as the whole; each level even has clearly labelled parts operating at that level. A meaningless stimulus could be broken down in exactly the same way into similar parts, but would the parts retain the same hierarchy of structure as the meaningful stimuli? If a human being was an object that had never before been encountered, could it be structured in the same fashion as described earlier, but without the names, and based purely on shape? Findings in the global precedence literature suggest such hierarchical organisation operates post-perceptually, and so may not be available without top-down influences, such as meaning, to drive it (Boer & Keuss, 1982).

When children recognise objects they rely heavily on shape (Landau, Smith & Jones, 1998). Landau *et al* noted that as the children got older, the function of the object increased in importance for identification. By about ten years old, function could, if necessary, override form (e.g. Richards, Goldfarb, Richards & Hassen, 1989). Function may also serve a purpose in the early learning of words and labels (Nelson, 1974). Adults used function to aid identification of unfamiliar objects, and, although with familiar objects they did not use function *per se*, they were also reluctant to rely totally on shape at the expense of function information (Landau, Smith & Jones, 1998).

Other research has also developed the idea of form versus function (e.g.Boucart & Humphreys, 1997). Matching tasks involving local, physical properties (colour was used) were not affected by semantic information, but when the matching task became more global, semantic information was used. At some point during object recognition semantic information is integrated into the process. If it comes after the identification is completed then it will not have much capacity to affect the representation. If it is incorporated during the process, it may have more effect on the representation. The reason semantic information appears to be accessed before the name may be because it aids the process of recognition in meaningful objects.

The importance of the meaning of the stimuli has been demonstrated in the area of face perception. For example, Donnelly, Humphreys and Sawyer (1994) showed that during a categorization task normal faces were matched using a holistic comparison to a mental representation. In contrast, faces that had been re-arranged were processed using a sequential feature by feature comparison. Purcell and Stewart (1986) found that faces could be remembered better than the single features of which the face was composed; in subsequent research they suggested that being able to attach a meaning to a visual stimulus could enhance detectability (Purcell & Stewart, 1988). Suzuki and Cavanagh (1995) demonstrated that when a set of features were organized into a face, it impeded a search for the individual features, relative to when the same features were organized into a meaningless pattern. Therefore, it seems well documented that when meaning is removed from facial stimuli, the processing of those stimuli undergoes a change; stimuli that are more meaningful are more holistic in nature.

However, faces are highly specific stimuli that may have specialized cognitive processing that differs from processing of non-facial stimuli (e.g., see Farah, Wilson, Drain, & Tanaka, 1998). Studies comparing faces with non-face stimuli have produced mixed results. Davidoff and Donnelly (1990) found that chairs were represented and processed holistically. Tanaka and Farah (1993), on the other hand, found that houses were represented and processed segmentally. Donnelly and Davidoff (1999) concluded that houses were represented segmentally, but could be processed holistically. Furthermore, in the studies described above, the "meaningless" objects are scrambled versions of the normal objects. It is possible that the recognisable parts still point to the origin of the stimulus, and that the results are a result of unfamiliarity with the configuration of the parts presented in the task. As we have already discussed, there seems a clear link between familiarity and configural processing. Ngohayon, Kawahara, and Toshima (1999), looking at Japanese kanji, claimed that meaning acted to encourage holistic representations, and that this effect could extend to meaningless kanji that were similar to meaningful kanji.

To summarise, meaning certainly appears to influence the structure of memories for items, and their processes, but whether this is associated with an effect on the representation is not so certain.

THIS RESEARCH: PURPOSE AND AIMS

The purpose of this research is to examine the role played by meaning during processes relating to object recognition. To understand the processing, we need to understand the building blocks that form the processing, and this means we have to understand the mental representations.

What information do we need to understand mental representations? First, one central piece of information is whether the representations are stored in a holistic way, or a piecemeal way. Second, we need to know how the representations are processed. If they are processed holistically, then the representation is transformed as one unit. If they are processed in a piecemeal fashion, then the representation is transformed as several segments. The form of the underlying representation can place constraints on the type of processing which can be carried out. Meaning may play a role in determining this form. Third, the processing itself can take several forms, given the same type of representation. Meaning may, again, play a role in which process is selected. Finally, the output of the process, or series of processes, has to allow a response to be made. Meaning may determine how the output will be used.

The research presented here cannot hope to answer all these questions, but is meant to shed light on the kind of evidence that ultimately could. The broad purpose is to investigate whether meaning can affect the representations that are involved in object recognition, and, if so, how meaning does this. In order to arrive at a conclusion, however tentative, the research aims to compare processing between meaningful and meaningless objects, which are otherwise equivalent. It is assumed that the format of the underlying representation will affect the processing of that representation. The specific purpose of this research is to determine whether there is a difference in the representations of meaningful and meaningless objects, and if so, what that difference is. The particular difference addressed is whether the representations are holistic or segment-based; and, if one of the representations is holistic, to determine what type of holistic representation is occurring. Past research has produced three definitions for holistic representations. The first definition refers to a representation that is undifferentiated and consists of one unit. The second definition refers to a representation that has parts and configural information, but the importance of configural information outweighs the importance of the part information. Finally, the third definition refers to a part-based representation processed in parallel, and thus appearing holistic..

Looking at the representations directly is not possible. Several methods of testing for the type of representation have been used (see Farah, Wilson, Drain & Tanaka, 1998). This research aims to examine the issue using three different approaches. However, a common theme throughout the research presented here is the use of the complexity of the stimulus as an analytical tool. If the representation is predominantly part-based, then a complex stimulus will be represented by more parts than a simple stimulus, assuming similar sized parts. If the representation is mainly holistic, then the predominance of configural information will either remove or lessen the importance of individual parts in the representation. In addition, the parts in a holistic representation will be less independent than will those in a part-based representation.

To conclude, the question asked in this research is whether meaning influences the format of the representation of an object; specifically, whether meaning encourages a holistic or a segment-based representation. This question will be addressed by comparing the performances of two sets of stimuli, which differ on the meaningfulness, but are comparable on other parameters.

CHAPTER 2

STIMULUS DESIGN AND ASSESSMENT

ABSTRACT

The purpose of this chapter was to outline the design of the stimuli that were used in the tasks reported in the following chapters, and to validate the suitability of the stimuli for these tasks. First, the design of the stimuli was described. Two types of stimuli were produced, meaningful and meaningless objects, produced as two-dimensional line drawings. Each type of stimulus had two levels of complexity, simple and complex. These two factors, meaningfulness and complexity, were assessed. Meaningfulness was examined by asking participants to name the stimuli, then to rate them for confidence in naming, familiarity and representativeness. Complexity was measured by using a compactness value, calculated from the square root of the area of the stimulus, and the number of lines contained within it. There were significant differences in meaningfulness that were not affected by complexity. There were also significant differences in complexity that were not affected by meaningfulness. Therefore, the stimuli were deemed suitable for inclusion in the following experiments.

A version of this chapter has been included in an article submitted to *Psychonomic Bulletin and Review* under the title “The Role of Meaning and Familiarity in Mental Transformations”.

CHAPTER 2

INTRODUCTION

The research presented in this thesis was looking at whether the representations of meaningful and meaningless stimuli differ. Two sets of stimuli were needed, therefore; one of meaningful objects and one of meaningless objects. In addition, the studies used the complexity of the stimulus as a tool to examine the underlying representations. This means that for each type of stimulus, half the stimuli needed to be complex and half needed to be simple. The assessments for meaningfulness and complexity are described later in this chapter. First, suitable stimuli had to be chosen.

Several types of stimuli have been used before to assess the type of representation. The list of stimuli includes letters, geometrical two-dimensional and three-dimensional shapes, houses, and chairs. However, many of the previous studies looking at holistic and segmented processing were looking at the representational format within a specific context. For example, some studies have looked at the effect of increasing familiarity (e.g. Gauthier, Williams, Tarr & Tanaka, 1997, used three-dimensional shapes), or expertise (e.g. Diamond & Carey, 1986, used faces of people and dogs). Other studies have examined a particular process, such as mental rotation (e.g. Cooper and Podgorny, 1976, used two-dimensional geometric shapes), or the order of processing (e.g. Paquet, 1991, used hierarchical stimuli based on letters). Yet other studies have investigated specialist processing, often of faces, by comparing the representations of one set of objects with another (e.g. Donnelly & Davidoff, 1999, compared houses and chairs with faces). None of these sets of stimuli fulfilled the criteria for the purposes of the research presented here.

First, half the stimuli had to be meaningful to the participants. This meant that random geometric shapes would not be adequate; real objects would have to be used as the basis for the stimulus sets. The choice of which objects to select was also important. Many previous studies had used, appropriately for their research, small sets of stimuli from the point of view of the number of different objects included in the set. This allows many properties of the object to be controlled; properties such as size, general shape, features within the object, and so forth. However, such stimulus sets also impose limitations in the interpretation of the results. For example, if all the objects come from the same category, the homogeneity of the stimuli may be more influential in determining the representational format than the meaning of the stimuli. One way to avoid such issues is to extend the number of objects within the stimulus group, and represent several different categories within the set. This means that, as far is possible, we are not confounding the meaning with other, equally important, stimulus factors.

Second, as well as a set of meaningful stimuli, a set of meaningless stimuli had to be produced. This set had to be as equivalent as possible to the meaningful stimuli in every respect, except for the meaning. Previous studies have tended to re-arrange the parts of an object to produce a meaningless, or "scrambled", version. For example, several studies have taken faces, and re-arranged them to make "scrambled" faces, in which the configuration is disrupted (e.g. Davidoff & Donnelly, 1993; Tanaka & Farah, 1993). In these faces, eyes, nose, and mouth are still present, and clearly identifiable, but are put in an unfamiliar configuration. This often leaves the face still identifiable as a type of face, by recognising the features as facial components, but the face will not be a normal example of a face. Although this method is adequate for disrupting the normal configuration of an object, it is not sufficient to render the object entirely meaningless. For assessing the difference between meaningful and meaningless stimuli, a scrambled stimulus that can be recognised as an atypical example from its parts is not satisfactory. Therefore, not only did the parts of the meaningful objects have to be re-arranged; but also they had to be re-arranged in such a way that the parts themselves could not be recognised.

One issue with any stimulus set for the type of questions asked here, is that it is not certain exactly which part of the stimulus a person is using to complete the task. Some studies have used clear, meaningful parts within the task; for example, the nose has been used as a facial feature. This was seen as a problem within the context of the studies reported here. The meaningless objects were designed using the same parts as the meaningful ones, but not as functional features within themselves. It could be argued that this characteristic is what makes a meaningful object meaningful. However, deciding this *a priori* could result in the argument becoming circular. The studies here have been designed to use features that are not necessarily functional within themselves. This means that any difference between meaningful and meaningless objects have to be due to other factors within the representations. It also means that if there are no differences found between the two types of object, then this finding is inconclusive. It could be that there are no differences, but it could also be that the important differences involve the functional aspects of the parts of the object.

The third criterion was that, if complexity were to be used as a methodological tool, there had to be some control over the degree of complexity in the stimuli. The objects have been presented as two-dimensional line drawings, constructed using squares and right-angled triangles. Although this simplification results in the loss of some information, for example, shading, it allows the complexity of the objects to be measured and controlled. This level of control would have been far more difficult using three-dimensional drawings or photographs. It also allows meaningless geometric shapes to be formed from the meaningful drawings. The geometric figures have the same physical properties as the real figure, but contain no recognisable features. Again, the control of this factor would have been more problematical had stimuli that are more life-like been used. Although drawings tend to focus more on external details than internal ones, the drawings used here do incorporate internal elements, and this is equivalent across meaningful and meaningless stimuli.

Line drawings have been found to reduce configural processing in faces compared with photographs (Leder, 1996). Whether the same would be true of non-face objects is not certain. However, line drawings have been chosen here because of two main reasons. First, they allow a higher level of control in terms of complexity, arrangement of parts, and selection of parts for the search tasks. Second, it would be very difficult to produce convincing photographs or even three-dimensional drawings of the meaningless objects.

Production of the Stimuli

Twelve meaningful objects were chosen from pre-existing patterns available for artwork (Verso, 1988). They were classified into simple and complex objects, based on the number of lines in the object. The objects were matched across simple and complex stimuli, from various semantic categories. The final set was as follows: cat, iron, key, bell, ship, house, (simple); and hedgehog, scissors, water well, anchor, helicopter, aeroplane, (complex).

Two-dimensional line drawings of these objects were produced. All the drawings were formed from squares and right-angled triangles (half a square), with each square measuring 0.5 cm by 0.5 cm. The size of the drawings was kept comparable. The meaningless objects were then made from the meaningful objects, keeping the two sets equivalent. This was achieved by taking each meaningful object and re-arranging the component parts into a geometrical object (see Figure 2.01 for examples; see appendix 1 for a complete set of stimuli).

This is, in effect, an extreme form of scrambling. Although the basic parts were maintained across meaningful and meaningless objects, identifiable features were not. For example, the basic shape of the rotor blade is preserved in the meaningless version of the helicopter, but it is not presented as a rotor blade as such. Similarly, the funnels of the ship are still present in the meaningless version, but not as funnels. This was carried out to avoid the objects still being recognisable through the features themselves. Had the object been recognised on the basis of the features, but not the arrangement of the features, it could be that the factor being manipulated was some characteristic such as configuration rather than meaning *per se*.

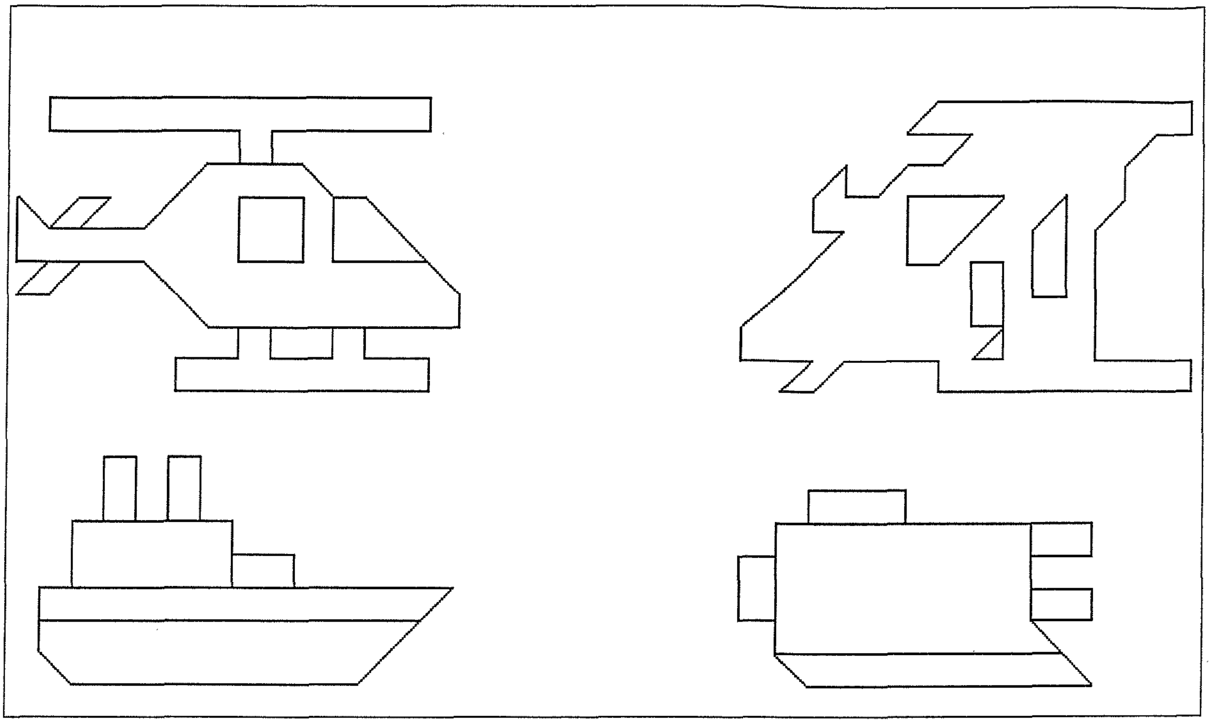

Figure 2.01: Examples of stimuli.

The top two objects are examples of complex stimuli, and the bottom two objects are examples of simple stimuli. The two objects on the left are meaningful; the two objects on the right are the meaningless versions of these stimuli. The same basic parts from the meaningful objects have been re-arranged into a meaningless object.

ASSESSMENT OF MEANINGFUL AND MEANINGLESS QUALITIES OF THE STIMULI

Method

Participants

The participants assessing the meaningfulness of the stimuli were 14 undergraduate students from the Psychology Department at the University of Southampton. There were 3 males and 11 females, aged between 19 and 23 (mean age 20.3). They were drawn from the same participant pool as those for the main experiments.

Materials

Each of the 24 objects was printed separately onto an overhead transparency, and labelled with a code number. A response sheet was given to the participants, enabling them to record a name for the object shown on the transparency, along with a confidence rating for the name, a familiarity rating, and a representativeness rating. At the top of the response sheet were the instructions and definitions for the task. Below the instructions were five columns. The first column contained the code number, in the order the stimuli were to be presented. The second column was blank for the participant to record a name for the object next to the code. The third to fifth columns contained the title of the decision they were to make (e.g. "familiarity"), and in each row a scale from 0 to 10 was drawn, for the participant to circle the appropriate score. All the ratings were on a scale of 0 to 10, and in each case 0 meant "not at all", 5 meant "unsure" and 10 meant "certain". The first decision was for the participant to rate how confident they were that the name they had given applied to the object shown. Familiarity was defined for the participants as "...the feeling that the object represented in the picture has been experienced previously, regardless of the context or form of that experience" (from Rugg, Schloerscheidt, & Mark, 1998). Representativeness was defined as "...the picture is a good and typical representation of the object...". These questions were designed to elicit whether the participants found the pictures meaningful; in this context meaningful was defined by whether they recognised clearly and easily what the drawing was depicting, and whether the depicted object had meaning for them, based on past experience.

Procedure

The participants were tested simultaneously. They were seated in front of an overhead projector, and given a sheet of paper that contained the instructions, and a labelled spreadsheet to record their results. The instructions told them that they would have the opportunity to look at a number of pictures. They would be required to make four decisions about each picture. The first decision was to decide what object the picture represented and to name it. The participants were then asked to rate how confident they were about their first decision; how familiar the object was; and how representative the picture was of the object.

The order of presentation of the stimuli was random, with the constraint that no meaningless stimulus was presented within three presentations of its meaningful counterpart. This was to prevent the participants from associating the two stimuli.

The experimenter, who was unaware of what the stimuli may or may not represent, presented each stimulus manually, leaving the stimulus on the screen for 30 seconds. It was then removed and, after a further 15 seconds, replaced with the next one.

Results

The ratings of confidence, familiarity, and representativeness were tabulated and averaged across the stimuli. Suitable results for the objects to be defined as meaningful would be a high confidence in naming the object and high ratings for familiarity and for representativeness. "High", in this instance, would be an average value of at least 7 on the 0-10 scale. Suitable results for the objects to be defined as meaningless would be low scores on confidence, familiarity, and representativeness. Here, each rating should receive an average value of less than 3 on the 0-10 scale.

Table 1: Mean values for confidence, familiarity and representativeness ratings of the stimuli

		meaningful			meaningless		
		conf	fam	rep	conf	fam	rep
complex stimuli	helicopter	9.93	8.07	9.14	2.07	2.29	1.36
	scissors	8.86	7.93	8.71	2.86	2.29	2.50
	aeroplane	9.50	8.14	8.71	2.07	2.93	1.50
	hedgehog	9.86	8.57	9.36	1.54	3.54	1.85
	anchor	6.07	4.93	5.07	1.86	2.86	1.64
	water well	9.34	8.00	8.57	2.34	2.93	1.93
	mean	9.15	7.61	8.70	2.03	2.73	1.42
	S.D.	0.82	0.78	0.93	1.20	1.08	0.91
simple stimuli	ship	9.79	8.00	9.57	0.50	1.86	0.36
	iron	9.29	8.71	9.50	1.15	3.54	0.54
	cat	8.71	6.64	7.71	2.64	3.86	2.07
	house	9.79	7.86	9.29	2.57	3.36	1.71
	bell	7.71	6.78	7.43	3.79	2.71	2.71
	key	9.64	7.64	8.71	1.50	1.07	1.14
	mean	8.93	7.61	8.26	2.12	2.81	1.80
	S.D.	1.45	1.33	1.59	0.45	0.47	0.40
overall	mean	9.04	7.61	8.48	2.07	2.77	1.61
overall	S.D.	1.13	1.04	1.26	0.86	0.80	0.70

The columns signify: Conf = confidence in naming; fam = whether the stimuli has been encountered before, in whatever format; and rep: how representative this version is of the object it is depicting.

Overall, the meaningful objects gave mean values of 9.04 for confidence, 7.61 for familiarity, and 8.48 for representativeness (see appendix 2 for comments). This was significantly different from the values obtained for the meaningless objects, which were 2.07 for confidence (t_{22} = 16.97, $p<.001$), 2.77 for familiarity (t_{22} = 12.97 $p<.001$), and 1.61 for representativeness (t_{22} = 16.47 $p<.001$). This suggests that the meaningfulness differed significantly between the two stimulus sets.

Furthermore, the difference between the meaningfulness of the stimulus sets was also significant when only the complex stimuli (confidence t_{10} = 12.05, $p<.001$; familiarity t_{10} = 8.95, $p<.001$; representativeness t_{10} =

13.69, $p<.001$), or only the simple stimuli (confidence $t_{10} = 10.97$, $p<.001$; familiarity $t_{10} = 8.33$, $p<.001$; representativeness $t_{10} = 9.64$, $p<.001$) were compared.

Finally, the data demonstrated that within each stimulus set there were no significant differences in meaningfulness between the simple and complex stimuli. This applied to both the meaningful stimulus set (confidence $t_{10} = 0.34$, $p>.1$, familiarity $t_{10} <0.01$, $p>.1$, representativeness $t_{10} = 0.59$, $p>.1$), and the meaningless stimulus set (confidence $t_{10} = 0.19$, $p>.1$, familiarity $t_{10} = 0.15$, $p>.1$, representativeness $t_{10} = 0.92$, $p>.1$).

Therefore, there was a significant difference in the perception of meaningfulness between the meaningful and the meaningless stimulus sets. This difference applied to both simple and complex stimuli, as well as to the stimuli as a whole. In addition, there was no significant difference in the meaningfulness between the simple and the complex stimuli.

ASSESSMENT OF SIMPLE AND COMPLEX QUALITIES OF THE STIMULI

Method

Several criteria for complexity have been used previously. One such measurement was a compactness value. This was calculated by taking the square root of the area, and then dividing this by the number of edges (or angles) in the perimeter (from Attneave, 1957; see also Cooper & Podgorny, 1976; Podgorny & Shepard, 1983). The lower the compactness value, the more complex the object is. However, the objects used in this thesis also had lines and features within the shape, as well as in the perimeter. Therefore, the number of internal lines was added to the number of lines in the perimeter. This figure then replaced the number of edges to give the measure of compactness used to assess the stimuli. The compactness was therefore calculated as follows: (√number of 0.5 cm squares) / number of lines in the stimulus.

Analyses

The compactness values were tabulated and averaged across the stimuli (see appendix 3). Simple stimuli give a higher compactness value than complex stimuli. Overall, there was a significant difference in compactness between the complex objects, mean value 0.20, and the simple objects, mean value 0.46 ($t_{11.64} = 4.86$, $p<.001$). This difference remained when the meaningful and meaningless objects were considered separately. The mean compactness value for the meaningful complex objects was 0.19, and 0.45 for the simple stimuli, ($t_{10} = 3.76$, $p <.05$). The compactness values for the meaningless stimuli were also significantly different; the complex stimuli had a mean value of 0.21, and the simple stimuli a mean value of 0.47 ($t_{10} = 2.95$, $p<.05$). The compactness values for the meaningful and meaningless stimuli were not significantly different from one another, for either complex stimuli ($t_{10} = 1.35$, $p>.05$), or simple stimuli ($t_{10} = 0.20$, $p>.05$).

This reflects that the complex objects are significantly more complex than the simple objects; furthermore, the levels of complexity between meaningful and meaningless objects are equivalent.

These assessments had shown that there were significant differences between the levels of meaningfulness. There were also significant differences between the levels of complexity. Finally, these differences were orthogonal to each other. Hence, the stimuli were appropriate to use in the experiments that follow.

CHAPTER 3

THE EFFECT OF MEANING ON PERFORMANCE IN A MENTAL ROTATION TASK

ABSTRACT

The aim of the study was to examine whether there was a difference in the representations of meaningful and meaningless objects, by comparing the pattern of errors produced by the stimuli during a visual mental rotation task. Three sources of error could be specified. The first source was the angle of rotation. More errors would occur with rotation through greater angular distances. This source of error would affect both holistic and piecemeal rotation processes. The second source of error was the complexity of the object, and the resultant representation. Complex representations, with more information, would lead to higher errors than simple representations. The third source of error was the reconstruction of the shape from its rotated parts. A holistic process would not be susceptible to this error source, as the whole representation would have been rotated as one unit. Therefore, there should be no interaction between the complexity of the stimulus and the angle of rotation. If a piecemeal process was used, with more parts in a complex stimulus than a simple stimulus, then there should have been an interaction between the complexity of the stimulus and the angle of rotation. Meaningful objects showed an interaction between the complexity and angle of rotation. In contrast, meaningless objects showed no such interaction. This suggests that the meaningful stimuli were rotated by a piecemeal process, whereas the meaningless objects were rotated holistically.

A journal article based on this chapter has been submitted to *Psychonomic Bulletin and Review* under the title “The Role of Meaning and Familiarity in Mental Transformations”.

CHAPTER 3

INTRODUCTION

The mental rotation of a representation shares similarities with the physical rotation of an object. One important commonality is that the time needed to complete the task increases as the degree of rotation becomes greater. The linear slope of increased response time as a function of angle of rotation in mental imagery (Shepard & Metzler, 1971) suggests that, like physical rotation, getting from one point to another involves passing through the intervening positions (Morgan, 1983). Mental rotation can be understood, within this context, as a series of transformations of mental representations. One issue in the area pertains to whether the mental transformations are holistic or piecemeal processes. In a holistic process, the entire representation is rotated as a single unit. In contrast, in a piecemeal process, the representation is rotated part by part; the "whole" is re-established through the spatial relations among the parts. It is assumed in this study that the processing will be determined by the underlying representation (Marr, 1982). A holistic representation will lead to holistic processing, and a segment-based representation will lead to piecemeal processing. If the type of processing during a mental transformation can be determined, then the underlying representation can be inferred.

The main approach of previous research has been to examine whether response time data show an interaction between the complexity of the representation and the angle of rotation. Cooper and Podgorny (1976) introduced the logic behind this approach. They suggested that the rate of rotation of the representation of an object would depend upon the number of parts that had to be rotated; the more parts, the longer the rotation would take to complete. They argued that if a representation is transformed holistically, then the entire representation would be rotated as a single unit, and the complexity of the representation should not, therefore, affect the rate of rotation. Conversely, if a representation is transformed through a piecemeal process, then complex representations, with more parts to rotate and re-establish, will show a different rate of rotation from simple representations (see Cooper & Podgorny, 1976). If reaction times are plotted against the angle of rotation, the difference will be demonstrated in the slope of the line. Holistic processing will produce parallel lines between stimuli of varying complexities (or, indeed, a single line if other factors, such as ease of encoding, are also equivalent across complexities), whereas piecemeal processing will produce steeper lines with increasing complexities. Using polygons of varying complexity Cooper and Podgorny found that the rate of rotation was not affected by the complexity of the representations, and concluded that a holistic process was involved. However, other researchers (e.g. Folk & Luce, 1987; Yuille & Steiger, 1982) have found steeper slopes for rotating complex representations relative to simple representations. These studies concluded that the rotation process was piecemeal rather than holistic.

The second approach for investigating the nature of the rotation process was to use hierarchically structured stimuli of large letters composed of smaller letters (Robertson & Palmer, 1983). The researchers assumed that if holistic processing was used, then the rotation would be performed using, predominantly, global information (the larger letters). If piecemeal processing was used, the predominant information guiding the rotation would be local (the smaller letters). Robertson and Palmer found an advantage in response times when the task was based on the larger letters, and concluded that rotation was a holistic process. However, this approach has also been inconclusive. Paquet (1991), using similar tasks and stimuli, concluded that a piecemeal rotation process could explain the data equally well.

The third approach for exploring the nature of mental rotation was to compare the rotation rates of possible and impossible objects (Dror, Ivey, & Rogus, 1997). Impossible objects appear more difficult to encode holistically than possible objects (Schacter, Cooper, & Delaney, 1990; Schacter, Cooper, Delaney, Peterson, & Tharan, 1991). Dror *et al* reasoned, therefore, that impossible objects were more likely to be transformed using a piecemeal process than a holistic one. Following this logic, if the rate of rotation of impossible and possible objects was comparable, then this would suggest that possible objects were rotated in a piecemeal fashion, just like the impossible objects. Alternatively, if the rate of rotation were different, then this would suggest that the possible objects were rotated holistically. Dror *et al* found that the rates of rotation of the possible objects were identical to those of the impossible objects, and concluded that a piecemeal process underlined image rotation. However, the results of Dror *et al* only demonstrated that possible and impossible objects were rotated using the same process; the conclusion that this process was piecemeal rests on the assumption that the impossible objects were not represented and rotated holistically.

To summarise, numerous studies have examined the nature of mental rotation and have produced inconclusive, and even conflicting, results. One way to interpret these contradictions is to assume that both holistic and piecemeal transformations are possible. Other studies have suggested that the type of rotation may be dependent on particular circumstances (Cochran, Pick, & Pick Jr, 1983). For example, whether the transformation is holistic or piecemeal may be influenced by the degree of familiarity (Bethell-Fox & Shepard, 1988), or the amount of practice with the task (Voyer, 1995). Both familiarity and practice appear to encourage holistic processing; an “expertise effect” can be accompanied by a move toward holistic representation and processing. For example, Bethell-Fox and Shepard (1988) found that unfamiliar geometric shapes produced piecemeal mental rotation; when the same shapes were made familiar, holistic processing occurred. Larson (1985) has also produced findings consistent with this interpretation.

If the rotation process is sensitive to the difference in the underlying representations, then it can serve as a tool to distinguish any differences between representations of meaningful and meaningless objects. This study aims to use the nature of the mental rotation process to examine the type of mental representations formed of the stimuli. Comparing performances with meaningful and meaningless stimuli can highlight any differences associated with the underlying meaning of the object.

Past research in mental rotation has almost exclusively examined response times. As described earlier, holistic representations would not entail greater response times with increased complexity, but part-based representations would. However, the pattern of results that can be found is open to alternative interpretations. For example, if the parts in a segment-based representation are processed in parallel, the resulting response times may not increase with the complexity of the objects, and thus will give a pattern of results that is identical to that of a holistic representation.

This study circumvents this difficulty by using error rates as the dependent measure instead of response times. Holistic and piecemeal rotation processes predict distinct patterns of errors, which are very similar to those found with response times. Three sources of error can be distinguished in mental rotation. The first source is the angle of rotation. As the representation is rotated greater distances, more processing is required, and more errors are likely to occur as processing demands increase. This source of error will affect both the holistic and piecemeal rotation processes. Both will predict a main effect of angle of rotation, in which errors will increase as the angle of rotation increases.

The second source of error is based in the complexity of the object that is to be represented. In both types of processes, an effect of complexity can be predicted, because there is more information in a complex representation than in a simple one. This effect could be due to increased difficulty in encoding, in retrieval, or in the comparison of a complex representation relative to a simple one. However, the errors will not be dependent upon the angle; the consistent difference across all the angles will not predict an interaction between complexity and angle.

The third source of error varies between holistic and piecemeal processing, because the variation is directly dependent upon the number of segments in the representation. When a representation is rotated, the spatial arrangement of the parts has to be maintained through the process if errors are to be avoided. If the rotation is a piecemeal process, it involves moving the different segments of the representation, one by one, and then realigning them to maintain the proper spatial relations among them. The errors in this type of movement are dependent on the number of segments in the representation; the more parts there are to be moved, the more likelihood there is of an error occurring (see also Rock, 1974). More errors will be expected to occur, therefore, in representations of complex objects, with a larger number of segments, than in the representations of simpler objects. Furthermore, this type of error will also be dependent upon the angle of rotation. The further a part has to be transformed before it can be re-aligned, the greater the opportunity is for misalignment. Therefore, in a piecemeal process, an interaction between the angle of rotation and the complexity is predicted. A holistic process, maintaining the spatial arrangement by rotating the representation as one unit, will not show an effect of this source of error.

Combining these three sources of error leads to a prediction that in a piecemeal process, where all the sources are operating, there will be an interaction between the angle of rotation and the complexity of the object. In a holistic process there will be no interaction between angle of rotation and the complexity of the object. In contrast with the first and second sources of error, the third source of error provides a way of distinguishing between parallel and serial piecemeal processes that response time does not. More errors can be expected with an increased number of segments regardless of whether those segments are processed in parallel or not.

In summary, the aim of the study was to examine whether there was a difference in the representations of meaningful and meaningless objects. This was to be achieved by comparing the pattern of errors produced by the two types of objects during a mental rotation task. If a holistic process was used, then there should be no interaction between the complexity of the stimulus and the angle of rotation. If a piecemeal process was used, then there should be an interaction between the complexity of the stimulus and the angle of rotation.

CHAPTER 3

EXPERIMENT 1: ROTATION OF MEANINGFUL AND MEANINGLESS STIMULI

Method

Design.

This experiment used a mixed design. Complexity of the stimuli and the angle of rotation were within-subject factors, and the type of stimulus was the between-subjects factor. There were two levels of complexity (simple and complex), and four angles through which the stimuli were rotated (0°, 50°, 100°, and 150°). There were two groups of participants, who each performed the same rotation task using a different set of stimuli. The first group performed the task using simple and complex meaningful stimuli, and the second group performed the task using simple and complex meaningless stimuli.

Participants.

The participants were 56 undergraduate students from the Psychology Department participant pool at the University of Southampton. In the group rotating meaningful stimuli, there were 8 males and 20 females, with an age range of 18 to 33 (mean age 20.9). In the group rotating the meaningless stimuli, there were 8 males and 20 females, with an age range of 18 to 49 (mean age 22.5). All the participants had normal or corrected-to-normal vision, and were right-handed. None of the participants had any experience with the stimuli or task prior to the experiment.

Materials.

Four stimulus sets were used. There were two meaningful sets and two meaningless sets; each set contained either six simple stimuli or six complex stimuli. Four additional stimuli, one for each set, were made for practice purposes and were not included in the main experiment.

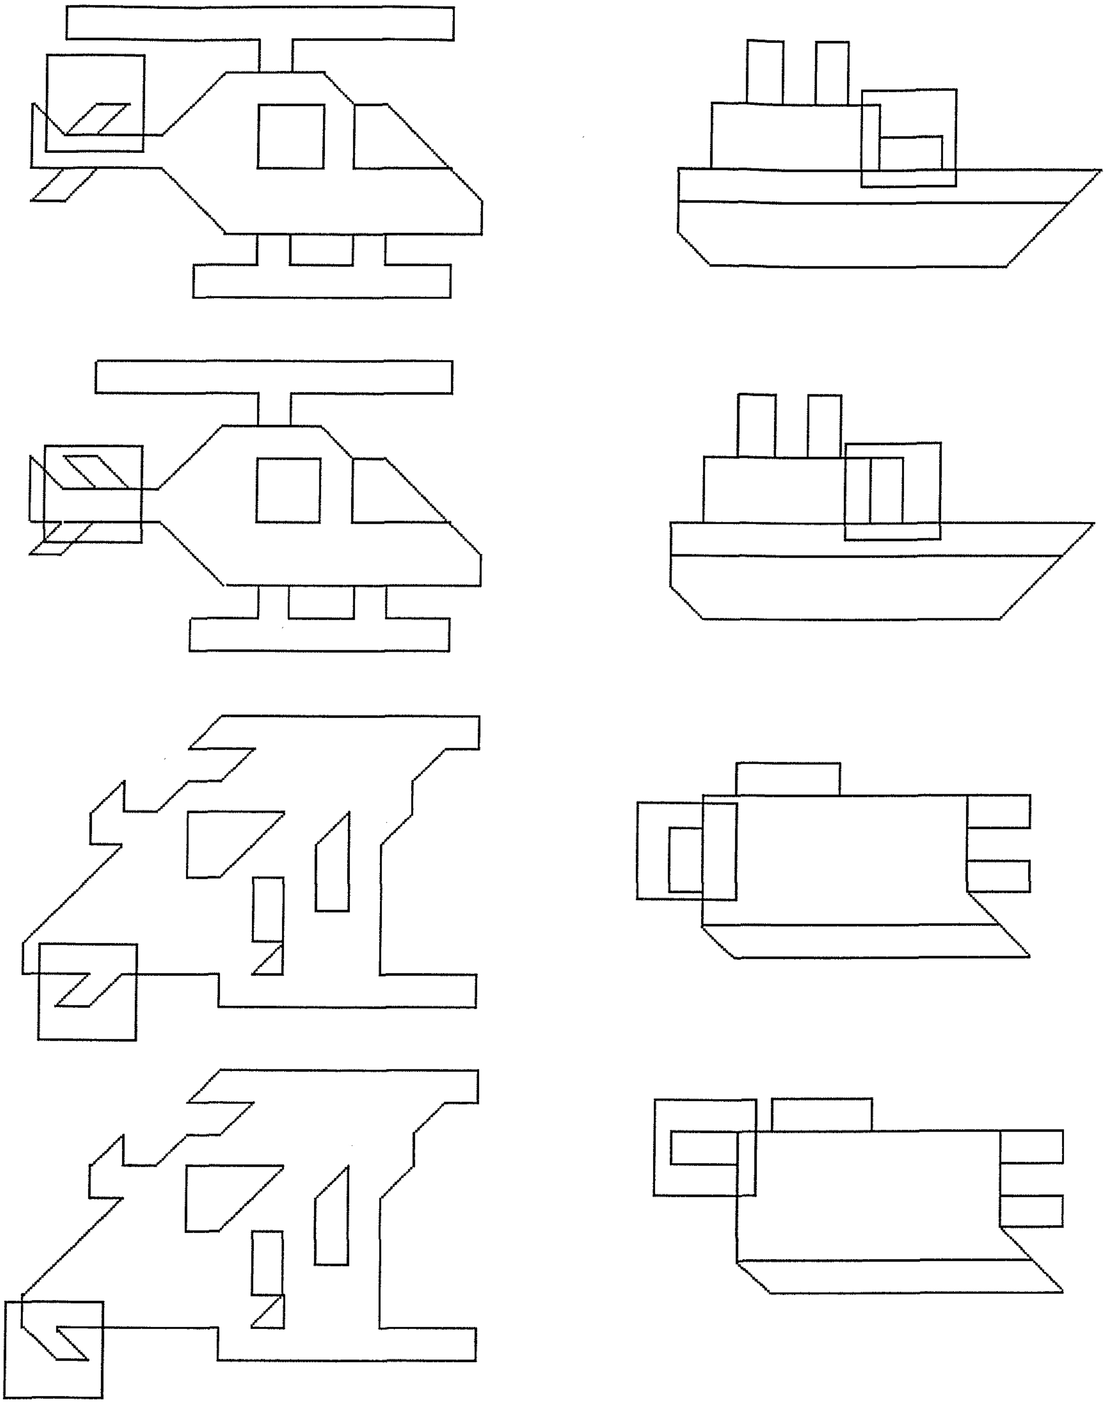

Figure 3.01: Examples of the Modified Stimuli

Examples of complex stimuli are shown on the left, and simple stimuli on the right. The top four stimuli are the meaningful stimuli, and the bottom four stimuli are meaningless. Each unmodified stimulus is shown with an example of a modified stimulus directly beneath it. The original part and the modifications are outlined.

First, the upright objects were produced, as outlined in Chapter Two. The next stage was to produce modifications for the stimuli, to be used for the "different" trials. A modification consisted of moving a small feature of the stimulus to an adjacent location. This meant that although the modified stimulus was different from the original stimulus, the total area and complexity of the two stimuli remained the same (see Figure 3.01 for examples, and appendix 4 for the full set of modified stimuli). For each stimulus four modifications were made, each at a different part of the object. The use of four modifications was to prevent the participants from anticipating either the type or location of the modifications.

Modifications of the meaningless stimuli were then produced in the same way as for the meaningful stimuli, and matched to those of the meaningful stimuli. Particularly salient areas of the meaningful stimuli were not modified. Finally, each of the original stimuli and their four modifications were rotated clockwise in the picture plane to produce stimuli at 0°, 50°, 100°, and 150°. Rotations of 180° or larger were not selected, to avoid ambiguity over whether to rotate clockwise or anti-clockwise, and to discourage mental rotation in the third dimension, outside the picture plane (Murray, 1997).

Two tasks were constructed; one contained only the simple stimuli and the other contained only the complex stimuli. Each task consisted of 192 trials: six stimuli were presented eight times at the four different orientations. The eight different presentations of each stimulus at each orientation included four "same" trials, using unmodified stimuli, and four "different" trials, using one of the four modified stimuli. The trials were organised into eight blocks of 24 trials. Each block included four presentations of each stimulus, once at each of the four angles of rotation. Two of these presentations were "same" trials and two were "different" trials, giving 12 "same" and 12 "different" trials within each block. The trials in a block were organised so that the number of trials in each orientation was equal, and that half the trials were "same" and half were "different". The order of the trials within a block was randomised.

The stimuli were presented using a Power Macintosh 8600/250 computer with a 17-inch Apple Vision monitor (running at 1024 x 768 resolution). The experiment was designed and administered using the commercial software Superlab, version 1.4 (Cedrus Corporation, 1991).

Procedure.

Each participant performed the two experimental tasks using the type of stimulus designated to their group. The tasks were identical in all respects (e.g., procedure, trials, counterbalancing, and so forth) except that they used a different set of stimuli, either complex or simple. Half the participants performed the task with the simple stimuli first, and then, after a short break, performed the task with the complex stimuli. The other half of the participants performed the two tasks in the opposite order.

In each trial the word “ready” was presented on the screen, and remained until the participant pressed the spacebar with their non-dominant hand. Following the spacebar press, a blank screen was presented for 500 ms, and then an upright version of a non-modified object appeared on the screen. Participants were instructed to study this representation carefully and only to continue when they were ready; at this point they should press the spacebar. When the participants pressed the space bar a mask appeared for 500 ms; the mask consisted of horizontal, vertical, and diagonal straight lines that completely covered the area occupied by the first representation. The mask was followed by the second object. This object was either the same as the first object, or one of the modifications of it, and was presented at one of the four angles. It remained on the screen for 750 ms. After this, the words “same or different?” appeared on the screen to prompt a response. The participant was instructed to respond immediately that they saw this screen. If the participant thought the two representations were identical, except for orientation, they were instructed to respond by pressing the key labelled "same" (the “b” key on a standard keyboard). If they thought the two representations were different from one another, they were to press the key labelled "diff" (the "n" key). The key presses were made using the forefinger or middle finger of their dominant hand. Following their response, a blank screen appeared for 750 ms, then a new trial began with the word "Ready?" (see Figure 3.02).

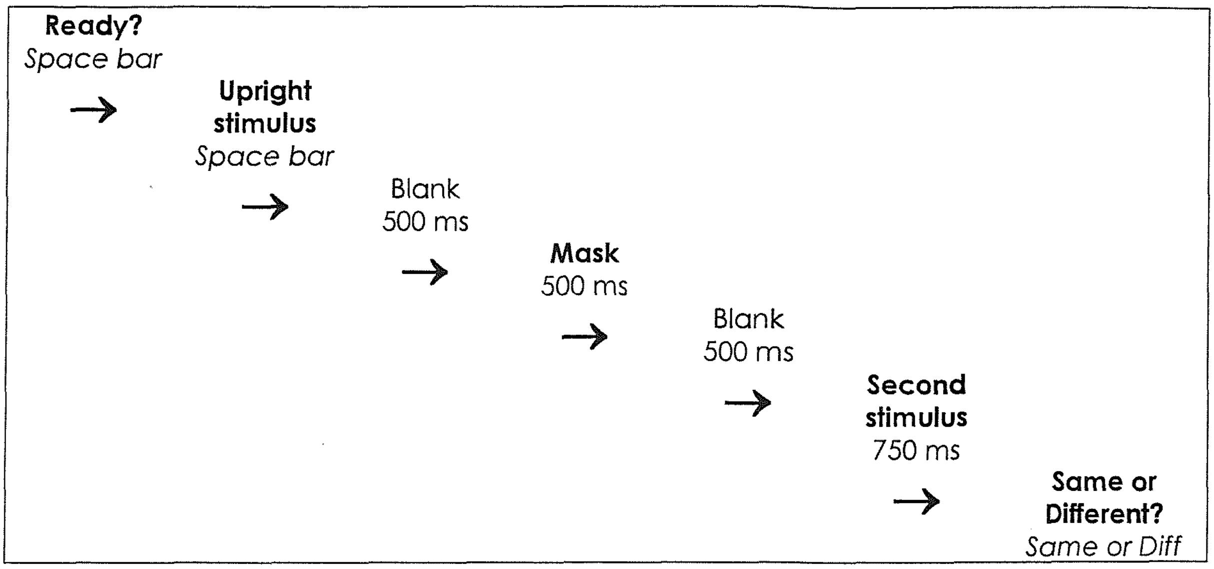

Figure 3.02: Events within a Trial of the Experiment.

The events visible to the participants are presented in bold. The first stimulus was an unmodified, upright version. The second stimulus could either be unmodified, or have any one of the four modifications, and be presented at any of the four angles. The items under each event are either the length of presentation or the response required from the participant.

Once the first experimental task had been completed, with either the simple or complex stimuli, the participants were encouraged to take a short break before continuing. When they were ready to go on to the next task a brief reminder of the instructions was presented, and then the trials proceeded exactly as in the previous task, except that the second set of stimuli was used.

Prior to performing the experimental task, the participants completed a set of practice trials using objects that were not included in the actual experimental task. First, they were presented with comprehensive instructions on the computer screen, including examples of what constituted "same" and what constituted "different" stimuli. Following the instructions, participants were given a chance to familiarise themselves with the response keys that they would need to use during the experimental task. The word “same” or “different” was presented on the screen, and participants had to press the appropriate key immediately. Feedback was provided; if the incorrect response key was pressed, a “beep” sounded. The words “same“ and “different“ were randomly presented 32 times (16 "same" and 16 "different”).

Next, using objects that were not included in the main task, they performed 64 trials on only upright (0°) stimuli. In each of the 64 trials an object was presented in the upright position, followed by a second object, also in an upright position. The second object was either the same (on half the trials), or a modification of the first object (on the other half of the trials). The sequence of events within each trial was identical to the trials in the main experimental task. The participant responded "same" or "diff", and was given feedback saying whether the response was correct or incorrect. Finally, 16 trials were given in which the rotation manipulation was introduced; the second representation was presented at any of the four possible orientations. Again, feedback was given following the response. The gradual practice sequence allowed the participants to become familiar and comfortable with the experimental task and procedure, and to gain a thorough understanding of what they needed to do, without seeing any of the stimuli they would be tested on in the main task. Once the practice tasks were completed, participants began the experimental tasks. No feedback was provided during the main experiment.

Results

General Analysis

For the main analysis, the percentage error rates (see table 3.1) were converted into d' values. The participants may not have been using identical strategies and processes in the same and the different trials, and using d' values as opposed to percentage errors would avoid conflating two potentially different sets of errors. A high d' value is generally associated with a low proportion of errors.

Table 3.1: Error Rates for the Meaningful and Meaningless Stimuli

Angle of	Rotation	0°	50°	100°	150°
Meaningful stimuli	simple	Mean 21.36	Mean 25.18	Mean 26.71	Mean 26.71
		S.D. 8.61	S.D. 7.93	S.D. 6.60	S.D. 7.88
	complex	Mean 20.63	Mean 28.73	Mean 30.81	Mean 32.38
		S.D. 9.47	S.D. 8.09	S.D. 8.08	S.D. 8.58
Meaningless stimuli	simple	Mean 21.71	Mean 27.44	Mean 31.11	Mean 29.76
		S.D. 7.78	S.D. 7.26	S.D. 9.13	S.D. 8.06
	complex	Mean 28.05	Mean 33.73	Mean 36.90	Mean 37.27
		S.D. 10.24	S.D. 7.20	S.D. 7.79	S.D. 8.43

The percentage errors for meaningful and meaningless stimuli are given for each angle of rotation, across simple and complex stimuli. The mean values and standard deviations are shown.

The d' values were subjected to an ANOVA, with complexity (simple and complex) and angle (0°, 50°, 100°, and 150°) as within-subject variables, and type of stimulus (meaningful and meaningless) as a between-subject variable. There was a significant effect of complexity ($F_{1, 54} = 36.03$, $p<.001$), a significant effect of angle ($F_{3, 162} = 71.51$, $p<.001$), and a significant effect of type of stimulus ($F_{1, 54} = 4.66$, $p<.05$). There were also significant interactions between complexity and type of stimulus ($F_{1, 54} = 5.01$, $p< .05$), complexity and angle ($F_{3, 162} = 5.17$, $p<.01$), and a borderline effect between complexity, angle and type of stimulus ($F_{3, 162} = 2.45$, $p=.07$). These findings were explored in more detail by examining the findings of each group of participants separately. Of particular interest in each case was the pattern of errors, and whether there was an interaction between errors and angle of rotation.

Meaningful Stimuli

There was a significant effect of complexity ($F_{1,27} = 7.67$, $p < .05$), a significant effect of angle ($F_{3,81} = 21.54$, $p < .001$), and a significant interaction between complexity and angle ($F_{1,81} = 8.08$, $p < .001$; see Figure 3.03).

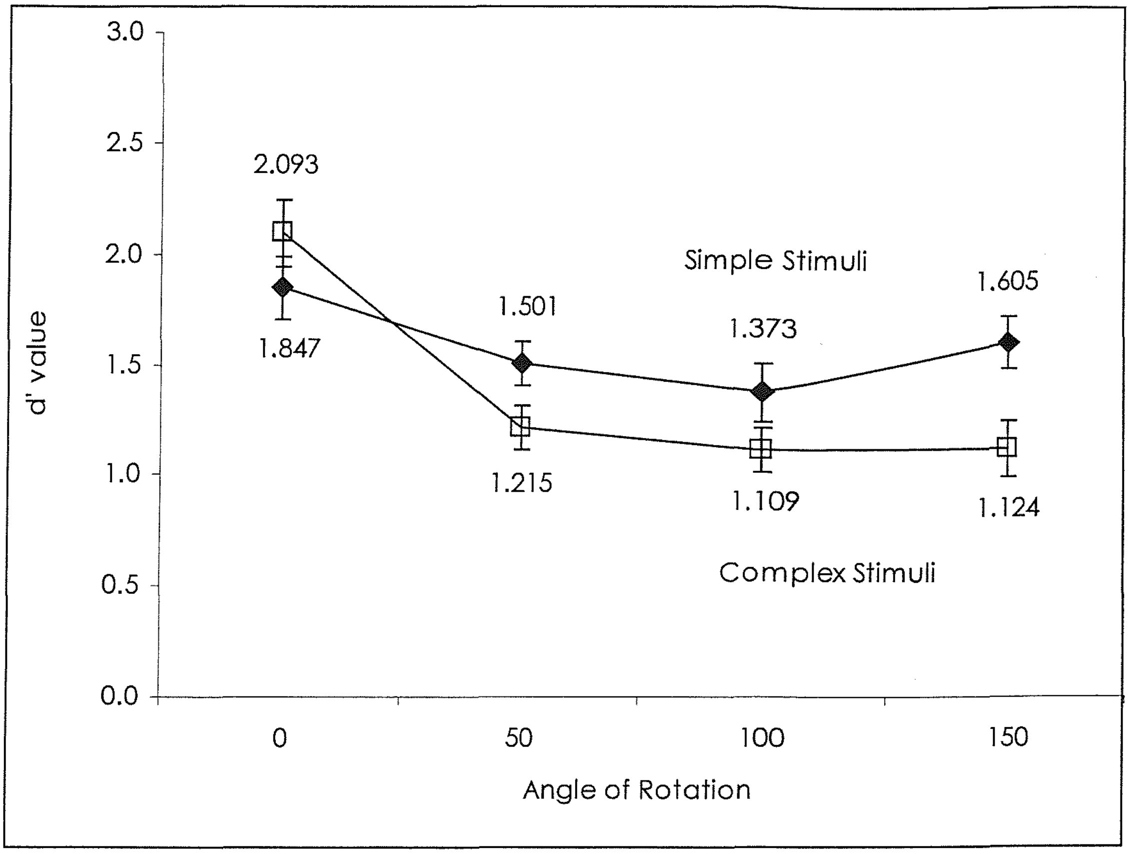

Figure 3.03: Mean d' values for the meaningful stimuli.

The simple stimuli are illustrated with closed diamonds, and the complex stimuli are illustrated with open squares. The simple stimuli demonstrate a higher d' than the complex stimuli for every degree of rotation except 0°. There is an interaction between angle of rotation and complexity. The mean value for each point is shown, along with standard error bars.

There were more errors made with complex stimuli than with simple stimuli, and hence the d' was higher for the simple stimuli than the complex stimuli. This was found for each angle of rotation, with the exception of 0°. Here, the number of errors was very similar, and the d' was actually higher for complex than for simple stimuli (0°: $F_{1,27} = 4.24$, $p<.05$; 50°: $F_{1,27} = 6.62$, $p<.05$; 100°: $F_{1,27} = 5.03$, $p<.05$; 150°: $F_{1,27} = 16.94$, $p<.001$).

Both complex stimuli ($F_{3,81} = 34.05$, $p<.001$) and simple stimuli ($F_{3,81} = 4.33$, $p<.01$) showed an effect of angle of rotation. However, the difference in d' from one angle to the next was only significant between 0°and 50° (simple: $t_{27} = 2.73$, $p<.05$; complex: $t_{27} = 6.63$, $p<.001$).

To check for a speed-accuracy trade-off, the response times were examined (see table 3.2). The response times mirrored the d' values, except for the results between simple and complex stimuli at 0°. There was a significant effect of complexity ($F_{1, 27} = 7.67$, $p<.05$), a significant effect of angle ($F_{3, 81} = 21.54$, $p<.001$), and no interaction between complexity and angle ($F< 1$). Longer response times were found with the complex stimuli than with the simple stimuli at every angle, including 0° (0°: $F_{1,27} = 8.65$, $p<.01$; 50°: $F_{1,27} = 5.99$, $p<.05$; 100°: $F_{1,27} = 4.31$, $p<.05$; 150°: $F_{1,27} = 9.58$, $p<.01$). The difference in response time from one angle to the next was only significant between 0° and 50° (simple: $t_{27} = 4.94$, $p <.001$; complex: $t_{27} = 4.50$, $p<.001$).

Table 3.2: Response times for the meaningful stimuli

	Angle of Orientation			
	0°	50°	100°	150°
Simple stimuli	Mean = 452.93	Mean = 563.46	Mean = 565.48	Mean = 568.25
	S.D. = 136.08	S.D. = 165.12	S.D. = 160.90	S.D. = 155.54
Complex stimuli	Mean = 553.84	Mean = 629.22	Mean = 638.02	Mean = 635.51
	S.D. = 161.42	S.D. = 174.73	S.D. = 187.01	S.D. = 184.22

The study times of the simple and complex stimuli were examined to ensure that the findings were not affected by variations in the study time across complexities. The length of the study time was comparable in the two conditions ($t_{27} = 1.0$, $p>.1$). To preclude the possibility that participants were guessing, a one-sample t-test was performed on the percentage errors in the most difficult condition, 150°. The errors were significantly below chance ($t_{27} = 15.64$, $p < .001$), reflecting that the responses were not based purely on guesses. There was no significant effect of the order in which the simple and complex stimuli had been used to perform the tasks ($F < 1$).

Meaningless Stimuli

There was a significant effect of complexity ($F_{1,27} = 31.52$, $p < .001$), a significant effect of angle ($F_{3,81} = 61.62$, $p < .001$), and no interaction between complexity and angle ($F < 1$; see Figure 3.04).

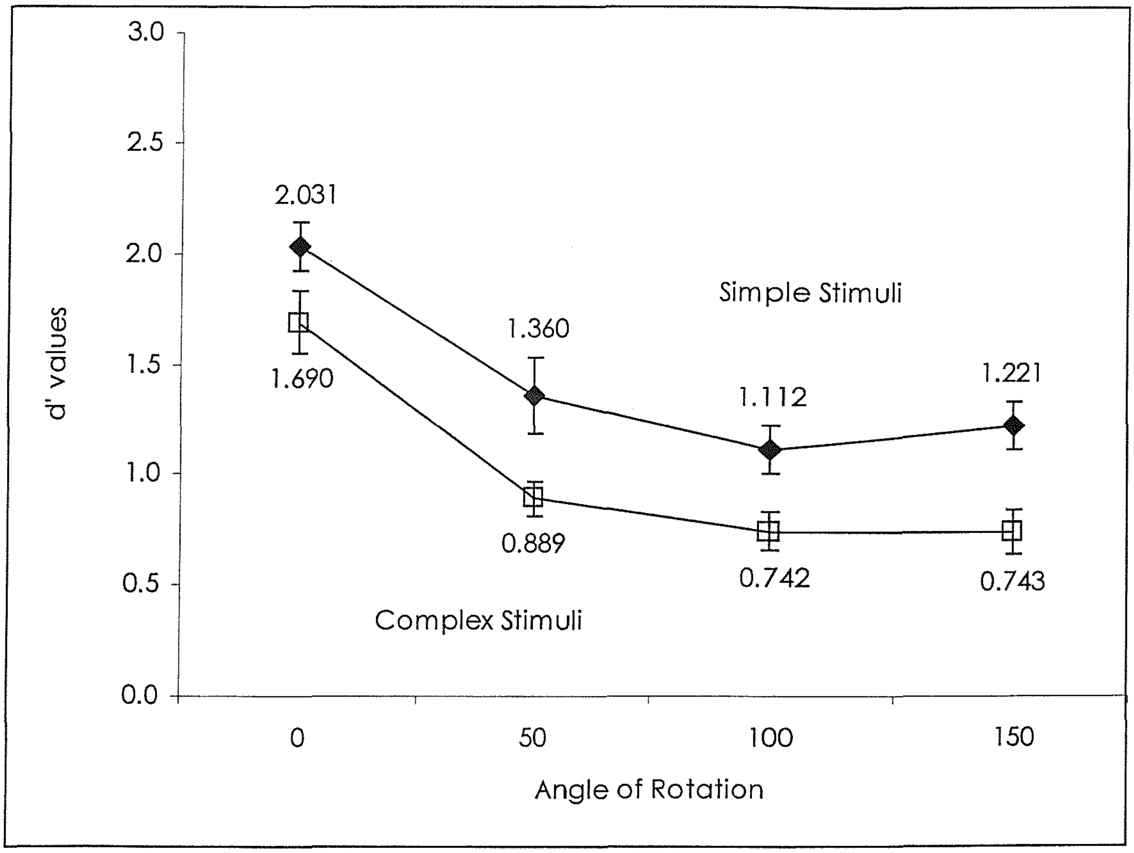

Figure 3.04: Mean d' values for the meaningless stimuli.

The simple stimuli are shown by closed diamonds, and the complex stimuli by open squares. The simple stimuli demonstrate a higher d' for every degree of rotation than the complex stimuli. There is no interaction between angle of rotation and complexity. The mean value for each point is shown, along with standard error bars.

There were more errors made with complex stimuli than with simple stimuli, with associated lower d' values. This was found for the three larger angles of rotation (0°: $F_{1,27} = 3.58$, p=.069; 50°: $F_{1,27} = 17.51$, p <.001; 10°: $F_{1,27} = 19.60$, p<.001; 150°: $F_{1,27} = 19.07$, p<.001).

Both complex stimuli ($F_{3,81} = 31.90$, p <.001) and simple stimuli ($F_{3,81} = 25.74$, p <.001) showed an effect of angle of rotation. The difference in d' from one angle to the next was significant between 0° and 50° (complex: $t_{27} = 6.55$, p<.001; simple: $t_{27} = 5.14$, p<.001), and borderline between 50° and 100° (complex: $t_{27} = 1.78$, p=.087; simple: t_{27}=2.02, p=.053).

To check for a speed-accuracy trade-off, the response times were examined. The response times mirrored the d' values. There was a significant effect of complexity ($F_{1,27} = 6.34$, p< .05), a significant effect of angle ($F_{3,81} = 23.21$, p< .001), and no interaction between complexity and angle ($F_{1,81} = 1.83$, p> .1). Longer response times were found with the complex stimuli than with the simple stimuli at every angle except 0°(0°: $F_{1,27} = 3.01$, p = .09; 50°: $F_{1,27} = 7.78$, p <.05; 100°: $F_{1,27} = 4.71$, p <.05; 150°: $F_{1,27} = 5.38$, p <.05). The difference in response time from one angle to the next was only significant between 0°and 50°(simple: $t_{27} = 3.92$, p <.001; complex: $t_{27} = 4.78$, p <.001).

Table 3.3: Response times for the meaningless stimuli

	Angle of Orientation			
	0°	50°	100°	150°
Simple stimuli	Mean = 454.13	Mean = 572.47	Mean = 578.92	Mean = 619.06
	S.D. = 154.08	S.D. = 194.21	S.D. = 180.80	S.D. = 196.54
Complex stimuli	Mean = 521.51	Mean = 721.65	Mean = 713.22	Mean = 731.32
	S.D. = 192.42	S.D. = 312.40	S.D. = 328.51	S.D. = 304.06

The study times of the complex and simple objects were compared. As with the meaningful objects, there was no significant difference between the study times for the complex and the simple stimuli ($t_{27} = 1.71$, p=.098).

To preclude the possibility that participants were guessing, a one-sample t-test was performed on the percentage errors at 150°; this analysis confirmed that the errors were significantly below chance ($t_{27} = 13.29$, $p<0.001$). There was no significant effect of order of completing the rotation tasks ($F< 1$).

Discussion

The meaningful stimuli showed an interaction between angle of rotation and complexity, demonstrating a piecemeal rotation process. In contrast, there was no interaction in the meaningless stimuli, suggesting that the rotation was holistic. Meaning would appear to affect how the rotation process occurs, encouraging a more part-based process.

However, there was a concern that the meaningful stimuli may not only be more meaningful than the meaningless stimuli, but also be more familiar (this was supported by the stimulus assessment described in Chapter Two). As previous research has suggested that familiarity can affect the rotation process, another group of participants was introduced to assess this factor. The third group used the meaningless stimuli in the rotation task. Prior to the rotation task, however, this group carried out a series of tasks to familiarise the participants with the stimuli, without practice with the rotation task itself. The pattern of errors in this group could determine whether the differences between the meaningful and meaningless groups could be due to the familiarity of the stimuli.

CHAPTER 3

EXPERIMENT TWO: ROTATION OF FAMILIAR, MEANINGLESS STIMULI

Method

Design

This experiment had a within-subjects design, with two factors, angle of rotation and complexity. As in the previous experiment, there were two levels of complexity (simple and complex), and four angles through which the stimuli were rotated (0°, 50°, 100°, and 150°).

Participants

There were 21 females and 7 males, with an age range of 18 to 43 (mean age 21.3). All the participants had normal or corrected-to-normal vision, and were right-handed. None of the participants had any experience with the stimuli or task prior to the experiment.

Materials

The materials used in this experiment were identical to those used in the meaningless condition of the previous experiment.

Procedure

The participants in this experiment performed an identical rotation task to the participants using unfamiliar, meaningless stimuli. However, prior to this they performed three familiarisation tasks. The familiarisation tasks were administered twice. On the first occasion, they were the only tasks performed and this session took place at least one day, but not more than seven days, before the rotation task. The second familiarisation session took place immediately prior to the rotation task. The tasks were chosen to provide experience and familiarity with the stimuli. The tasks consisted of imaging the objects and rating how vivid the image was, comparing each object to other objects, and comparing objects using a memory component. Performing a task, rather than just examining the objects, encouraged the participants to pay attention to the items, with an associated fuller depth of processing. It also enabled an empirical measurement to be taken, to ascertain that the participants acquired familiarity with the stimuli.

In the first task, participants were presented with each stimulus twice, in random order. Each time an object was presented, the participants were to spend as long as they wanted studying the object. They were then asked to close their eyes, and to generate an image of the object. Once this had been achieved, they had to rate the vividness of the image on a scale of 0 to 9, where 0 was no clear image at all, and 9 was a realistically life-like image.

The second task was a same/different judgement. The participants were shown two stimuli, and asked to say whether they were the same or different, as quickly as possible. "Different" stimuli consisted of two different objects, and "same" stimuli consisted of two identical objects. Each stimulus was shown a total of 18 times. The stimuli were presented one after the other, each remaining on the screen for one second.

The third task was similar to the one above, but included a memory component. The participants were shown three of the stimuli, one after the other, for 1.5 seconds each. Then a fourth stimulus was presented for 1.5 seconds. The participants had to say whether the fourth stimulus was one of the previous three or not. If it was, they responded “same”, and if it was not, they responded “different”. Again, all the stimuli were shown an equal number of times, and were counterbalanced for position in the four-stimulus sequence.

These three tasks were designed to increase familiarity, without practice with the rotation task. Same/different tasks were chosen so that the participant would be using the keys in the same way for the rotation task as for the familiarization tasks, avoiding interference between the tasks. To prevent the participants learning the modifications that were to be used in the rotation task, only the unmodified stimuli were shown, and only at the upright orientation.

Participants were asked to complete the stimulus assessment ratings described in Chapter Two. These were given following the first set of familiarity tasks and following the main experiment. This allowed efficiency of the familiarity procedures to be ascertained. The participants had to name, to the best of their ability, each of the objects, and then rate how confident they were in this decision on a scale of 0 to 10, where 10 was "certain", 5 was "unsure" and 0 was "not at all sure". Participants also used the same rating scale to judge how familiar the objects were, and how representative they were of the name applied to them.

Results

Once again, for the main analysis, the percentage error rates (see table 3.2) were converted into d' values.

An ANOVA with angle of rotation and complexity as within-subject factors showed a significant effect of complexity ($F_{1,\,27} = 93.49$, $p<.001$), a significant effect of angle ($F_{3,\,81} = 109.0$, $p<.001$), and no interaction between complexity and angle. ($F<1$; see Figure 3.05).

More errors were made with the complex than with the simple stimuli, with associated lower d' values. This was found for each angle of rotation (0°: $F_{1,27} = 29.68$ $p<.001$; 50°: $F_{1,27} = 31.38$, $p<.001$; 100°: $F_{1,27} = 29.99$, $p<.001$; 150°: $F_{1,27} = 30.20$, $p<.001$).

Both complex stimuli ($F_{3,81} = 57.87$, $p<.001$) and simple stimuli ($F_{3,81} = 71.38$, $p<.001$) showed an effect of angle of rotation. The difference in d' from one angle to the next was only significant between 0°and 50°for the complex stimuli ($t_{27} = 9.66$, $p<.001$), and between 0° and 50° and between

100° and 150° for the simple stimuli (0° and 50° $t_{27} = 12.4$, $p<.001$; 100° and 150°. $t_{27} = 2.17$, $p<.05$).

Table 3.4: Error Rates for the Familiar, Meaningless Stimuli

Angle of	Rotation	0°	50°	100°	150°
Familiar, meaningless stimuli	simple	Mean 13.36 S.D. 7.57	Mean 24.70 S.D. 7.80	Mean 24.18 S.D. 8.71	Mean 27.64 S.D. 7.01
	complex	Mean 20.81 S.D. 9.71	Mean 32.68 S.D. 8.40	Mean 34.19 S.D. 7.61	Mean 35.10 S.D. 7.45

The percentage errors for meaningful and meaningless stimuli are given for each angle of rotation, across simple and complex stimuli. The mean values and standard deviations are shown.

The response times were examined. There was no significant effect of complexity ($F_{1, 27} = 3.18$, $p=.086$), a significant effect of angle ($F_{3, 81} = 16.36$, $p<.001$), and no interaction between complexity and angle ($F< 1$). The difference in response time from one angle to the next was only significant between 0°and 50°(simple: $t_{27} = 4.30$, $p<.001$; complex: $t_{27} = 4.01$, $p<.001$).

Table 3.5: Response times for the meaningless, familiar stimuli

	Angle of Orientation			
	0°	50°	100°	150°
Simple stimuli	Mean = 484.90	Mean = 575.26	Mean = 589.82	Mean = 594.75
	S.D. = 163.12	S.D. = 173.79	S.D. = 192.13	S.D. = 188.40
Complex stimuli	Mean = 496.23	Mean = 628.67	Mean = 625.94	Mean = 655.20
	S.D. = 122.20	S.D. = 205.65	S.D. = 176.69	S.D. = 256.80

To preclude the possibility that participants were guessing, a one-sample t-test performed on the percentage errors at 150° ($t_{27} = 13.29$, $p<.001$) showed that they were significantly below chance. Study times of the complex and simple objects were compared. The participants took longer to study the complex objects (mean time 1,587 ms, S.D. 698 ms) than the simple objects (mean time 1,354 ms, S.D. 468 ms, $t_{27} = 2.15$, $p<.05$). There was no significant effect of order of completing the rotation tasks ($F<1$).

These findings are very similar to those found when the meaningless objects were used without familiarity tasks. To ensure that the familiarity tasks had had an effect on the participants, the results from the tasks were examined. First, the three computer-presented tasks were analysed.

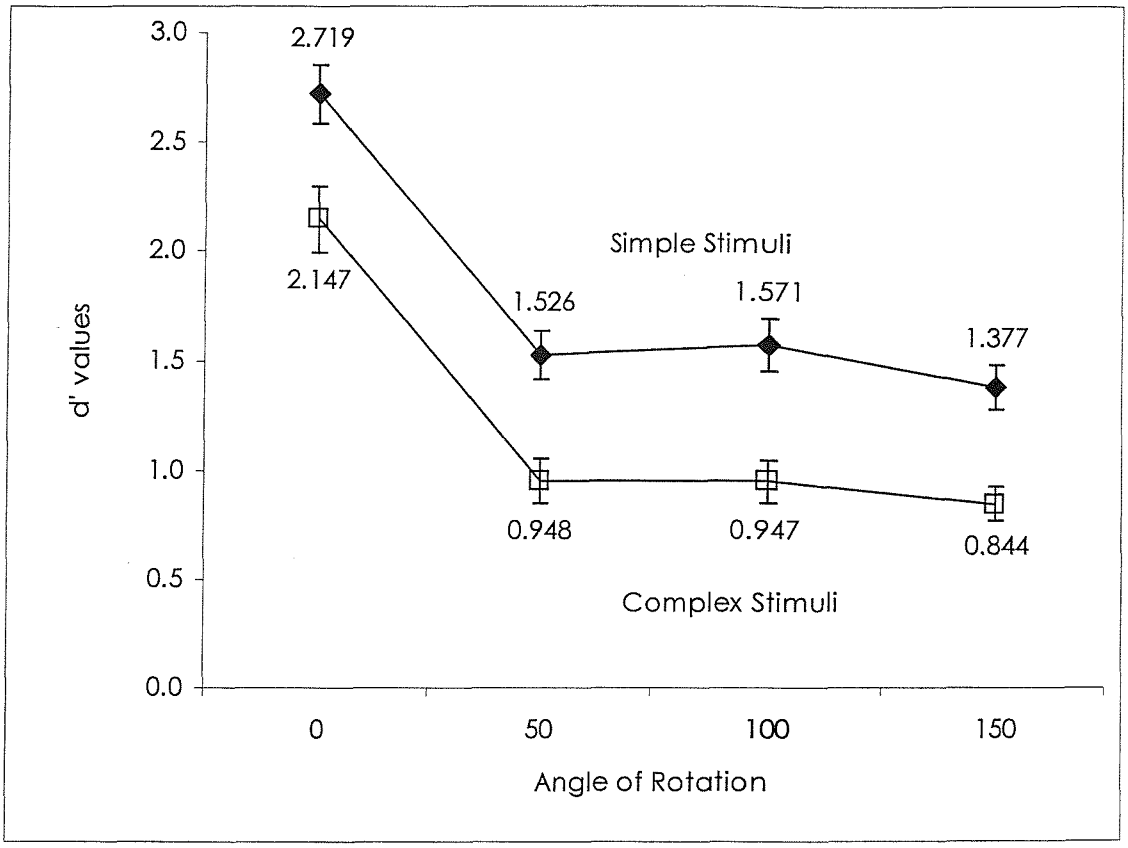

Figure 3.05: Mean d' values for the familiar, meaningless stimuli.

The simple stimuli are illustrated with closed diamonds, and the complex stimuli are illustrated with open squares. The simple stimuli demonstrate a higher d' for every degree of rotation than the complex stimuli. The pattern of results is very similar to the unfamiliar meaningless stimuli, with no interaction between angle of rotation and complexity. The mean value for each point is shown, along with standard error bars.

The response times on the first occasion were compared with those of the second occasion for the same/different and memory component tasks (see table 3.3). In both cases, there was a reduction in response times for correct answers from time 1 to time 2 (same/different task: $t_{27}=3.53$, p <.01; memory task: $t_{27}=3.55$, $p<.01$). For the vividness task, there was an increase in the ratings of vividness of the representation generated from time 1 to time 2 $t_{27}= 6.35$, $p<.001$).

Table 3.6: Mean Values of the Familiarity Tasks

Task	Performance at time 1	Performance at time 2
Vividness assessment (rating of vividness)	mean = 5.61 S.D. = 1.24	mean = 6.49 S.D. = 1.19
Same/different judgement (response times)	mean = 327.53 S.D. 139.46	mean = 263.1 S.D. = 113.39
Memory component task (response times)	mean = 340.52 S.D. = 114.13	mean = 290.5 S.D. = 103.41

Performances at time 1, after the first set of familiarisation tasks, and at time 2, after the second set of familiarisation tasks, are tabulated above.

In addition, the participants' subjective reports of familiarity from the first questionnaire (after the first session of familiarity tasks) were compared with those obtained in the stimulus assessment (see table 3.4).

The results of each participant were summed across stimuli, and three averages produced for the stimulus set. There were no significant differences between the values from questionnaire 1 and 2. These results were collapsed, and then compared with the judgements of the participants described in Chapter Two, who had completed no tasks at all.

The meaningful stimuli were more familiar than the familiar, meaningless stimuli ($t_{22} = 2.09$, $p<.05$); the familiar, meaningless stimuli were rated as significantly more familiar than the unfamiliar, meaningless stimuli ($t_{22} = 15.4$, $p<.001$). Although the ratings for representativeness and confidence in naming were significantly higher than for the meaningless stimuli (representativeness: ($t_{22} = 9.47$, $p<.001$; confidence $t_{22} = 9.01$, $p<.001$), they were significantly below those of the meaningful stimuli (representativeness: ($t_{22} = 7.95$, $p<.001$; confidence $t_{22} = 9.20$, $p<.001$).

Table 3.7: Mean Values of the Stimulus Assessment

Assessment	Questionnaire 1	Questionnaire 2
Confidence (0 = not at all; 10 = certain)	mean = 5.23 S.D. = 1.01	mean = 5.04 S.D. = 0.93
Familiarity (0 = not at all; 10 = certain)	mean = 6.82 S.D. = 0.59	mean = 6.91 S.D. = 0.48
Representativeness (0 = not at all; 10 = certain)	mean = 4.87 S.D. = 1.06	mean = 4.84 S.D. = 0.95

Questionnaire responses at time 1, after the first set of familiarisation tasks, and at time 2, after the second set of familiarisation tasks, are tabulated above.

GENERAL DISCUSSION

This study examined the type of transformation used when the participants were mentally rotating stimulus representations. The purpose of this was to access the format of the underlying representations of the stimuli, and, specifically, to assess whether there was a difference in representations between meaningful and meaningless stimuli. This was achieved by examining the pattern of error rates (expressed as d' values to avoid conflating errors made during same and different trials). Following the logic described in the Introduction of this Chapter, a piecemeal transformation predicted an interaction between complexity and angle of rotation, and a holistic transformation predicted no interaction between complexity and angle of rotation. It was assumed that the underlying representation would constrain the type of processing that can occur.

For the meaningful stimuli, there was an interaction; rotating the complex stimuli entailed a greater change in the proportion of errors made across increasing angle of rotation than that of rotating the simple objects. The interaction reflected that the representations were rotated part by part in a piecemeal process. For the meaningless stimuli, there was no interaction. The change in the proportion of errors made across increasing angles of rotation was the same for both simple and complex stimuli. The lack of interaction suggested that the representations were rotated holistically. The underlying representations would therefore appear to be segment-based in the case of the meaningful objects and holistic in the case of the meaningless objects.

However, as well as being more meaningful, the meaningful stimuli were also more familiar than the meaningless stimuli. The difference in the representations could be due to familiarity rather than meaning. To assess the role of familiarity a group of participants gained familiarity with the meaningless stimuli prior to the rotation task. In the familiar, meaningless stimuli, although there was an overall reduction in the number of errors made, there was again no interaction between the complexity and the angle of rotation. This suggests that it was the meaning of the stimuli, rather than familiarity with the stimuli, that encouraged a segment-based representation.

Analyses of the response times showed that the results were not explained by a speed-accuracy trade-off. The response times mirrored the proportion of errors except that there was no interaction in any of the three groups. This finding suggested that the processing of the separate parts in the meaningful stimuli might have taken place through a parallel process rather than through a serial process. Some of the contradictions in previous research may be explained through parallel processing of several parts being interpreted as holistic processing of only one part.

The findings of this study have been interpreted as the meaningful objects being represented by their parts and the meaningless objects being represented as a single unit. Although an apparent contradiction of other studies, this finding can be explained within current paradigms. One speculation is that the representations that already exist for a given object may interact and influence the type of processing that is used. The interaction could take two forms. These are not mutually exclusive, and may, indeed, be related (see Vecera & Farah, 1997). The first way is that the parsing of the objects may be influenced. Previous knowledge could cause the parsing to take place in such a way that the individual units are functional in themselves, and may therefore be easier to treat as separate units. Alternatively, previous knowledge may provide supporting information about the spatial relationships between the units, making it easier for the individual units to be re-aligned correctly. Meaningful objects have meaningful parts, and can be consistently segmented into such parts, enabling a piecemeal rotation process to be readily used on such a representation. In contrast, in meaningless objects the segments are not 'natural' entities and hence it may be easier just to rotate the entire representation in a holistic fashion.

Object recognition research can support this claim through viewpoint-independent theories; meaningful objects are represented by their constituting parts and their spatial relations to allow abstract viewpoint-invariant representations (e.g., Biederman, 1987; Biederman & Gerhardstein, 1993; Cave & Kosslyn, 1993). The finding that rotation can take place through both holistic and piecemeal processing supports the idea of flexible computations. Marr (1982) pointed out that several different algorithms would be able to perform the same computation. If one route cannot succeed, then an alternative can be selected. This flexibility allows accurate processing of information despite the constraints imposed by the circumstances; these constraints could include the nature of the representation (see also Cohen & Kubovy, 1993).

Although this seems a logical explanation for the findings, the focus is on the processing. We need to consider in more detail what these processes are implying about the underlying representations. First, the representations of the meaningless objects appear to be holistic. However, the "holistic" referred to in this account is the undifferentiated representation described as an early stage of representation by Palmer (1977). This inference arises because of the lack of evidence for any processing of parts. Second, although the representations of the meaningful objects have been described in terms of piecemeal processing, the account relies heavily on the presence of configural information to re-align the parts. This means that the representations of the meaningful stimuli could also be holistic, but in this case the term refers to the presence of strong configural information.

Bethell-Fox and Shepard (1988), in a study using random geometric shapes, suggested that greater familiarity with the stimuli results in the rotation process becoming more holistic. In the present study, both the familiar and unfamiliar meaningless stimuli appeared to be processed holistically. Familiarity with the meaningful objects did not make the representation holistic in the same sense, although it may have increased the available configural information. Bethell-Fox and Shepard achieved familiarity through exposure to the stimuli in the context of the rotation task; the present study achieved familiarity through separate, non-rotation tasks (for the meaningless stimuli) and previous experience with similar objects (for the meaningful stimuli). Therefore, it would seem that familiarity with the rotation task itself could have different effects from familiarity with the stimuli *per se*. If this is so, the process used during the rotation task may be constrained by other factors apart from the type representation formed of the object.

The conclusions drawn from the results of this study rest upon the assumption that mental rotation was taking place or, at least, some transformation was taking place where more effort is needed to complete the task at greater angles. If this assumption is not satisfied, then the conclusions drawn may have to be modified. The number of errors increased as the angle of rotation increased for both simple and complex stimuli in all three conditions. However, the biggest increase in errors occurred between the angles of 0°and 50°, and the increase was, generally, non-significant at greater angles. In the absence of a linear function between error rate and angle of rotation, we need to consider what may have been happening during the task in terms of processes and representations.

One possibility is that the deadline procedure prevented sufficient processing of the stimuli at the greater angles, and the participants were just guessing at rotations beyond 50°. Two findings make this an unlikely explanation. First, the responses were not at chance, even at 150° in the complex conditions, so pure guesswork would seem an unlikely option. However, given error rates approaching 40%, some level of guesswork would have been involved. Second, the response times showed a main effect of angle. This suggests that although the participants had been told to respond immediately, they did so only when they had a response to impart. Whether the participants were guessing or not, the task would seem more difficult when a rotated stimulus is involved, and this would result in more effort being needed to complete the task.

Another possibility is that some aspect of the task changed between 0° and the other angles, besides regular changes in angular distance. At 0° only a comparison task is performed. At the higher angles, a rotation and comparison task had to be performed. If the comparison task remained the same at each angle, this would provide a constant source of errors and latency. The difference between one angle and the next should then be due to rotation, and previous studies have suggested that rotation in the picture plane produces linear increases in performance measures. However, the relatively flat line after 50° suggested that it was not angle of rotation that was driving the difference. One possibility propounded by Murray (1995) was that it might be easier and quicker to rotate in the depth plane. This may not be successful for mirror representation decisions, but for imaging and naming, and for detecting one modified area of a stimulus, it could be a better technique in terms of speed and accuracy.

The other alternative is that once rotation was included in the task demands, there was change in how the comparison task was performed. One explanation, if this is so, is that the strategy for performing the task changed. Findings similar to the ones described here have also been described in other studies; notably, for example, in such studies as those carried out by Lawson and Joliceour (1998). In the first of these, for example, they found, during an identification task at several angles of rotation, a linear performance up to the angle of 60°, followed by a plateau (see Lawson & Jolicoeur, 1998, Experiment 1, Figure 1, p794). In their experiment, unlike the present study, there was no deadline procedure or pressure to respond, but, similar to the present study, the stimulus was only presented for a very brief time. Lawson and Jolicoeur suggested that a view-specific code was used to identify the stimulus, using a transformation-followed-by-match strategy. This was adequate, in their experiment, for angles up to 60°. However, if the presentation time of the stimulus prevented this strategy at higher angles, then the participant had to change to a view-independent strategy. This has been used to explain similar findings in other naming studies (e.g. Jolicoeur & Milliken, 1989; Murray, 1995). A viewpoint-independent computational model of object recognition also showed a similar function across angle of rotation (Hummel & Biederman, 1992).

A viewpoint-independent strategy would probably involve looking for a diagnostic part of the shape and using this to provide the response. However, the task in the present study was not an identification task; in fact, the comparison was across two stimuli that (certainly in the case of the meaningful stimuli) would share the same identity. The only diagnostic part of the shape was the modification, or, in the case of the same trials, an absence of a modification. The task, therefore, may encourage a viewpoint independent strategy, but the differing locations of the modification may have made this strategy difficult to use.

Carpenter and Just (1978) claimed that mental rotation consisted of two stages. One stage involved a search for diagnostic parts that could be used as "landmarks" during the process. The other was the actual rotation process itself. Both stages were dependent upon the orientation of the object, but the search was relatively quick whereas the rotation was relatively slow. If the rapid search for landmarks could provide the information needed to make a response, then there would be no need to wait for the relatively slow rotation process to be completed. The brief presentation time could also make the first stage difficult to accomplish, leading to errors in the responses.

An alternative possibility to overcome the demands of the brief presentation time is that only part of the stimulus was rotated, giving the appearance of a holistic rotation. In fact, if the part was rotated in a single unit, it would be a holistic rotation, but of the part, rather than of the complete stimulus. To take this argument a stage further, it may be that the entire stimulus is rotated as one unit. This implies that the rotation process, at least in the task presented here, was holistic in nature for all three types of stimuli, because the performance was comparable during the angles where rotation would be occurring. The significant effect from 0° to 50° could then be explained by a change in representation from a format that facilitated the comparison to another that was compatible with the rotation process. If the rotation process was holistic, this could mean a change from a part-based representation to a holistic representation. If only part of the representation was rotated, it would mean a change from a holistic representation to a part-based one, to allow a small portion to be isolated.

The main effect of complexity was not unexpected, reflecting the increased difficulty in one or more aspects of making the comparison with the complex stimuli. The expected outcome of this was a main effect of a constant difference between the simple and complex stimuli across all angles of rotation. However, although there appeared to be a constant difference across most angles in all three stimuli, this was not found for the meaningful stimuli at 0°. In fact, the d' values indicate that there was an advantage for the complex stimuli. This suggests that whatever change takes place between the task at 0° and the task at 50° and beyond is more profound in the meaningful stimuli than the meaningless stimuli, and more profound in the complex, meaningful stimuli than in the simple meaningful stimuli.

To take this argument further, if the difference between simple and complex stimuli was due to the comparison task, and this difference was constant across all angles for the meaningless stimuli, it implies that the comparison task itself, and, therefore, the representations involved, remained constant. Furthermore, for the meaningful stimuli, there appears to be a change in the comparison task, and the representations involved. This adds support to the speculation that there is an optimal representation for the comparison task, which appears related to the meaning of the stimulus. There is also an optimal representation for the rotation task, and if the two are not compatible, the representation has to be converted to allow the task to be performed. The conversion would appear to be either greater, or qualitatively different, in the meaningful stimuli. Moreover, this difference does not depend upon the familiarity with the stimuli.

The change in representation, if it occurs, is found in the meaningful stimuli, where a different pattern of results is found from the comparison task alone to the rotation and comparison task together. An alternative to a change in representation could be a change in the use of the information available in the representation. If both parts and configuration are present, then the comparison task could use the part information to detect one different segment in the whole stimulus. In contrast, the rotation task could use the configural information to keep the parts aligned correctly. This implies that these two sources of information are not readily accessible in the representations of the meaningless stimuli.

There is another explanation for the results found in the meaningful stimuli at 0° that differs slightly from the one above. The complex, meaningful stimuli are somehow "special" compared with the other stimulus types, including simple, meaningful stimuli. Whatever representational format the other stimulus types produce, meaning can convert if a given criterion, such as a task demand, warrants it. In this task, the conversion results in a significantly more accurate performance in the task, although not a reduced response time. The conversion could be towards a more part-based representation, which may allow better detection of the modifications among the stimuli. Equally, if the comparison involved more of a template matching strategy, the conversion could be to a more holistic representation, perhaps through the use of configural information.

The research began with the assumption that there was a critical link between representations and their subsequent processing. Specifically, it was assumed that the representation formed of an object would determine later processing. Although the study has shown a difference between the meaningful and meaningless stimuli, which can be interpreted through differences in representational format, it is not so clear that subsequent processing is determined by that format. An alternative possibility is that although there are "natural" formats for an object, and meaning of the object can influence this, the type of mental processing may over-ride this initial representation. This may make higher mental transformations, such as mental rotation, an unsuitable vehicle for determining the type of underlying representations.

To conclude, the results from this study demonstrate that meaning affects the pattern of errors that occurs during a mental transformation. This has been interpreted as meaningful stimuli being represented as part-based representations, although strong configural information among the parts can not be ruled out. Meaningless stimuli have holistic representations, in the sense of being one, undifferentiated unit. Familiarity does not appear to drive this difference. It is less clear whether the format of the representation actually guides the process itself, or whether the demands of the process guide the format of the representation, particularly in the meaningful stimuli.

CHAPTER 4

THE EFFECT OF MEANING ON PERFORMANCE IN A PART SEARCH TASK

ABSTRACT

The aim of this study was to use a part-search task to examine whether there was a difference in the representations of meaningful and meaningless objects. In the search task used here, for both part-based and holistic representations, a greater response time was predicted for complex stimuli than for simple stimuli. However, a holistic representation would produce a greater difference between simple and complex stimuli than a part-based representation. This is because in a holistic representation there is more dependence among the parts used than in a part-based representation, and this dependence will interfere with the separation of the parts. In addition, when the response times are combined with the pattern of the error rates, the different types of holistic representation can be distinguished. Both parallel processing and an undifferentiated representation will lead to less difference in response times between complexities than a representation with strong configural properties. However, because of the difficulty in isolating a single part, an undifferentiated representation is likely to produce performance close to chance. The participants studied a stimulus for as long as they wanted. They were then shown two parts, and had to make a two-alternative forced choice about which of the two had been part of the stimulus. Both meaningful and meaningless objects showed a significant effect of complexity; there was also a significant interaction between complexity and type of stimulus. The meaningful objects showed a greater effect of complexity than the meaningless objects. This suggests that the meaningful objects were represented more holistically than the meaningless objects.

INTRODUCTION

The representational format of a stimulus can have an effect on the processing of an individual part of the stimulus. If meaning has an effect on the type of representation, then meaning will affect how a single part is processed. For example, a word-superiority effect (WSE) has been a long-established finding. Letters in meaningful words were processed more readily than individual letters forming meaningless strings (e.g. Wheeler, 1970). The WSE has been explained as a memory effect, in which it is more efficient to rehearse one word, rather than several letters (Massaro, 1973). The word could be seen as an organisational schema for the letters. This would only be possible where the letters of the word could combine into a meaningful word. Later work in computer modelling has shown the importance of top-down factors in producing the effect. The familiarity with the word, and the understanding of allowable and unallowable letter combinations were offered as an explanation for the facilitatory effect (e.g. McClelland & Rumelhart, 1981).

Similar facilitatory effects have been found for objects. An object superiority effect (OSE) has been demonstrated in several domains. For example, Weisstein and Harris (1974; see also Lanze, Maguire, & Weisstein, 1985) showed that a line was more readily detected when it formed part of a three-dimensional object, than if the same line formed part of an equivalent two-dimensional geometric shape. The effect has also been demonstrated in faces (e.g. Homa, Haver & Schwartz, 1976), where a face superiority effect (FSE) was found. Facial features were easier to detect in a normal face, than in a scrambled face containing the same features. This has been expanded by Purcell and Stewart (1991) who found that a stimulus was detected and classified (as normal, inverted or scrambled) more quickly when it formed part of a three-dimensional object, than when it formed part of an equivalent two-dimensional geometric shape.

However, improved processing associated with meaning is not always found. In some cases, processing is actually impeded by the organisation of the material into a meaningful and coherent form. This is known as an object inferiority effect (OIE), and has been described in tasks where a small part of the whole stimulus has to be detected and processed (Banks & Prinzmetal, 1976; Prinzmetal & Banks, 1977; Suzuki & Cavanagh, 1995).

The OSE and OIE, in which segments are processed with more or less ease when they are part of an object, has been likened to an analogous effect between scenes and objects in the scene (e.g. Davidoff & Donnelly, 1990). Here, too, it has been shown that objects are processed more readily when they form a consistent and meaningful part of a scene that when they are incongruent (Biederman, 1972; but see also Pezdek, Whetstone, Reynolds, Askan, & Dougherty, 1989).

It would seem, therefore, that meaning could operate on the processing of individual parts of an object in two ways: either to facilitate the processing or to impede it. The improved processing of letters found in words suggests that meaning may play an organisational role, allowing the use of top-down processing. To understand OSEs and OIEs we need to explain: First, what property of the stimulus leads to the effect and, second, how the effect is realised. The explanation must be able to account for both the superiority effect and the inferiority effect. The relevance here is whether these effects can be attributed to the underlying representation; and, if so, whether such a task could be exploited to differentiate between the types of representations.

The properties appear to focus around the idea of the organisation of the material into meaningful patterns; the resulting representations of the object then affect performance in later tasks. Earlier work suggested several properties as responsible for the OSE. Originally, the three-dimensionality of the stimulus appeared to promote an OSE (Lanze, Maguire, & Weisstein, 1985), although it was not clear exactly how this factor operated. One possibility put forward was that the context could aid the detection of the target (Biederman, 1972), perhaps through the provision of a schema or perceptual set (Smith & Haviland, 1972). Again, this may relate to the organisation of the material; clear organisation may help in the rehearsal or accurate recall of the information. Context has also been posited as a guide for attentional mechanisms to locate the relevant components for successful completion of a task (Chun & Jieng, 1998). However, it is not universally agreed that context is any aid to detection at all (Hollingsworth & Henderson, 1998). Some studies have found that the incongruent information is better-recalled (Pezdek, Whetstone, Reynolds, Askan, & Dougherty, 1989).

Homa, Haver, and Schwarz (1979) considered that although context was important, context alone could not account for the superiority effect. They further considered that the effect was more than just organisation; the whole that resulted from such organisation in terms of an OSE was greater than good organisation of the component parts could produce. In this spirit, they suggested the Gestalt principles as possible important factors. The roles of connectedness and closure have been supported by several researchers (e.g. Pilon & Friedman, 1998; Prinzmetal & Banks, 1977). Connectedness is important as an organisational property of hierarchical representations of an object (Kimchi 1998). Furthermore, the role of the parts and their locations in representations of objects is closely related to the property of connectedness (Saiki & Hummel, 1998a; Saiki & Hummel, 1998b). Such properties, therefore, could allow the part to be located for further processing, possibly by using the relationship the part holds with respect to the whole to cue attention.

However, the organisational factors described above also prevent successful processing in other situations. Although these ideas are useful in explaining an OSE, they do not address what is happening during an OIE. Furthermore, the same stimulus can be found associated with both an OSE and an OIE (Mermelstein, Banks & Prinzmetal, 1979); therefore any property within the stimulus is not sufficient explanation alone.

Mermelstein, Banks and Prinzmetal (1979) suggested that rather than the organisational factor being at the centre of the difference, the explanation lies within the role of memory. The critical factor, they posited, was the degree to which the task relies on memory for the whole stimulus. If successful completion of the task required a high level of memory for the whole stimulus, then organisation of the stimulus would facilitate the performance in the task. If successful completion of the task did not require a high level of memory, and particularly if a reliance on more perceptual processes was involved, then the organisation of the stimulus would impede the performance in the task. This does not provide a complete explanation of the findings, however. Purcell and Stewart (1986), for example, found an OSE in a detection task, in which a memory component was not involved. However, the basic idea of task demands interacting with representational constraints is still sound.

The overall picture emerging from this research is that an organisational property in the representation interacts with the processing of the task in such a way that performance is sometimes facilitated and sometimes impeded. The issue of relevance to this study is whether the organisational property can be related to the holistic or part-based constituents of the representation, and whether, therefore, it could be used to distinguish between them. The answer could lie with an interaction between the format of the representation and the task to be performed. If the interaction is synergistic, then an OSE will result. If the interaction is antagonistic, then an OIE will result. If the stimulus had a strong holistic representation, then it could be difficult to ignore the whole in order to focus on a single part. If successful performance involved accessing a single part, then an OIE would result. However, if successful performance involved the whole stimulus, then the holistic properties could aid in the performance of a task, and an OSE would result.

Tanaka & Farah (1993) have described how hierarchical levels can relate directly to the representation. The object is parsed into parts, which results in a structural description. The number of parts can vary from one to many. Configural information is the relationship among the parts of the representation. As the configural information becomes more dominant, the parts lose their integrity. Eventually, there is no difference between within-part information and between-part information. At this point an object is represented at a high level, that is, holistically. The individual parts are not represented independently, but as part of the whole. In holistic processing, therefore, both part-based and configural information is used together (Tanaka & Sengo, 1997), in such a way that the integrity of the part-based information is lost as specific information in its own right.

It has been suggested that comparison tasks involve a top-down analysis, conveying an advantage to objects with higher-level hierarchical representations (Ankrum & Palmer, 1991). The role of hierarchical representations where the higher-levels may be given processing dominance is also explored in the idea of global processing showing precedence over local processing (Navon, 1977), and the use of an "outside-in" processing in analysing shape (Earhard & Walker, 1985). If this is the case, then an object that can provide a high-level representation will be at an advantage in such comparison tasks. Other tasks may be disadvantaged by the existence of high-level properties if their completion relies on detecting or processing the low-level, or part-based, properties.

Suzucki and Cavanagh (1995) compared the roles of low-level features (individual parts from which the stimulus was composed) and the conjunction of features (the configural pattern that emerged from the stimulus as a whole; the high-level properties). They looked at feature and configural searches of facially organised stimuli compared with the same features organised into a non-face. Their findings showed that there was only access to the higher level configurations during rapid search processes. This would mean that in material with strong configural information, the individual features could not be accessed, and the task would be rate-limited by the holistic search, regardless of which level of information would produce a more efficient process. Suzucki and Cavanagh suggested that object superiority will be produced if "...global features of the stimulus set are more rapidly processed for target discrimination than the constituent low-level features, whereas object inferiority occurs if the low-level features are more rapidly processed" (p 911).

Performance will be affected, therefore, by two factors. The first factor is whether the representation is holistic or part-based, and the second is whether the task is dependent on the holistic information for successful performance, or whether it depends more upon the parts for success. To take this a stage further, the task itself can determine whether strong configural or strong part-based information confers an advantage. Therefore, if success depends upon only a single part being processed, such as a feature search, then the performance should differentiate between a part-based and a holistic representation.

The processing of one part of a representation of a stimulus should not be impeded by the presence of other parts if the parts of the representation are independent of each other. That is, if the parts exist in a part-based representation, then isolating one part should be relatively easy. However, if the representation is holistic then the processing of one part will be impaired (Pomerantz, 1981; Tanaka & Farah, 1993; Tanaka & Sengco, 1997). Their relations with each other will affect the processing of the individual parts; isolating one small area from the whole will be relatively difficult. This would suggest that a search for a small segment of a part-based representation should take less time than a comparable search of a holistic representation.

The complexity of the stimulus will also influence the time taken. The more complex the stimulus is, then the more parts it can be broken in to. The more parts that a stimulus can be broken in to, then the longer it will take to perform a serial search of the parts. In a part-based representation, the search of a more complex stimulus would take longer than that of a simple stimulus, because there are more parts to search. However, as all the parts are independent from each other, the relationships among them will not interfere in the search. In a holistic representation, not only would a more complex stimulus have more potential parts, but also the more potential parts there were, then the more relations there would be among the parts. The dependence of the parts in a holistic representation would interfere with the processing of any one part. Therefore, there should be a greater difference in the response times (assuming a flat error rate across conditions) between simple and complex stimuli in a holistic representation than in a part-based representation. If the representation is an undifferentiated, single unit, then there may be little difference between performance with the simple and complex objects, because in both cases a part has to be created. However, in this case it will be very difficult to isolate a single part (e.g. see Donnelly & Davidoff, 1999), and consequently, performance measured by error rates will be close to chance. A parallel search would also lead to little difference across complexities, but in this case the errors should be below chance level.

The aim of the study was to assess whether the meaning of the stimulus influences the response times in a part search task. Both part-based and holistic representations could be expected to produce some increase in response times in a more complex stimulus than a simpler stimulus. However, a holistic representation would produce a greater difference between complex and simple stimuli than would a piecemeal representation. This would not be the case if an undifferentiated, holistic representation was formed, or the parts were processed in parallel. No effect of complexity on the response times would be predicted in either of these cases. However, in addition, an undifferentiated representation would produce a high error rate.

Method

Design

This study was a within-subjects design, incorporating two factors. There were two levels of meaning (meaningful and meaningless) and for each level of meaning, there were two levels of complexity (simple and complex). In each of these four conditions, the response times of the correct responses and the error rates in a part search task were measured.

Participants

Twenty-eight participants were recruited from the Psychology Department participant pool at Southampton University. There were 22 females and 6 males, with ages ranging from 19 to 43 (mean age 21.5). All the participants had normal or corrected-to-normal vision. None of the participants had the seen the stimuli before.

Materials

There were 24 stimuli; 12 meaningful stimuli (six simple and six complex) and 12 meaningless stimuli (six simple and six complex), which were described in previous chapters. A unique part was derived from each of these stimuli. The part consisted of a geometric line drawing, composed of five 0.5 cm lines, taken directly from a part of the stimulus; the area selected varied across stimuli. The parts were designed to match as closely as possible across meaningful and meaningless stimuli, and across simple and complex stimuli (see Figure 4.01; see also appendices 5 and 6).

The parts were determined by several considerations. A recognisable part of the stimulus was avoided, for two main reasons. First, the object represented in the stimulus should not be recognisable from the part alone (this would circumvent the need for a search at all). Second, such a part would advantage the meaningful stimuli above the meaningless stimuli. The response would then follow from immediate recognition for the meaningful stimuli, and a search for a part for the meaningless stimuli. The same sized part was used across complex and simple stimuli, to reinforce the predicted effect of complexity. A part should not appear in any other object but the one it was taken from, to avoid any ambiguity in the task when choosing between two parts.

Procedure

Each participant performed two tasks. The tasks were identical, except that one task used the meaningful objects, and the other used the meaningless objects. The order of the tasks was fully counterbalanced across participants.

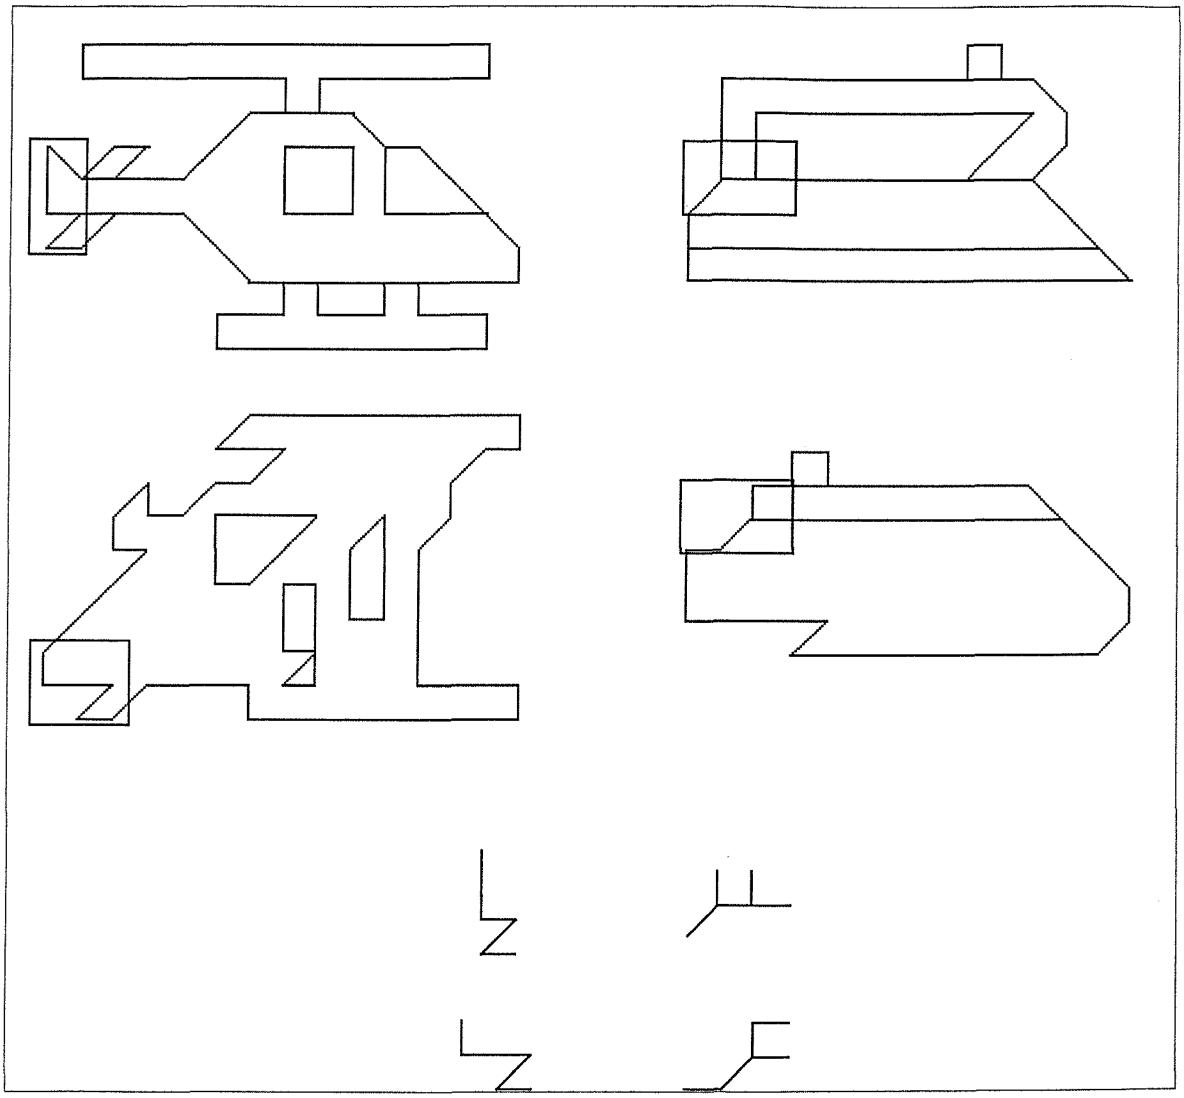

Figure 4.01: Examples of Stimuli and Parts

An example of a complex stimulus is shown on the left and a simple stimulus on the right. The top stimuli are meaningful, and the bottom stimuli are meaningless. The parts used in the feature search are outlined, and are shown in isolation underneath.

The search task was a two-alternative-forced choice task. First, a stimulus was presented in the centre of the screen for as long as the participant wanted to study it; when the participants were ready to continue, they pressed the space bar with the fingers of their non-dominant hand (see figure 4.02).

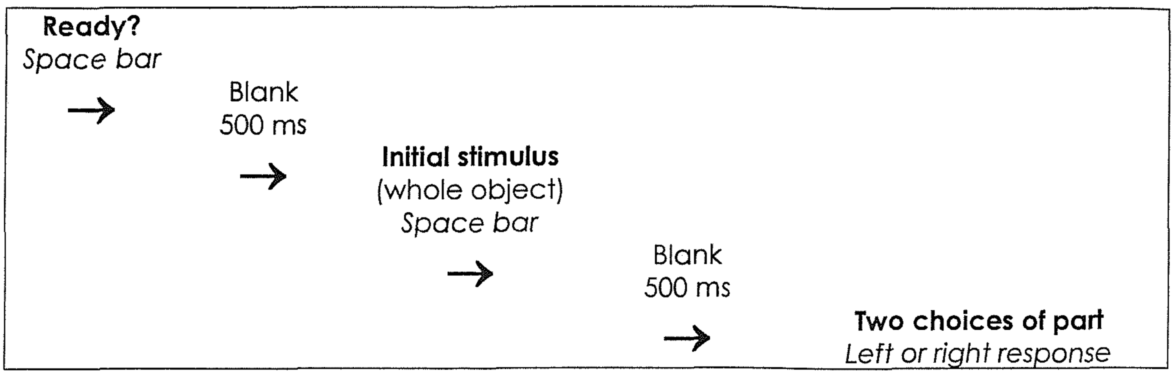

Figure 4.02: Events within a Trial of the Experiment

The events visible to the participants are presented in bold. The first stimulus is a whole object. The second stimulus consists of two parts; one was present in the whole object, and one was not. The items under each event are either the length of presentation or the response required from the participant.

The study time was determined by the participant to aid the memory component of the task. The aim of the task was to perform a search of the representation. Giving the participant as long as they wanted to study the stimulus should have ensured adequate encoding of the representation. A blank screen was then shown for 500 ms. Following this, two different parts were presented; one on the right and one on the left. One of these parts was present in the first stimulus shown in the trial. The other part was not present in the shape, although it may be similar to a part that was. The participant had to choose which part had been present in the whole stimulus, as quickly as possible, while maintaining accuracy. The participants indicated their answer by pressing keys labelled “R” (the “n” key on a standard keyboard) or “L” (the “b” key) with the first two fingers of their dominant hand. The parts remained on-screen until a response was recorded.

There were 48 trials in each of the two tasks. There were 12 parts, each one from a different stimulus. All the parts were presented the same number of times. In half the trials in which a part was presented it was the correct choice, and in the other half of the trials it was the incorrect choice. Each part was also presented on the right in half the trials, and presented on the left in the other half. The second experimental task was identical to the first task, but the other set of stimuli was used. If the meaningful stimuli were used in the first task, then the meaningless stimuli were used in the second task, and *vice versa.*

Before the two main tasks, each participant performed a baseline search task, using only the parts. The baseline task was very similar to the experimental tasks, differing only in the first presentation. Instead of presenting an object, however, one of the parts was presented. Following this, two different parts were presented, one on the right and one on the left. One of these parts was the same as the first part, the other was not. The participant had to choose which of the two parts was presented initially, as quickly and as accurately as possible. The parts remained on-screen until the participant responded. The purpose of the part-alone task was to provide a baseline that would check that the parts were discriminable from each other, and that response times reflected the search for the part rather than an effort to distinguish between the two parts.

Some of the parts might be more salient than others, and therefore recognisable without a search. Although the parts were designed with this problem in mind, it seemed prudent to check whether the design had prevented the problem, especially in the simple objects where the number of potential parts is less. The participants were tested after the baseline task. The participants were shown each of the parts, and had to say what object they thought the part might have been taken from. They were told not to spend too much time thinking about it, but to go with their first impressions. They were given the option of saying “don’t know” if they had no idea. This test was repeated at the end of the experiment to ensure that the participant had not learned to associate the correct part with the object, and could therefore complete the task without performing a part search.

A within-subjects design had been chosen to provide statistical power, and for increased control for participant variables in the task. However, the nature of the design gave rise to a potential concern. The participants were tested on both sets of stimuli, meaningful and meaningless. There was a possibility that the participant might recognise the meaningless stimuli as a version of the meaningful stimuli, particularly in the group who performed the task with the meaningful stimuli first. To examine this possibility, the participants were tested after the first and second tasks, referring only to the stimuli they had just seen in the preceding task. In each case, each stimulus was shown, and the participant had to name the stimulus, rate their confidence in the name, the familiarity, and the representativeness (as in previous stimulus assessments). The difference in this assessment is that they were given the option of "don't know" if they had no idea what the stimulus was.

Results

The data of interest were the response times of the correct responses (see Figure 4.03). In particular, the question was whether there was a difference between the meaningful or meaningless stimuli, and how meaning interacted with stimulus complexity. In a part-based representation, the prediction was that there would be less difference between the simple and complex response times than in a holistic representation. The error rates would also be informative in determining whether performance was at chance or not.

Analysis of Response Times

The response times were analysed using an ANOVA, with the type of stimulus and the complexity of the stimulus as within-subject factors, and order of presentation as a between-subjects factor. There was a main effect of complexity ($F_{1,26} = 15.97, p < .001$), and a significant interaction between the type of stimulus and the complexity ($F_{1,26} = 5.22, p < .05$), but no main effect of type of stimulus ($F < 1$). There was no main effect of order ($F < 1$).

Further analyses examined the nature of the interaction. In both the meaningful and meaningless stimuli, response times involved in a search of the complex stimuli were longer than those in a search of the simple stimuli (meaningful: $F_{1,27} = 11.52, p < .01$; meaningless: $F_{1,27} = 6.17, p < .05$). This difference was greater for meaningful stimuli than for the meaningless stimuli. Neither the simple stimuli ($F_{1,27} = 2.57, p > .1$) nor the complex stimuli ($F_{1,27} = 1.27, p > .1$) showed a difference between meaningful and meaningless stimuli.

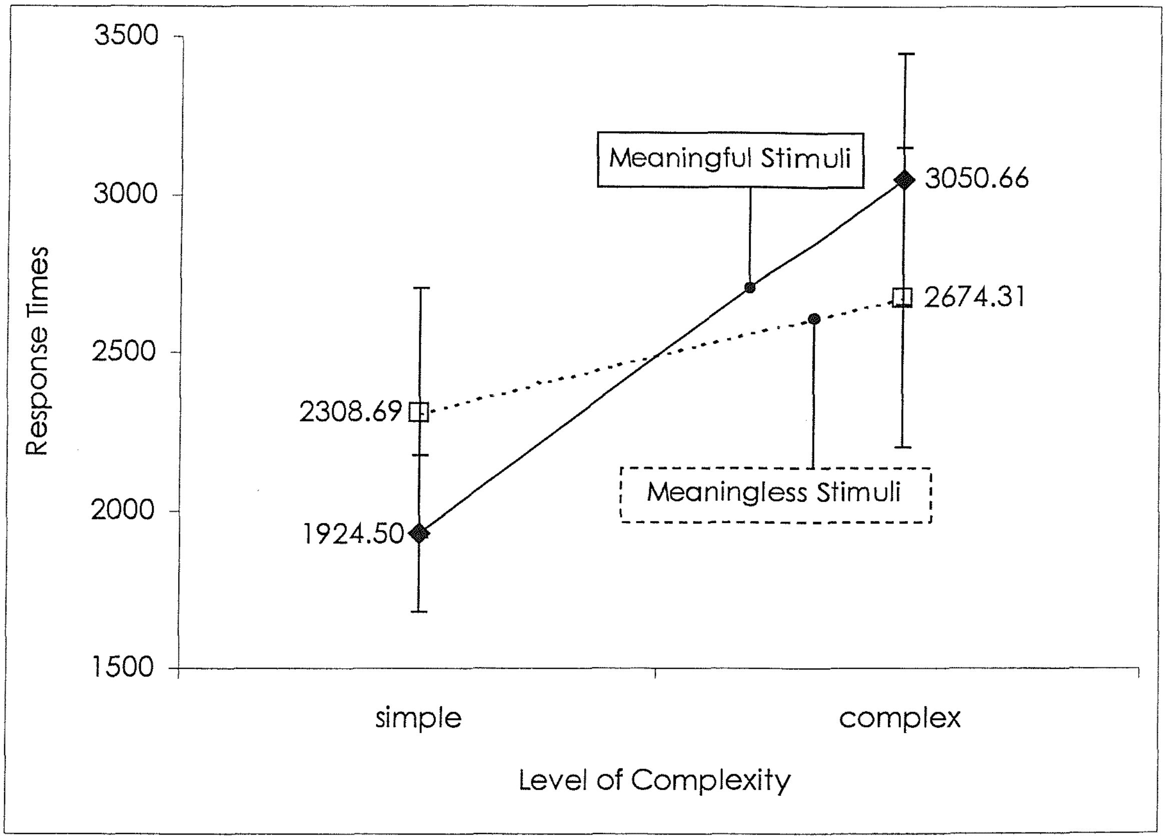

Figure 4.03: Mean Response Times of Meaningless and Meaningful Objects across Complexity.

The meaningful stimuli are shown with closed diamonds, and the meaningless stimuli are shown with open squares. The complex stimuli show a longer response time than the simple stimuli; this difference is greater in the meaningful stimuli. The mean values and standard error bars are given.

Analysis of Error Rates

The analysis was repeated with the error rates (see figure 4.04) to ensure that a speed-accuracy trade-off had not operated, and also to check the level of performance. First, a one-sample t test revealed that the percentage errors were significantly below 50% in all four conditions (simple, meaningful: $t_{27} = 11.3$, $p < .001$; complex, meaningful: $t_{27} = 9.79$, $p < .001$; simple, meaningless: $t_{27} = 13.97$, $p < .001$; complex, meaningless: $t_{27} = 9.16$, $p < .001$). This suggested that the representations were not completely undifferentiated.

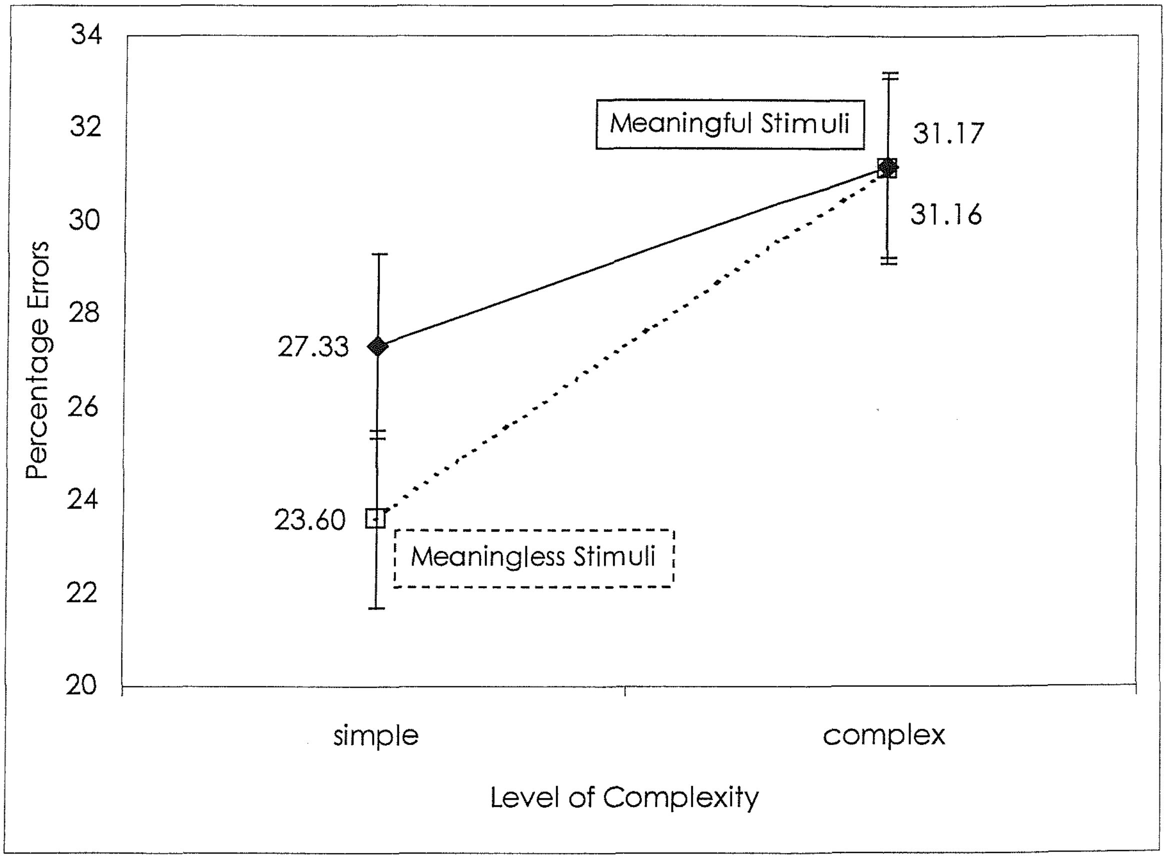

Figure 4.04: Mean Error Rates of Meaningless and Meaningful Objects across Complexity.

The meaningful stimuli are illustrated with closed diamonds, and the meaningless stimuli are illustrated with open squares. The complex stimuli show a greater percentage errors than the simple stimuli. The mean value for each point is shown, along with standard error bars.

The logs of the error rates were used in the ANOVA. There was no effect of type of stimulus ($F_{1,26} = 1.76$, $p > .1$), a significant effect of complexity ($F_{1,26} = 10.70$, $p < .01$), and no interaction between type and complexity ($F_{1,26} = 1.35$, $p > .1$). There was no order effect ($F < 1$).

More errors were made with the complex meaningless objects than with the simple meaningless objects ($F_{1,27} = 8.48$, $p < .01$). This difference was not significant for the meaningful objects ($F_{1,27} = 2.54$, $p > .1$). Although there was no difference between the meaningful and meaningless complex stimuli ($F = <1$) there was a borderline difference between the two type of stimuli at the simple level ($F_{1,27} = 3.16$, $p = .087$). The meaningful stimuli showed a lower response time than the meaningless in the simple condition, but also showed a higher error rate. This difference was not significant for the response times, and was only borderline for the error rates.

Analysis of the Baseline Values

The parts were made equivalent across the conditions during the design of the stimuli. However, despite this consideration, one possibility was that the parts themselves were harder to discriminate in some of the conditions than in others. Examination of the baseline values checked this (see table 4.1). The baseline values were derived from the trials when a part was presented as the initial stimulus.

Table 4.1: Baseline values for all stimulus conditions

	Percentage errors		Response times	
	Simple	Complex	Simple	Complex
Meaningful stimuli	Mean = 2.53 S.D. = 9.51	Mean = 3.58 S.D. = 7.75	Mean = 639.23 S.D. = 223.82	Mean = 677.24 S.D. = 302.49
Meaningless stimuli	Mean = 2.08 S.D. = 8.75	Mean = 4.63 S.D. = 6.86	Mean = 645.43 S.D. = 226.33	Mean = 663.01 S.D. = 217.48

An ANOVA on the baseline response times showed no effect of type ($F<1$), a borderline effect of complexity ($F_{1,27} = 4.125$, $p= .052$), and no interaction between type and complexity ($F<1$). The effect of complexity was greater in the meaningful stimuli ($F_{1,27} = 3.82$, $p= .061$) where there was a longer response time for the complex stimuli.

Error rates for all three conditions were low, with mean error rates of less than 5%. An ANOVA on the baseline error rates showed no effect of type ($F<1$), an effect of complexity ($F_{1,27} = 7.621$, $p< .05$), and no interaction between type and complexity ($F_{1,27} = 1.91$, $p> .1$). The effect of complexity was localised to the meaningless stimuli ($F_{1,27} = 6.54$, $p< .05$) where the mean error rate from the complex stimuli (mean = 4.6) was greater than the error rate from the simple stimuli (mean = 2.1).

The most important information from the baseline analysis was that there was no effect of stimulus type, or interaction involving it. The effect of complexity was a cause for minor concern. The original data were corrected using the baseline values, by subtracting the baseline response times from the object response times for each participant. This would result in a response time for the part search and extraction of the part from the object, without the time for the comparison between the two parts.

The results confirmed the original response times. This suggests that any difference in baseline results was not the cause of the results described earlier. There was no effect of stimulus type ($F<1$), a significant effect of complexity ($F_{1,27} = 13.87$, $p< .01$), and a significant interaction between type and complexity ($F_{1,27}= 4.80$, $p< .05$). For both meaningful stimuli ($F_{1,27} = 10.73$, $p< .01$) and meaningless stimuli ($F_{1,27} = 5.76$, $p< .05$) the complex stimuli produced longer response times than the simple stimuli. The difference between the simple and complex response times was greater for the meaningful than for the meaningless stimuli.

Analysis of the Study Times

The study times were analysed (see table 4.2). The relatively long study times suggested that the participants were trying to encode the stimuli prior to performing the search task, although there was considerable variation about the mean.

Table 4.2: Comparison of Study Times Across all Conditions

	Meaningful Stimuli		Meaningless Stimuli	
	Simple	Complex	Simple	Complex
Study Time	Mean: 2700.79 S.D.: 1593.40	Mean: 4982.68 S.D.: 3299.35	Mean: 3892.46 S.D.: 2908.3	Mean: 5121.84 S.D.: 3798.16

There was a borderline effect of stimulus type ($F_{1,27} = 3.31, p = .08$), a significant effect of complexity ($F_{1,27} = 22.49, p < .001$), and a significant interaction between type and complexity ($F_{1,27} = 5.47, p < .05$). For both meaningful stimuli ($F_{1,27} = 20.24, p < .001$) and meaningless stimuli ($F_{1,27} = 12.79, p < .01$) the complex stimuli produced longer study times than the simple stimuli. The meaningless stimuli were studied for longer than the meaningful stimuli in the simple condition ($F_{1,27} = 7.49, p < .05$), although there was no difference in the complex condition ($F < 1$).

Stimulus Assessment

Recognition of the Parts

Further tests had checked whether the parts were identifiable as being from a stimulus, without the necessity of a part search. This was assessed by analysing the responses given to the test taken after the baseline task, and at completion of all the tasks. The participants were asked to identify the stimulus from which the part had been derived. The number of correct answers was expressed as proportion recognition (see appendix 7). A response was considered correct if it corresponded to the name that the participant had given the object. This was important for the meaningless objects, which had no standardised name. It was also important for the meaningful objects where the participant might have used a different name from the standard (e.g. candelabra for the anchor).

Before the search task, participants were poor at recognising the parts in the meaningless objects. As the participants had never encountered the stimuli before, any other result would have been surprising. Following the search task, three of the twelve stimuli showed some degree of recognition. The number of correct identifications made remained low. However, the participants may have had difficulty identifying the object without the benefit of an established name for the object. This would result in a low estimate of recognition.

The meaningful objects showed higher recognition. Before the task, recognition was low, apart from the ship stimulus. After the task, several parts had been reliably associated with a particular stimulus. Three stimuli were identified as being particularly problematical: the ship, the house and the hedgehog. To assess the effect that the recognition may have had on the overall results, the trials in which one of these parts were involved were removed. The trials where the part was present, but incorrect, were removed as well as the trials where the part was correct, because these could be responded to by using a process of elimination, which still depended upon recognition of the part. For example, if the ship part was presented with the iron as the object, the other part could be chosen because the ship part was recognised, even if the iron part was not known. In both cases, a part search would not be necessary. Once the stimuli that included strongly recognised parts had been removed, the analyses of response times and error rates were repeated.

The pattern with the error rates in the meaningful stimuli showed a reduced number of errors in the simple condition, but an increased number of errors in the complex condition (see figure 4.05). An ANOVA using the logs of the error rates was performed. There was a main effect of complexity ($F_{1,27} = 31.10$, $p < .001$), a main effect of type ($F_{1,27} = 7.84$, $p < .01$), and a significant interaction between type and complexity ($F_{1,27} = 15.16$, $p < .01$). As with the response times, for both meaningful ($F_{1,27} = 25.68$, $p < .001$) and meaningless objects ($F_{1,27} = 8.48$, $p < .01$) there were more errors with complex than with simple stimuli. There were more errors with the meaningless stimuli than with meaningful stimuli in the simple conditions ($F_{1,27} = 12.77$, $p < .01$). There was no difference between the types of stimuli in the complex condition ($F_{1,27} = 3.00$, $p > .1$).

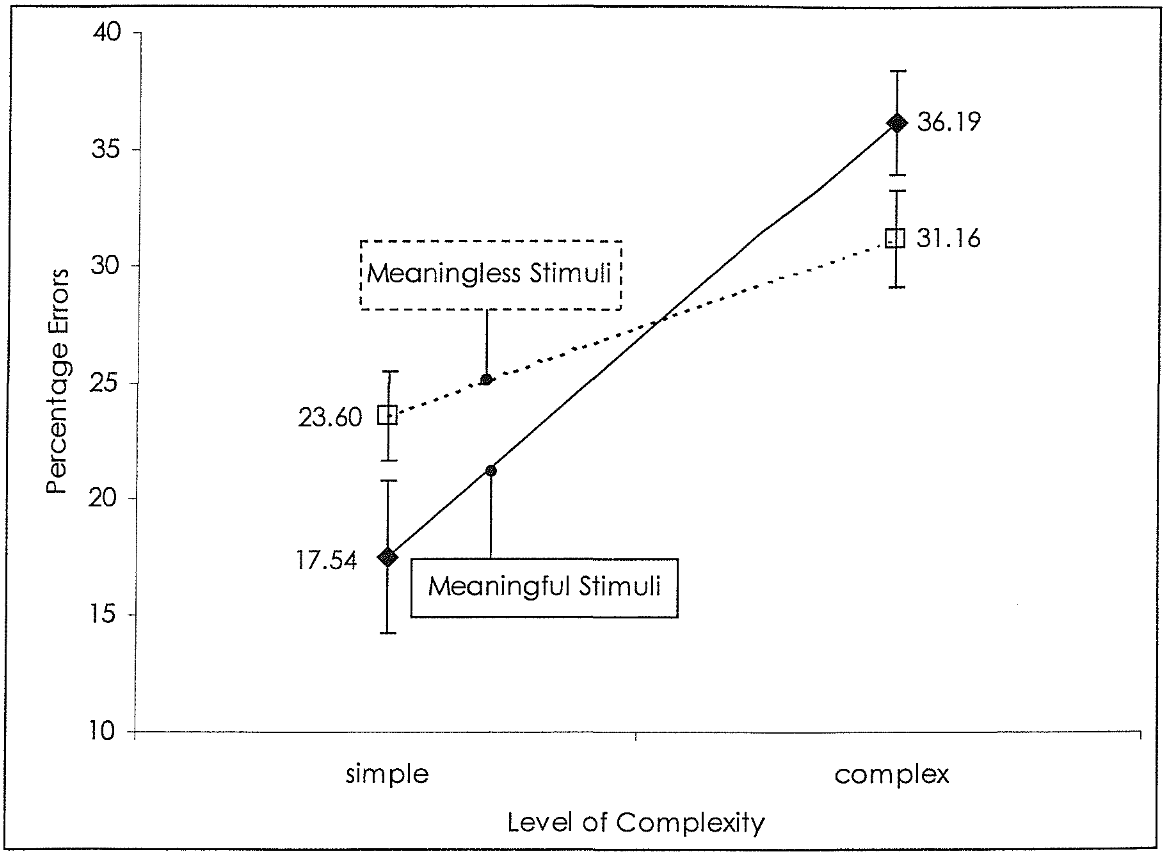

Figure 4.05: Mean Errors of Meaningless and Meaningful Objects after Correction for Recognisable Parts.

The meaningful stimuli are illustrated with closed diamonds, and the meaningless stimuli are illustrated with open squares. Trials in which the parts may have been recognised without the need for a search have been removed. The complex stimuli show a greater percentage errors than the simple stimuli; this difference is greater in the meaningful stimuli. The mean value for each point is shown, along with standard error bars.

The pattern of response times remained unchanged. There was a main effect of complexity ($F_{1,27} = 16.81$, $p < .001$) and a significant interaction between type and complexity ($F_{1,27} = 6.26$, $p < .05$). There was no main effect of type, ($F<1$). For both meaningful ($F_{1,27} = 12.56$, $p < .01$) and meaningless stimuli ($F_{1,27} = 6.17$, $p < .05$) the response with the complex stimuli took longer than that with the simple stimuli. There was no difference in response times across the type of stimulus at either the simple level ($F_{1,27} = 2.43$, $p > .1$) or the complex level ($F<1$). The difference between the simple and complex levels remained greater in the meaningful stimuli.

Stimulus Difficulty

The removal of trials in which the part may have been recognised increased the errors in the complex, meaningful stimuli, but decreased the errors in the simple, meaningful stimuli. In addition, although an effect of recognition was not found in the meaningless condition, this could have been a function of difficulty in providing a consistent name for the object. Furthermore, there is an assumption that recognising what shape the part came from during a questionnaire task would be predictive of performance in a timed computer task. It may be that although the participant could recognise the part in the questionnaire, this information was not available or not useful during the search task. In some trials, the task may have been very difficult for the participant to perform successfully.

To assess this, an analysis of the stimuli was carried out, with the purpose of excluding stimuli that caused difficulty in the search task. The number of correct responses for each stimulus was calculated, and any stimulus that produced a correct response rate of less that 65% was removed from the analysis. Three stimuli were removed from the complex meaningful set, one from the simple meaningful set, three from the complex, meaningless set, and none from the geometric, simple set. Once these stimuli had been removed, the means of the error rates were recalculated for each set (complex meaningful: mean = 26.21, S.D. = 14.22; simple meaningful: mean = 21.46, S.D. = 12.10; complex meaningless: mean = 22.33, S.D. = 13.67; simple meaningless: mean = 23.32, S.D. = 10.28). The ANOVA of the log of the error rates showed no effect of type of stimulus ($F < 1$), no effect of complexity ($F < 1$), and no interaction ($F < 1$).

The response times for each set were also recalculated (see figure 4.06). The ANOVA of response times showed no effect of type of stimulus ($F_{1,27} = 1.38, p > .1$), a main effect of complexity ($F_{1,27} = 8.01, p < .01$), and no interaction ($F_{1,27} = 2.3, p > .1$). For both meaningful ($F_{1,27} = 5.49, p < .05$) and meaningless stimuli ($F_{1,27} = 4.1, p = .05$) the complex stimuli showed longer response times than the simple stimuli.

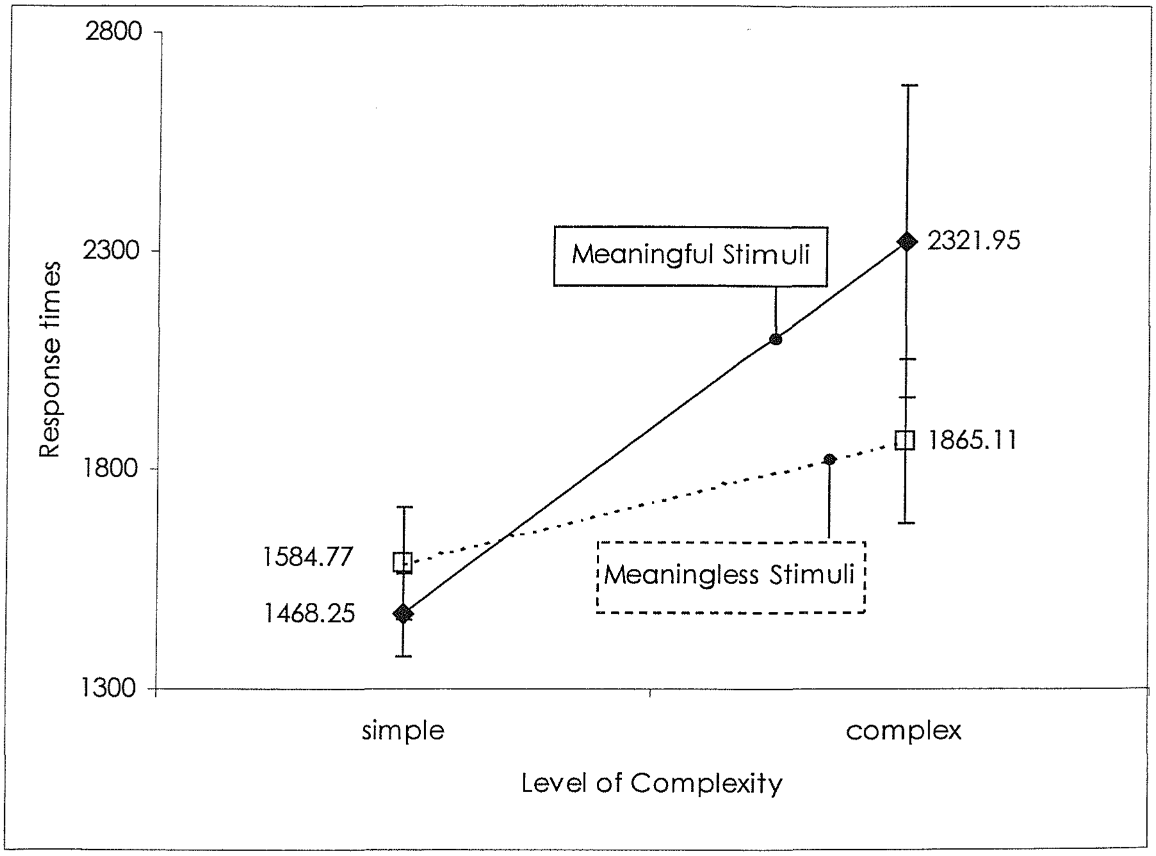

Figure 4.06: Mean Response Times of Meaningless and Meaningful Objects across Complexity after Correction for Difficult Stimuli.

The complex stimuli show a longer response time than the simple stimuli. The mean values and standard error bars are given.

Assessment of the Meaningfulness of the Stimuli

The participants had rated the stimuli, to ensure that they did not find the meaningless stimuli recognisable or representative, but that they did find the meaningful stimuli so. The participants did not appear to recognise the meaningless stimuli with any degree of ease. They found the stimuli difficult to name, and gave low ratings for representativeness and familiarity (see table 4.3). These results stand in contrast to those of the meaningful stimuli, where there were higher ratings for all three judgements.

Table 4.3: Assessment of Meaningful and Meaningless Stimuli

	Meaningful Stimuli		Meaningless Stimuli	
Judgement	Simple	Complex	Simple	Complex
Confidence	Mean: 9.07	Mean: 9.30	Mean: 1.77	Mean: 1.47
	S.D.: 1.39	S.D.: 0.67	S.D.: 1.09	S.D.: 0.50
Representativeness	Mean: 7.71	Mean: 8.39	Mean: 1.47	Mean: 1.02
	S.D.: 1.50	S.D.: 1.00	S.D.: 0.96	S.D.: 0.55
Familiarity	Mean: 7.43	Mean: 7.69	Mean: 1.40	Mean: 1.27
	S.D.: 1.54	S.D.: 0.97	S.D.: 0.72	S.D.: 0.38

Four judgments were made. In the first (not shown here), the participant had to name the depicted object. The second judgement was to rate their confidence in the name they had chosen. The third judgement was to rate how representative they thought the stimulus was of the type of object, and finally they had to rate how familiar they thought the object. The last three judgements were made on a scale of 0 to 10, where 0 was defined as "not at all", 5 was "unsure" and 10 was "completely".

These findings suggest that there is an effect of complexity for both meaningful and meaningless stimuli. A part search of a complex object shows a poorer performance, in terms of response times and error rates, than the performance in the same task with simple objects. More importantly, the effect of complexity is greater in the meaningful stimuli than in the meaningless stimuli.

CHAPTER 4

DISCUSSION

This study looked at whether there was a difference in the representations of meaningful and meaningless objects. In contrast to the last study, which used a mental transformation to assess the difference, here a part search task was used. Both segment-based and holistic representations could be expected to produce an increase in response times in a complex stimulus relative to a simple stimulus. In both cases, the complex stimulus contains more information relative to the part than the simple stimulus. However, a holistic representation would produce a greater difference between complex and simple stimuli than would a segment-based representation. This is because not only does the representation have to be searched for the part, but also, as there is more interference among the parts than in a more independent, segment-based representation, it would take longer to extract a single part. Therefore, in a part-based representation, the prediction was that there would be less difference between the simple and complex response times than in a holistic representation. Parallel processing of a number of parts would give no effect of complexity. An undifferentiated holistic representation would also show no effect between simple and complex conditions, but in this case error rates would be high (probably close to chance).

With regard to response times, there was an effect of complexity for both meaningful and meaningless stimuli. Although there was no overall effect of meaning, the effect of complexity was greater with the meaningful stimuli than with the meaningless stimuli. There was a greater difference between the response times in the simple and complex meaningful stimuli than between the simple and complex meaningless stimuli. This suggests that meaningful stimuli have a more holistic representation than the meaningless stimuli. Furthermore, neither representation was likely to be individual parts processed in parallel, or an undifferentiated representation, because there was an effect of complexity. These findings are consistent with meaningful objects producing more holistic representations; the holistic representation is defined in terms of having parts, but with the emphasis on the configural information rather than the part information.

A further possibility that needs to be considered is that there may be a speed-accuracy trade-off occurring, the most likely of which would probably be with the meaningless complex objects. In this case, the task may have been so difficult with some of the stimuli that the participants just gave up and guessed, thus giving an artificially low response time. This was resolved by removing all stimuli where there was an error rate of 35% or greater, and repeating the analysis. The pattern with the response times remained the same. Furthermore, based on the number of stimuli that had to be removed, the meaningless, complex condition was no more difficult than the meaningful, complex condition.

Alternatively, some of the parts could have been associated with the stimulus, eliminating the need for a search and, again, leading to falsely low response times. This may be accompanied by a higher than expected error rate if the participants thought they knew what object a part came from, but they were, in fact, mistaken. This would lead to a short response time, but a high error rate. This was resolved by assessing which stimuli were recognisable from the part, and removing all trials containing this part. Although this affected the pattern of the error rates (making it more like the pattern with the response times), the pattern with the response times remained largely unaffected.

The findings presented here suggest that the effect of meaning is to make the representation of that object more holistic than that of an equivalent meaningless object. This is compatible with other findings. For example, previous research has found that whereas faces and chairs are represented holistically, scrambled faces and chairs are not (Davidoff & Donnelly, 1990). Other studies (Tanaka & Farah, 1993) have found that non-face objects are not represented holistically. However, this is not necessarily a contradiction. Representations are probably best thought of as lying on a continuum from very holistic to very part-based (Tanaka &

Farah, 1993). If this is so, then, certainly in a task such as the one presented in this study, how holistic a representation appears will depend upon what it is compared with. Therefore, it may appear less holistic in one context, and more holistic in another, depending upon the format of the representation that it is contrasted with. The issue here is not whether the representations of objects are more or less holistic than those of faces, but rather whether a representation of a meaningful object is more or less holistic than that of a meaningless equivalent.

Tanaka and Sengco (1997) claim that in a holistic representation both configural and part-based information is present. This could fit in with the results found in this study and the previous one. In this study, a part was presented in isolation, and a search of the representation had to be made. A strong configural element would interfere with the extraction of a single part. In the previous study, the comparison was between two whole stimuli, and here the configural information may aid in the detection of a single part by acting as a cue, as described in the introduction. Kimchi (1994) suggested that, all things being equal, configural properties are used in discrimination rather than component properties.

One assumption in the interpretation of these results is that in a holistic representation, the parts are not independent, and consequently one factor in the response time is the need to break the part away from the surrounding areas. In this way, a part-based representation has an advantage. Another possibility is that the segment-based representation is not naturally parsed in the way that corresponds to the parts used in the task (Reed & Johnson, 1975). Meaningless parts were used for two reasons. First, using a functional part could allow recognition of the object without the need for a feature search. Second, the meaningless objects did not have particularly functional parts.

In a segment-based representation, the parts will already have been formed. It would seem unlikely that the natural parts would correspond to the ones presented in this study. In that case, the cognitive system would have to take the parts that were already parsed, break these into components, then recombine the components into the "unnatural" parts used in the task. If there were no parts to begin with, in an undifferentiated representation, then whether or not the parts were natural (or figurally good) would not matter. Fresh parts would be formed in either case. The segment based representation would be at a disadvantage compared with a holistic representation. Cave and Kosslyn (1993) also looked at this issue. They concluded that the parts were not important in the task. What they considered of more relevance than the parts was whether the parts used in the task disrupted the invariant information in the representation. Furthermore, comparable parts had been used previously (see Ankrum & Palmer, 1991) during a similar task.

Whether the poor structure of the parts is a real disadvantage can be called into question following recent research on feature creation. It has been shown that features (equivalent to the parts described here) can be created through categorisation, and, presumably, through other tasks in response to the task demands (Schyns & Rodet, 1997; Schyns, Thibault & Goldstone, 1999). However, although these explanations are very plausible, several questions do need to be addressed before they can readily be accepted as an explanation here. First, the authors are vague on exactly how flexible the process is. If the object has already been parsed, it is not clear how this will interact with the task demands. Second, the authors explain the process is determined by the "perceptual experience ....of the individual" (Schyns & Rodet, 1997, p681), but the nature of the experience is not clarified. For example, it is possible that the effects of an immediate perceptual experience, that is the demands of a short-term task, will be different from longer-term experience, that is, familiarity and familiarity of use.

It can be seen that, overall, there is no evidence for either an OSE or an OIE at the complex level, and only weak evidence of an OSE at the simple level. The weakness of the findings could be due to the unlimited exposure times in this experiment. Although long exposures may prevent their occurrence, OSEs have been recorded following presentations of up to two seconds (Davidoff & Donnelly, 1990). The mean study times here, however, were all above two seconds. If superiority effects are an encoding phenomenon rather than a representational one (Tanaka & Farah, 1993), then the conditions of study presented here may not be ideal for their production. If so, then the long study times may be able to counteract the effects of the interaction between the task demands and the representation. This would mean another explanation would have to be sought for the difference between meaningful and meaningless stimuli.

One consideration that could be explored in this light is that there might be a change in format from simple to complex stimuli. The greater change across complexity for the meaningful stimuli could be due to meaning effecting a change in the representation from simple to complex stimuli. The change improves performance with the simple stimuli, suggesting a move to a more part-based representation, and impedes performance with the complex stimuli, suggesting a move to a more holistic representation.

A further explanation can be offered. Research looking at the schemata for simple and complex scenes found that the complex scenes, when briefly presented, were encoded only by the information necessary for the schema (Pezdek, Maki, Valencia-Laver, Whetstone, Stoeckert, & Dougherty, 1988). Irrelevant details were missed out. This could apply here. Although the participants were allowed to study the stimuli for as long as they wanted, this might still result in a "standard" representation, based on the name of the object. When the part to be searched for was not part of this representation, then the task would become more difficult, and performance would drop. This strategy could not occur with the meaningless stimuli, because there was no meaningful name to help in the encoding. It would also be more pronounced in the complex, meaningful stimuli because encoding the complex stimuli would be more difficult than encoding the simple stimuli.

In summary, the study presented here suggests that meaningful objects are represented holistically, whereas meaningless objects are represented by their parts. The form of holistic representation was that of a structural description, with parts and configural information included, but with a stronger importance awarded to the configural information. It also considers whether there may be a change in representational format from simple to complex stimuli. It would be worthwhile, therefore, comparing the amount of configural processing across meaningful and meaningless stimuli, and across simple and complex stimuli.

CHAPTER 5

THE EFFECT OF MEANING ON WHOLE AND PART PROBE TASKS

ABSTRACT

The aim of this study was to use two probe tasks to assess the relative contributions of configural and part information in the representations of meaningful and meaningless objects. A part probe consisted of a small section of the stimulus. A whole probe contained the same small section within an entire stimulus. If configural information is of greater importance, then performance should be better when the whole probe is presented, because this information is absent when only a part is presented. If part information is of prime importance, then performance should be better in the second condition. In this situation, all the redundant information has been removed, leaving just the diagnostic information. Two types of part were used for the probe. One was "figurally good" and the other was not. A whole stimulus was presented for 275 ms, followed by two probes. The participant had to decide which of the two probes matched the initial stimulus. In half the trials the probes were whole probes, and in the other half they were part probes. A better performance was produced in the whole probe task relative to the part probe task; this was found in both meaningful and meaningless, and both simple and complex stimuli. The type of part affected the meaningful stimuli, but not the meaningless stimuli. The complex, meaningful stimuli were particularly affected. The findings can be explained within the framework of a hierarchy of representations. Representations of meaningful objects have a better structure than representations of meaningless objects, in terms of being parsed into "good" parts, with strong configural information about the relations among the parts. This is more noticeable in the complex, meaningful stimuli.

CHAPTER 5

INTRODUCTION

The studies described in the previous chapters demonstrated that a difference exists between the representations of meaningful and meaningless stimuli. One candidate to explain the difference lies with the role of the configural information within the representation. This issue was addressed directly in this experiment. Specifically, the study examined whether the difference found between meaningful and meaningless stimuli could lie within the availability and use of configural and part information. The idea that different representations could have different emphases on configural and part information has been put forward before (e.g. Tanaka & Farah, 1993). This leads to a concept of the representation lying on either a continuum or a hierarchy, with very segment-based representations at one end, and very configurally-based, or holistic, representations at the other.

The idea of a hierarchy of representation was clarified by Navon (1977) who suggested that a scene or object was "decomposed" by finer and finer analysis. The higher levels of the hierarchy were more global in their properties. An advantage was found for the higher levels in terms of more accurate and faster responses to the global information than to the local information (see Kimchi, 1992, for review). Global information also interfered with local information more than local information interfered with global information. Navon suggested that the higher levels were processed first, an effect known as global precedence. The global advantage was dependent upon several factors for expression (see Kimchi, 1992 for more details); this suggested that the order of processing did not necessarily determine the optimum representation for a given task. Several alternatives to this idea have been suggested. One was that elemental features were built up into the whole representation (e.g. Marr, 1982), or even that the representation was developed from the middle out (Kinchla & Wolfe, 1978). Palmer (1977) suggested that the first representation was an unstructured whole. This was parsed into parts, followed by the parts being combined into a structured whole. The structured representation, therefore, contained well-defined parts and their arrangement.

There was some debate about whether higher levels corresponded to holistic representations and lower levels corresponded to parts (Kimchi, 1994). Kimchi stated that configural properties facilitated classification tasks at both global and local levels, and when configural properties were paramount for the task, neither a global nor a local advantage would emerge. When non-configural properties were paramount, then a global advantage occurred. Although this suggested that the global and holistic properties were separate, it did give some support to the idea that the two properties were linked through configural information.

Research within other fields supported the idea of configural information relating to a continuum of representation. Some researchers (e.g. Tanaka & Sengco, 1997) considered that the one holistic unit may actually be parts, but that the parts are strongly dependent upon each other. The dependence is postulated as arising from a strong role played by the configural information, which encompasses the relationship among the parts of an object. Tanaka and Sengco claim that in a holistic representation both the parts and the relationship among the parts are important. Donnelly and Davidoff (1999) considered that the crucial difference between holistic and piecemeal processing occurred between the relative importance of the configural information and the part information. A high importance for configural information signified a holistic process, whereas a high importance for part-based information signified a piecemeal process. It would therefore appear worthwhile exploring the different levels of representation that may exist for meaningful and meaningless objects; in particular, the contribution of configural information during performance of the tasks may be worth investigating in more detail.

Tanaka & Farah (1993) claimed that holistic representations were one, single unit. As such, if a section of the whole had to be processed in any way, then performance would be more successful if the section remained part of the whole than if it was removed and isolated. Conversely, if a representation were segment-based, the section of the whole would be easier to process when it was isolated than when it was contained in the whole.

This paradigm has been used to compare the use of part-based and configural information during the processing of objects. If configural information is predominant, then providing this information should improve performance compared with a situation when the information is absent. This can be assessed by comparing performance in a recognition or matching task when the component is presented alone (known as a part probe) and when it is presented within the context of a whole object (a whole probe). An increase in performance with the whole probe (a whole probe advantage, or WPA) suggests that configural properties were an important source of information used in the performance of the task.

Davidoff and Donnelly (1990) used a same-different task and a two-alternative forced choice task to look at the representations of faces and chairs. In both tasks, they found a distinct whole probe advantage for both faces and chairs, which was removed when the objects were scrambled. They concluded that both faces and chairs had holistic representations. Ankrum & Palmer (1991) used a similar task, with line drawings of two-dimensional geometric figures as their stimuli. They also found an advantage for the whole probe over the part probe. When the stimulus and whole probes were separated into segments, however, this advantage was removed.

Tanaka and Farah (1993) used a recognition memory task to look at the representation of houses and faces. They found that there was a whole probe advantage for faces, but not for houses. Their conclusion was that only faces were holistic; other objects, houses in particular, were segment-based. Donnelly and Davidoff (1999) attempted to resolve this contradiction by making the tasks and stimuli between the two studies more equivalent. They then concluded that faces were holistic, but that although houses could be processed holistically, they were represented in a part-based fashion.

The contradictory findings across the studies could lie with several sources. First, the studies described used a single source for comparison. This means that not only was there only one category of object, but there was only one level of complexity within each study. If complexity affects the representation, then the effect will be missed. Second, all the objects were compared with faces. If we accept that faces are a strong example of the use of configural processing, and if we further accept that there is a continuum from part-based to configural, then most objects will appear less holistic than faces. For the purpose of this study, we needed to compare meaningful with meaningless objects.

This experiment aimed to extend the findings of Donnelly and Davidoff (1999), and specifically incorporated a comparison of meaningful and meaningless objects. Within these categories, simple and complex objects were used. It also aimed to examine the possibility that differing roles for configural and part-based information could provide at least part of the explanation for the results described in previous chapters. In addition, this study provided another condition. One of the criticisms of the experiment in the last study related to the goodness of the parts during the search task. This study addressed the issue, by having some of the parts with high figural goodness, and some with low figural goodness. Performance can then be compared across these parts.

To summarise, the present study aimed to assess the use of configural information during a probe matching task. To test this we based our methodology on that of Donnelly & Davidoff (1999). An object was presented, followed by two alternative probes. In one condition, this choice consisted of two whole stimuli; in the other condition the choice was between two parts. The task was to say which of the two alternatives matched the original object.

If configural information plays an important role, then performance should be better in the first condition, because configural information is completely absent when just a part is presented. If part-based information is of prime importance during this task, then there will be no advantage in having the configural information present. In fact, it is possible that performance would be better in the second condition. In this situation, all the redundant information has been removed.

Method

Participants

Thirty-two participants were recruited. Twenty-four Psychology undergraduates from the departmental participant pool took part; other students from an introductory psychology course were paid £4 for participating. There were 24 females and 8 males, with ages ranging from 17 to 30 (mean age 20). All had normal or corrected-to-normal vision. None of the participants had seen the stimuli before.

Materials

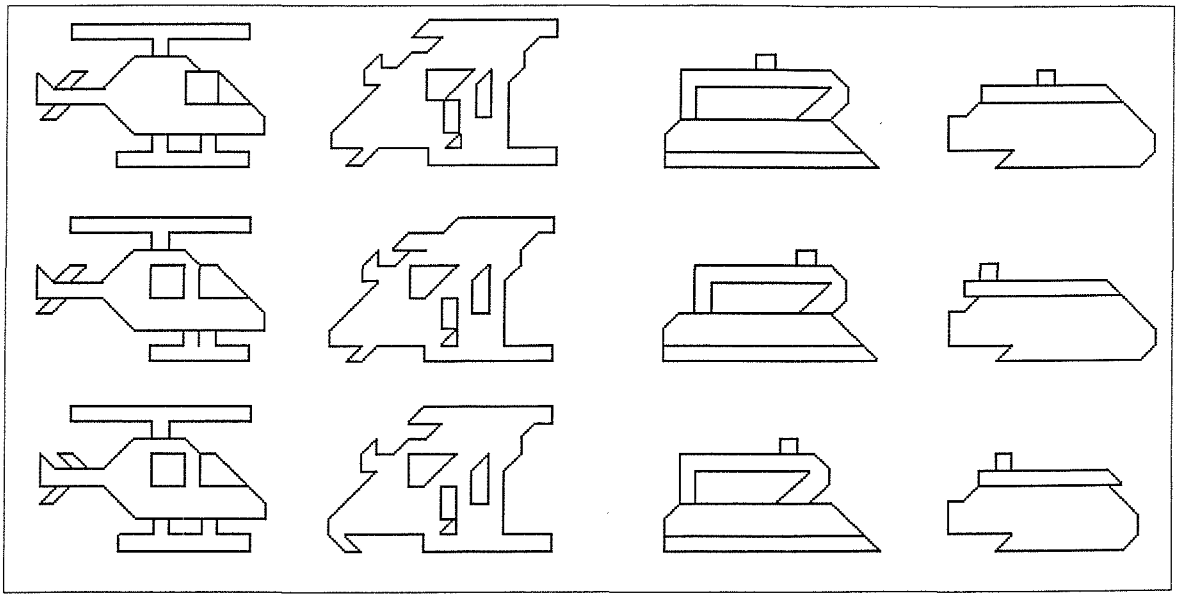

Figure 5.01: Examples of the Different Sets of Stimuli

The columns show, from the left, complex, meaningful; complex, meaningless; simple, meaningful; and simple, meaningless stimuli. The different rows show the different versions of a stimulus set.

The materials were taken from six meaningful stimuli and six meaningless stimuli used in the previous experiments. Three of the six objects were simple, and three were complex. These objects formed the prototypes. The prototype was modified to produce a set of three stimuli (see figure 5.01 and appendix 9 for the full set of whole probes). Each modification affected one small area of the stimulus, and was kept as similar as possible across the meaningful and meaningless stimuli, and across the complex and simple stimuli. The modifications were selected from different areas of the object. These stimuli were used for both the initial presentation, and the whole probe.

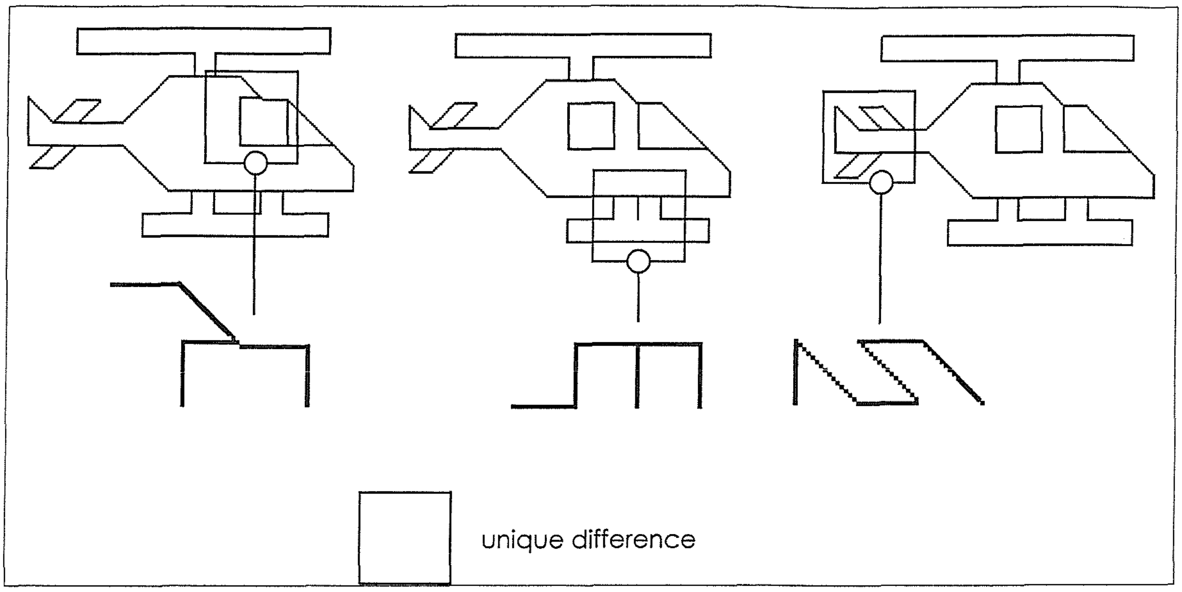

Figure 5.02: Examples of the Differences among the Three Modified Stimuli in a Stimulus Set.

The three whole probes of a stimulus set are shown above, with the associated part probe shown below.

The unique section of the object was then used to produce a part probe. In each case, the part probe consisted of six 0.5 cm lines, joined together (see figure 5.02 for examples, and appendix 8 for the full set of part probes). To keep the part probes equivalent across the stimuli, the same type of part was used in all stimuli, and kept as similar as possible. In the previous study, one criticism was that the parts themselves appeared quite abstract, and may not relate to the way the parts would be parsed naturally. In this study, one of the parts (the abstract part) was similar in nature to the ones used in the previous study (see the left part in figure 5.02). This part had low figural goodness. Another of the parts (the concrete part), however, incorporated a more structured area of the stimulus (see the right part in figure 5.02; see appendix 10 for the relationship of the part probes and the whole probes). This part had high figural goodness.

As before, the stimuli were presented using a Power Macintosh 8600/250 computer with a 17--inch Apple Vision monitor (running at 1024 x 768 resolution) and the experiment was designed and administered using the commercial software Superlab, version 1.4 (Cedrus Corporation, 1991).

Procedure

Each participant performed two search tasks. The whole stimulus was shown, followed by two probes. The task was to say which of the two probes matched the initial stimulus. Although the task was essentially the same in both tasks, the tasks differed by the type of probe that was used. In one task, the participant saw two stimuli from the same stimulus set; these were known as whole probes. In this case, they had to say which probe was identical to the first stimulus. In the other task, only the difference between the two stimuli was shown; these were known as part probes. The task here was to say which part had been present in the initial stimulus. The difference between the two tasks was the information presented. In the part probe task, just the necessary part information was presented, and in the whole probe task this information was presented within the context of the configural information. The order of presentation of these two tasks was counterbalanced across participants.

These two tasks were performed with both the meaningful objects, and the meaningless objects, the order of which was also counterbalanced across participants. Each participant performed four tasks; two with the whole probes, one using meaningful and one using meaningless stimuli, and two with the part probes, again one with meaningful and one with meaningless stimulus. The participant was encouraged to take a short break between tasks.

For all the tasks, the word "ready" was presented on the screen, and remained there until the participant indicated they were ready by pressing the space bar, with the fingers of their non-dominant hand. After a pause of 500 ms one of the whole stimuli was presented. This initial stimulus was presented for 275 ms. (see Donnelly & Davidoff, 1999). After a pause of 500 ms, two probes were presented, from the same set, which remained on the screen until the participant responded. The participant indicated their choice by pressing either the key labelled "L" (signifying the probe on the left; the "b" key of a standard keyboard) or the key labelled "R" (signifying the probe on the right, the "n" key of a standard keyboard), using the first two fingers of their dominant hand (see figure 5.03).

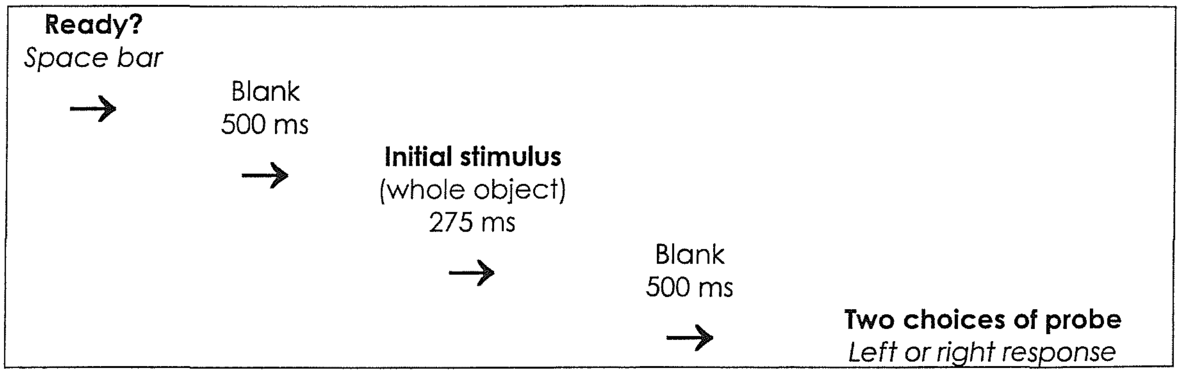

Figure 5.03: Events within a Trial of the Experiment

The events visible to the participants are presented in bold. The first stimulus is a whole object. The second stimulus consists of two probes; one was present in the whole object, and one was not. In the whole probe task, these are also the whole stimuli. In the part probe tasks they are a small part of the object. The items under each event are either the length of presentation or the response required from the participant.

Each of the four tasks consisted of 72 trials. There were six sets of stimuli for a task (three simple and three complex), with three stimuli in a set. All the stimuli were used as the initial stimulus, and also produced two probes; a whole probe, and a part probe. Every probe was paired with the other two of the same type in the set, appearing on the left in half the trials, and on the right in the other half. This made a total of six probe pairs per task. Each of the probe pairs was presented twice, once with one probe as the correct choice, and once with the other probe as the correct choice.

Before the main four tasks, training was given to ensure the participant understood the task, understood the types of features and information that would be presented and was familiar with the use of the keys. The training lasted approximately twenty minutes. The tasks were identical to those in the main task, but used completely different sets of stimuli, and had fewer trials.

Results

General Analysis

The data of interest were the error rates, and these were analysed using an ANOVA to assess the type of stimulus and complexity in the two probe tasks. The response times in the tasks were also checked to eliminate a speed-accuracy trade-off as an explanation for any results. Next, the performances across the two types of part were compared. Finally, the whole probe advantage was quantified, and compared across the conditions of type of stimulus and complexity.

First, a general analysis was performed using the logs of the percentage error data (see figure 5.04). Complexity (simple and complex), stimulus type (meaningful and meaningless) and probe type (whole and part) were within-subject factors. The presentation order of the probe tasks, and the presentation order of the type of stimuli, were between-subject factors. There was a main effect of stimulus type ($F_{1,28} = 9.95$, $p < .01$), a main effect of complexity ($F_{1,28} = 68.74$, $p < .001$), and a main effect of probe type ($F_{1,28} = 238.8$, $p < .001$). There was also a significant interaction between the stimulus type, the probe type and the complexity ($F_{1,28} = 4.47$, $p < .05$). There was no effect of probe order or stimulus type order (both $F < 1$).

The results of this analysis indicated that the performance differed across the two probe tasks; performance was better in the whole probe task.

The performances did not show a consistent pattern across stimulus type and complexity. The findings were broken down for a more detailed analysis.

Analysis of the Whole Probe Task

Complexity and stimulus type (meaningful or meaningless) were within subject factors in the analysis. There was no main effect of stimulus type ($F_{1,31} = 2.42$, $p > .1$). Closer analysis showed that although this was the case in the simple stimuli ($F < 1$), in the complex stimuli more errors were made with the meaningless than the meaningful stimuli ($F_{1,31} = 5.05$, $p < .05$). There was also a main effect of complexity ($F_{1,31} = 48.93$, $p < .001$), where more errors were made with the complex stimuli than the simple stimuli for both the meaningful ($F_{1,31} = 11.54$, $p < .01$) and the meaningless ($F_{1,31} = 31.08$, $p < .001$) stimuli. There was no interaction between stimulus type and complexity ($F < 1$).

Analysis of the Part Probe Task

The analyses were then repeated with the data from the part probe task. It was noted that, in the part probe task, some of the error rates were high (over 40%). Therefore, the error rates in this task were first analysed using a one-sample t-test to exclude a floor effect. All the values were significantly below chance (simple, meaningful: $t_{31} = 15.68$, $p < .001$; complex, meaningful: $t_{31} = 5.08$, $p < .001$; simple, meaningless: $t_{31} = 6.78$, $p < .001$; complex, meaningless: $t_{31} = 4.92$, $p < .001$).

The error rates were then analysed using an ANOVA as before. There was a main effect of stimulus type ($F_{1,31} = 7.95$, $p < .01$), a main effect of complexity ($F_{1,31} = 23.53$, $p < .001$), and a significant interaction between stimulus type and complexity ($F_{1,31} = 14.18$, $p < .01$). There were more errors made in the meaningless condition. In contrast with the whole probes, the effect of stimulus type was found to be significant in the simple stimuli ($F_{1,31}$

= 14.44, $p < .01$), but not for the complex stimuli ($F > 1$). More errors were made with the complex stimuli than the simple stimuli. This was found in both the meaningful condition ($F_{1,31} = 42.83$, $p < .001$), and in the meaningless condition ($F_{1,31} = 4.76$, $p < .05$).

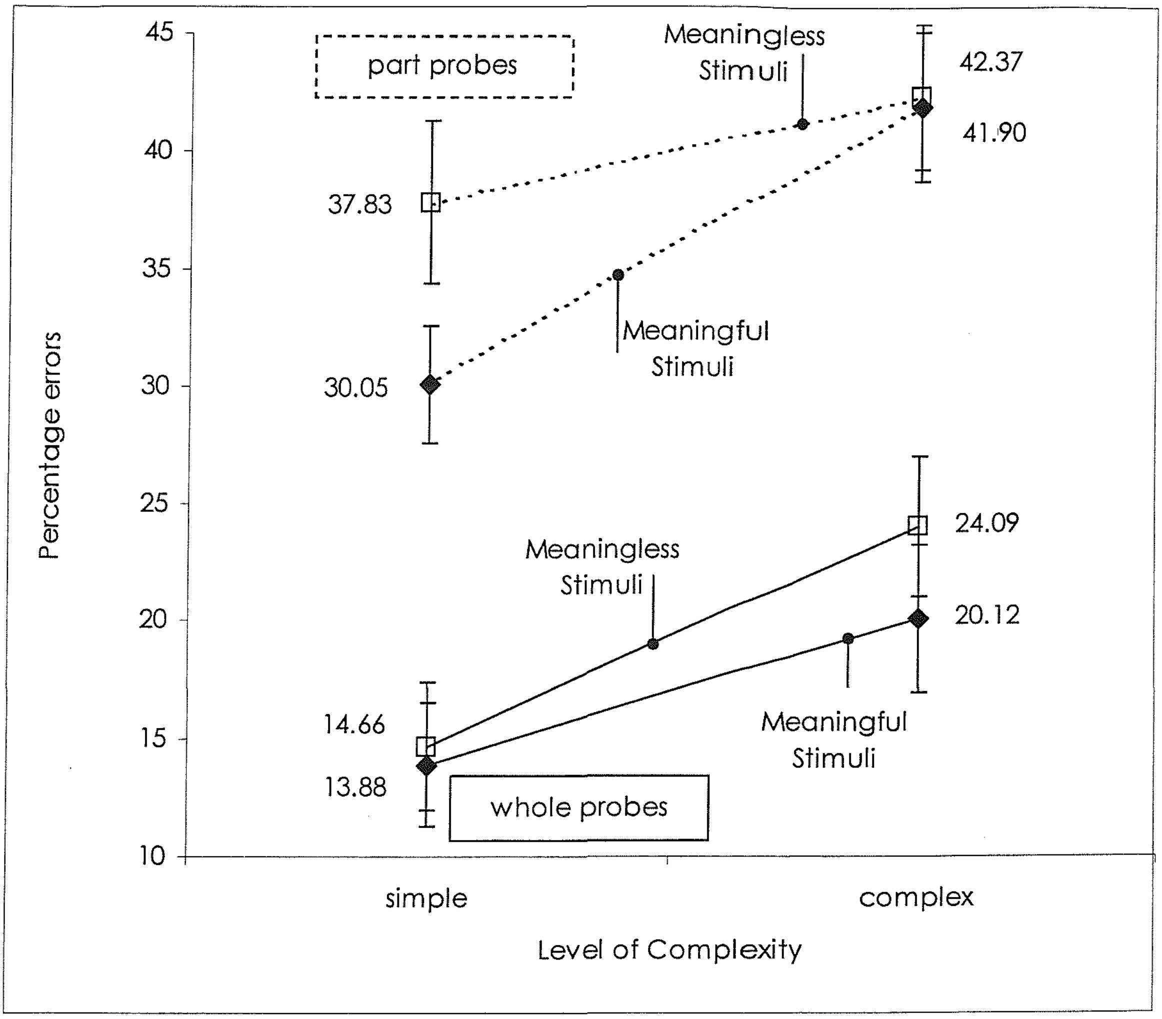

Figure 5.04 Error Rates for the Whole and Part Probes.

There was a better performance with the whole probes than with the part probes. The complex stimuli show higher error rates than the simple stimuli. The pattern of errors differed across the two probe tasks. The mean value for each point is shown, along with standard error bars.

Analysis of the Response Times

Although the error rates were below chance, they were still high in the part probe tasks. Therefore the response times were examined, to check whether a speed-accuracy trade-off had occurred. Only the times of the correct responses were used.

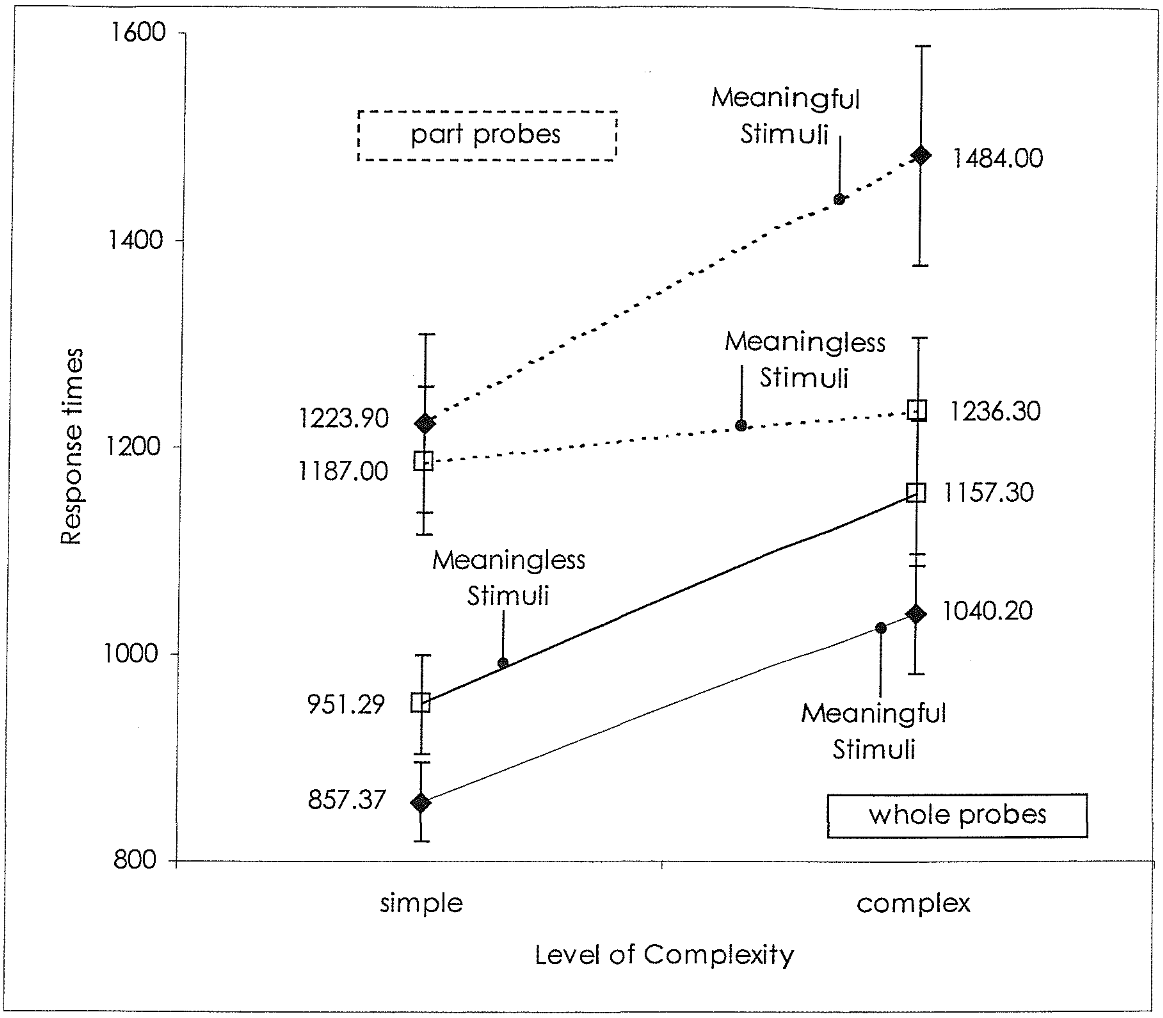

Figure 5.05 Response Times For Whole and Part Probes.

There was a better performance with the whole probes than with the part probes. The response times in the whole probe task support the findings with the error rates. The response times with the part probes tasks were lower that expected in the meaningless condition. The mean value for each point is shown, along with standard error bars.

For the whole probe task, there was a main effect of stimulus type ($F_{1,31} = 13.38$, $p < .01$). Meaningless stimuli produced longer response times in both simple ($F_{1,31} = 9.63$, $p < .01$) and complex ($F_{1,31} = 5.22$, $p < .05$) conditions. There was also a main effect of complexity ($F_{1,31} = 27.26$, $p <$ .001), where more time was taken to produce a response with the complex stimuli than the simple stimuli for meaningful ($F_{1,31} = 24.66, p < .001$) and meaningless ($F_{1,31} = 12.86, p < .01$) stimuli. There was no interaction between type and complexity ($F < 1$). The response times support the findings with the error rates.

In the part probe task, the response times showed a main effect of stimulus type ($F_{1,31} = 6.64, p < .05$), a main effect of complexity ($F_{1,31} = 20.22, p < .001$), and a significant interaction between stimulus type and complexity ($F_{1,31} = 6.79, p < .05$). Longer response times were found with the meaningful stimuli than the meaningless. This was significant for the complex stimuli ($F_{1,31} = 15.03, p < .01$), but not for the simple stimuli ($F_{1,31} = 2.14, p > .1$). More time was taken to produce a response with the complex stimuli than the simple stimuli in the meaningful condition ($F_{1,31} = 13.14, p < .01$), but not in the meaningless condition ($F < 1$). The findings do not support the error rates, and are considered at greater length in the discussion.

Analysis of the Type of Part

The error rate data were broken down, and the errors calculated for the abstract parts and the concrete parts. The data from one of the participants had to be removed because she had no valid data for the abstract parts. To assess if the nature of the part had any effect on performance of the task an ANOVA was performed using the type of probe, the type of stimulus, the complexity and the type of part as within-subject factors. There was no main effect of the type of part ($F < 1$), but there were significant interactions. The type of part interacted with the probe type and the complexity ($F_{1,30} = 15.5, p < .001$); the type of stimulus and the complexity ($F_{1,30} = 8.12, p < .01$); and the type of probe, the type of stimulus and the complexity ($F_{1,30} = 6.38, p < .05$).

When separate ANOVAs were performed on the whole probe data and the part probe data, the three way interaction was maintained for the whole probes ($F_{1,30} = 10.23, p < .01$), but not for the part probes ($F < 1$). Four paired t-tests compared the data across the concrete and abstract parts in the whole probe task. Only one pair was significant; the complex, meaningful stimuli showed lower error rates with the concrete parts than with the abstract parts ($t_{30} = 2.87$, $p<.01$; p=.012 was taken as the value for alpha using a Bon Ferroni correction). There was a borderline finding in the simple, meaningful stimuli ($t_{30} = 2.59$, $p=.015$). Here, error rates were lower with the abstract parts.

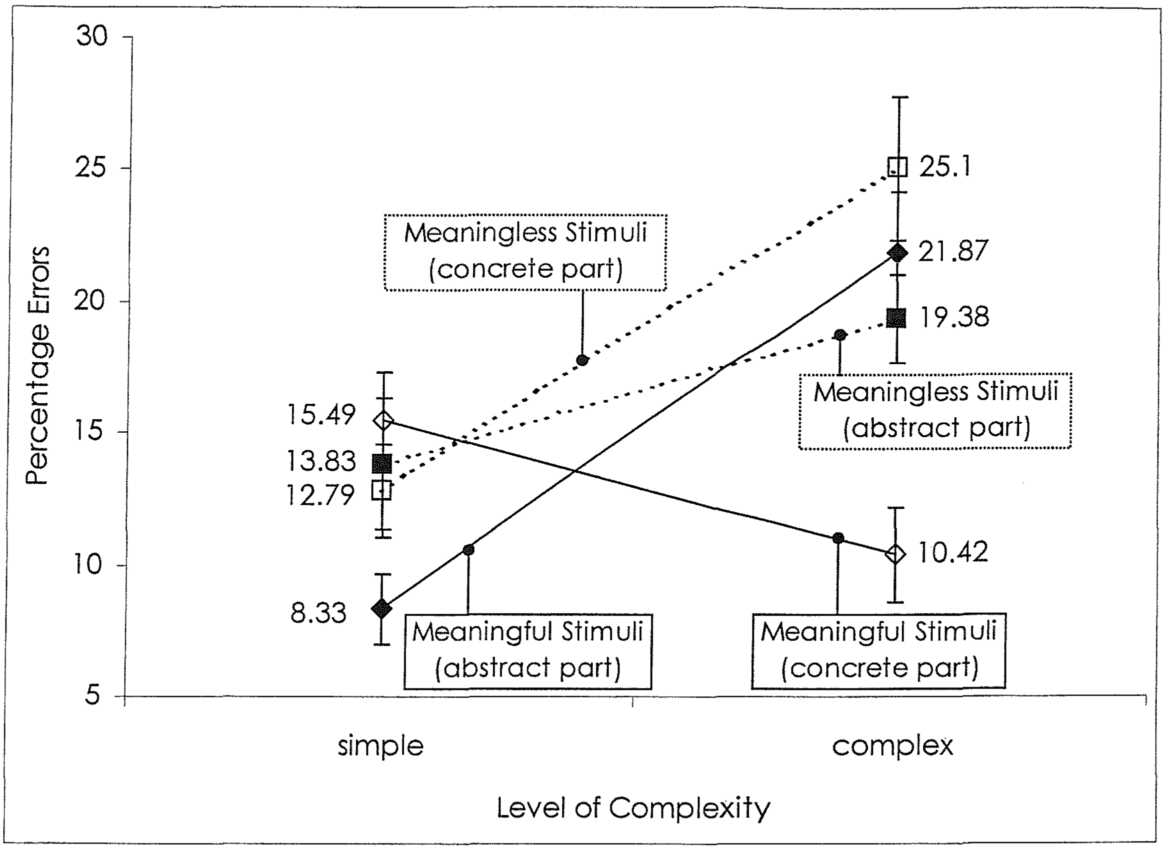

Figure 5.06: Error Rates for Different Parts used as Probes in the Whole Probe Task.

The effect of the two types of part is compared across the meaningful and meaningless stimuli. The concrete parts are figurally good, whereas the abstract parts are figurally poor. The mean value for each point is shown, along with standard error bars.

Given this difference, the ANOVA examining the effect of complexity and stimulus type on the error rates in the whole probe task was repeated, once with the results from abstract parts and once with the results from the concrete parts. The pattern of results for the two types of parts was different (see figure 5.06).

For the abstract parts, there was no effect of type of stimulus ($F_{1,30}$ = 1.43, p> .1), a main effect of complexity ($F_{1,30}$ = 24.71, p< .001), and no interaction (F< 1). The performance with the complex stimuli showed higher error rates than with the simple stimuli for both stimulus types (meaningful: $F_{1,30}$ = 14.06, p< .01; meaningless: $F_{1,30}$ = 9.39, p< .01).

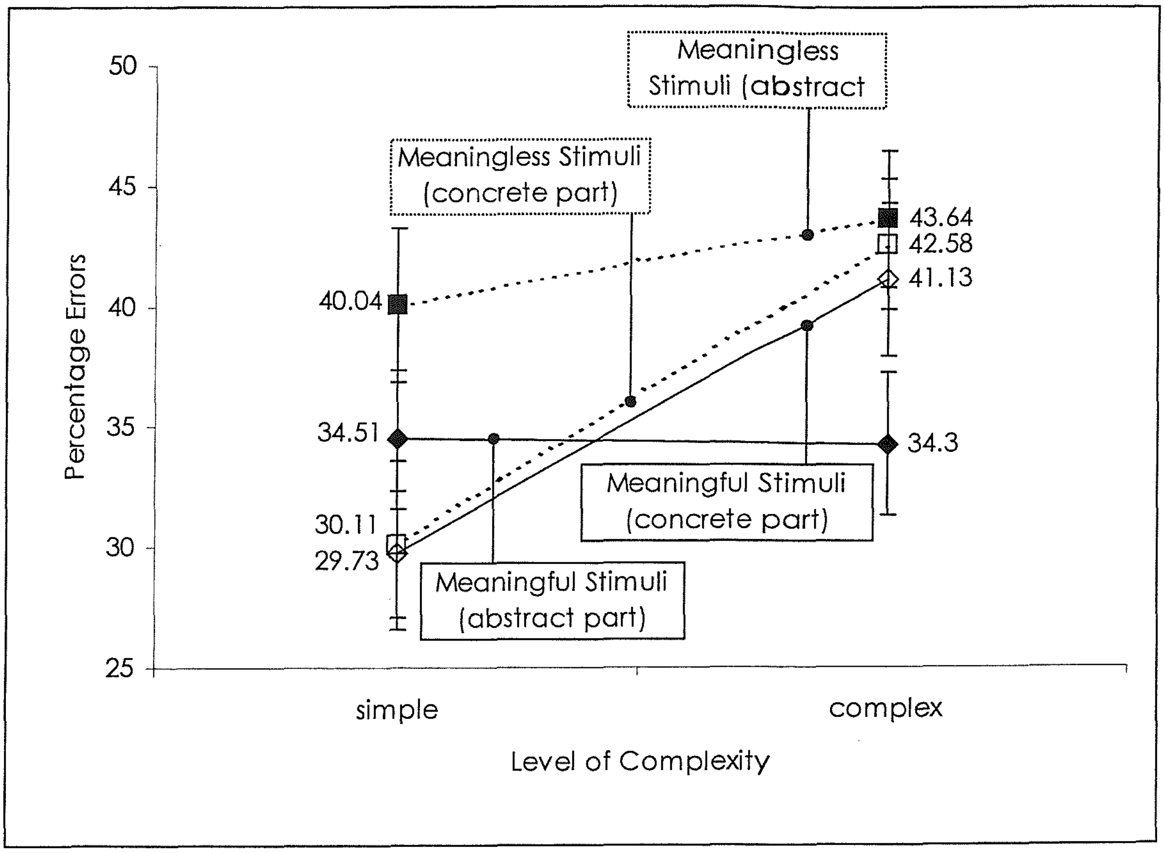

Figure 5.07: Error Rates for Different Parts used as Probes in the Part Probe Task.

The effect of the two types of part is compared across the meaningful and meaningless stimuli. The concrete parts are figurally good, whereas the abstract parts are figurally poor. The mean value for each point is shown, along with standard error bars..

For the concrete parts, there was a main effect of type of stimulus ($F_{1,30}$ = 6.9, p< .05), no main effect of complexity ($F_{1,30}$ = 1.04, p> .1), and a significant interaction ($F_{1,30}$ = 16.65, p< .001). The errors with the meaningless stimuli were higher than with the meaningful stimuli in the complex condition ($F_{1,30} = 27.8$, $p < .001$), but there was no difference in the simple condition ($F < 1$). The errors in the complex condition were higher than those in the simple condition for the meaningless stimuli ($F_{1,30} = 20.19$, $p < .001$), but this difference was only borderline with the meaningful stimuli ($F_{1,30} = 3.51$, $p < .071$).

Closer examination of the data from the part probe task (see figure 5.07) revealed that for the abstract parts, there was a main effect of type of stimulus ($F_{1,30} = 5.97$, $p < .05$), no effect of complexity ($F < 1$), and no interaction ($F < 1$). The error rates were lower in the meaningful stimuli for the complex condition (meaningful: $F_{1,30} = 6.06$, $p < .05$) but not for the simple condition ($F_{1,30} = 1.45$, $p > .1$).

For the concrete parts, there was no effect of type of stimulus ($F < 1$), a main effect of complexity ($F_{1,30} = 15.87$, $p < .001$), and no interaction ($F < 1$). The errors with the complex stimuli were higher than with the simple stimuli for the meaningless stimuli ($F_{1,30} = 11.95$, $p < .01$), but this difference was only borderline with the meaningful stimuli ($F_{1,30} = 3.62$, $p < .067$).

Comparison of the Whole Probe Advantage across Conditions

The percentage of correct responses from each individual in the whole probe task was divided by the percentage of correct responses produced by the same individual in the part probe task. This ratio of whole probe data to part probe data produced a measure of the whole probe advantage, or WPA. This value provided a quantitative measure of the degree of the advantage in the different conditions. A WPA of more than 1 would indicate an advantage with for the whole probes; a WPA of less than 1 would indicate an advantage for the part probes. A WPA of 1, or very close to 1, would indicate equivalent performance with the two types of probe.

A ratio was used rather than subtracting one set of data from the other, because the ratio better reflects the underlying relationship between the whole and the parts. An additive effect, where a subtraction would be appropriate, makes the assumption that the two factors, the part information and the configural information, are independent. If the part information is removed, then what is left is the advantage created by the configural information. However, if the part information is removed, then the configural information would no longer exist; one cannot specify the relationship among the parts if there are no parts. A multiplicative effect makes the assumption that the configural information interacts with the part information, and that independence of the two effects is not maintained (see Dosher & Rosedale, 1997). In this case, a ratio is more suitable.

The WPA was calculated as described above for both the abstract and the concrete parts. First, a one-sample t-test assessed whether there was a whole probe advantage present, that is, that the values were significantly above 1. A significant WPA was found in all cases (meaningful stimuli: simple, abstract $t_{30} = 6.32$, $p<.001$; simple, concrete $t_{30} = 5.00$, $p<.001$; complex, abstract $t_{30} = 3.79$, $p<.001$; complex, concrete $t_{30} = 5.82$, $p<.001$; meaningless stimuli: simple, abstract $t_{30} = 5.50$, $p<.001$; simple, concrete $t_{30} = 4.36$, $p<.001$; complex, abstract $t_{30} = 4.73$, $p<.001$; complex, concrete $t_{30} = 4.61$, $p<.001$).

The pattern of results for the two types of parts was again different (see Figure 5.08). These patterns were analysed with ANOVAs using type of stimulus, complexity, and type of part as within subject factors. There was no main effect of the type of stimulus ($F< 1$), complexity ($F_{1,30} = 1.29$, $p> .1$), or type of part ($F< 1$). There were significant interactions between the type of stimulus and the type of part ($F_{1,30} = 4.3$, $p< .05$), the complexity and the type of part ($F_{1,30} = 5.58$, $p< .05$), and the type of stimulus, the complexity, and the type of part ($F_{1,30} = 4.91$, $p< .05$). The two types of part were separated, and the pattern of results analysed for each, using an ANOVA with the type of stimulus and the complexity as within-subject factors.

For the abstract parts, there was a main effect of type of stimulus ($F_{1,30} = 4.46$, $p< .05$), no effect of complexity ($F< 1$), and no interaction ($F_{1,30} =$

1.55, $p > .1$). There was no difference in the WPA between the meaningless and the meaningful stimuli in the simple condition ($F < 1$), but the meaningless stimuli showed a higher WPA in the complex condition ($F_{1,30} = 4.97, p < .05$). The WPA with the complex stimuli was lower than with the simple stimuli for the meaningful stimuli ($F_{1,30} = 5.58, p < .05$); but not for the meaningless stimuli ($F < 1$).

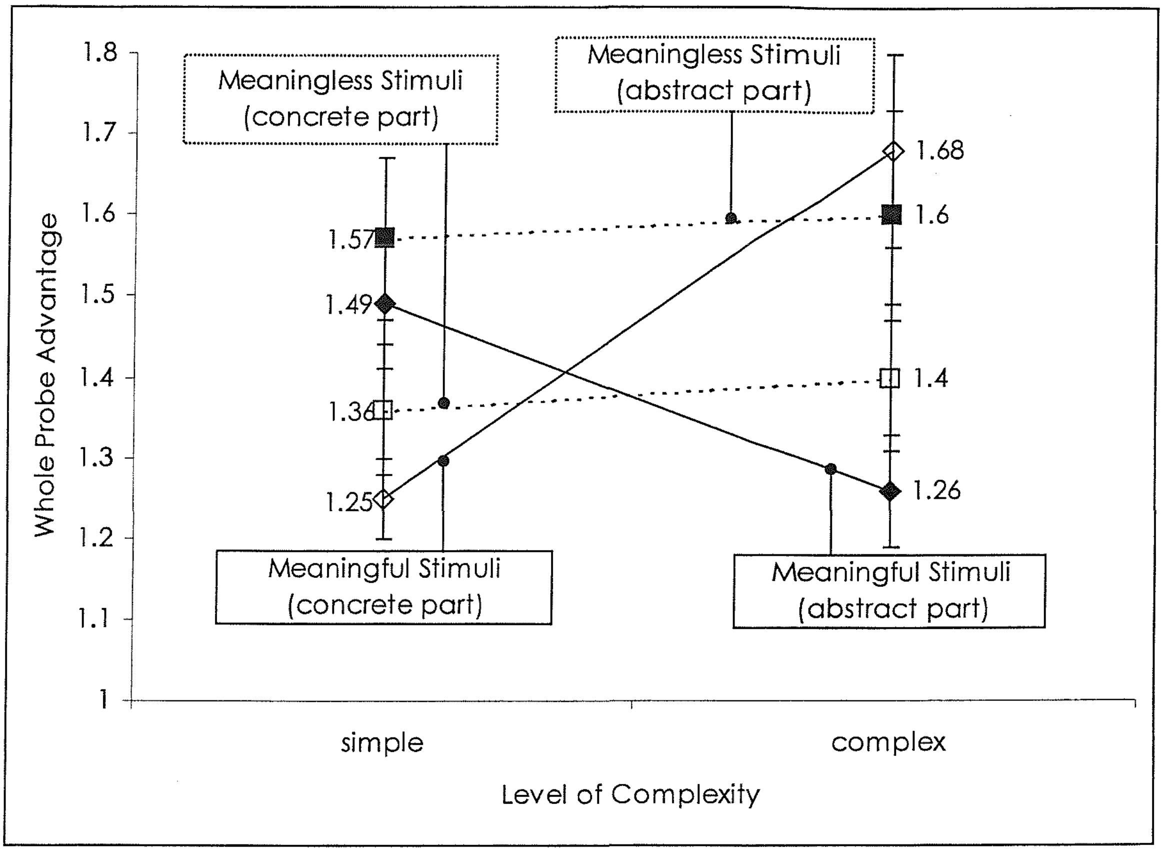

Figure 5.08: The Whole Probe Advantage (WPA).

The whole probe advantage for the meaningless stimuli was relatively unaffected by the type of part used as a probe, and by the complexity. This was not the case in the meaningful stimuli, where the type of part affected the WPA, and this interacted with the complexity. The mean value for each point is shown, along with standard error bars.

For the concrete parts, there was no main effect of type of stimulus ($F < 1$), a main effect of complexity ($F_{1,30} = 7.63, p < .05$), and a significant interaction ($F_{1,30} = 7.04$, $p < .05$). The WPA in the meaningless stimuli was lower than with the meaningful stimuli in the complex condition ($F_{1,30} = 4.93$, $p < .05$), but there was no difference in the simple condition ($F_{1,30} = 1.09$, $p > .1$). The WPA in the meaningful stimuli was lower in the simple condition than the complex condition ($F_{1,30} = 15.77$, $p < .001$). There was no effect of complexity in the meaningless stimuli ($F < 1$).

The two types of stimulus were also separated, and the pattern of results analysed for each, using an ANOVA with the type of part and the complexity as within-subject factors. For the meaningful stimuli, there was no main effect of complexity ($F < 1$), and no main effect of type of part ($F < 1$), but a significant interaction ($F_{1,30} = 21.68$, $p < .001$). For both types of part there was a significant difference in the WPA between the simple and complex conditions. With the abstract parts, the WPA was higher in the simple stimuli than in the complex stimuli ($F_{1,30} = 9.94$, $p < .01$), but with the concrete parts the complex stimuli showed the higher WPA ($F_{1,30} = 19.49$, $p < .001$). The WPA with the abstract parts was lower than with the concrete parts for the complex stimuli ($F_{1,30} = 9.68$, $p < .01$); but WPA with the concrete parts was lower than with the abstract parts for the simple stimuli ($F_{1,30} = 6.48$, $p < .05$).

For the meaningless parts, there was a main effect of type of part ($F_{1,30} = 4.94$, $p < .05$), no main effect of complexity ($F < 1$) and no interaction (($F < 1$)). When the effect of the type of part was examined more closely in the simple and complex conditions, the effect was not significant (Simple: $F_{1,30} = 2.31$, $p > .1$; complex: $F_{1,30} = 2.03$, $p > .1$).

In summary, the error rates were lower in the whole probe task than in the part probe task. In the whole probe task, although the pattern across the complexity was the same, there were fewer errors in the meaningful stimuli than the meaningless stimuli in the complex condition. In the part probe task both the meaningful and the meaningless stimuli produced the same number of errors in the complex stimuli, but the meaningless stimuli produced more errors in the simple condition. The response times suggested that the task with the meaningless stimuli might have been so difficult that the participants were giving up, and guessing. The type of part used for the probe had little effect on performance in the meaningless stimuli, but did have an effect on the meaningful stimuli, particularly in the complex condition.

DISCUSSION

This study looked at the relative roles of configural and part-based information in the representation of an object. Participants performed a two-alternative probe task, under two conditions. In both conditions, a whole stimulus was shown. In one condition, two isolated parts serving as probes followed this initial presentation. In the second condition, the parts were shown in the context of a whole stimulus. The role of configural information can be assessed by the difference in performance across the whole and part probe tasks. If configural information was relatively more important during this task, then the condition with the whole probes should show a better performance, because the configural information is not present in the part probes. This was compared across meaningful and meaningless stimuli, and across simple and complex conditions. Performance was also compared across figurally good and figurally poor parts in the probes.

Overall, performance in the whole probe task, in terms of the number of errors produced, was better than in the part probe task. This suggests that configural information was useful for all the four types of stimuli during the task. In the whole probe task the task difficulty in the simple stimuli was equivalent across stimulus types, but in the complex condition, the task was harder with the meaningless stimuli.

In the part probe task, again, the error rates were lower in the meaningful than in the meaningless stimuli. In this case, however, although the error rates were comparable, but high, at the complex level, at the simple level there were fewer errors with the meaningful stimuli. When the response times were examined, they were found to be shorter in the meaningless than in the meaningful stimuli. Taken with the pattern of errors, this suggested that the part probe task was more difficult than the whole probe task. In the meaningless condition, the task might have seemed so difficult that the participants were just giving up and guessing. In the meaningful condition, the response times suggested that they were attempting the task, although the error rates showed that this was not always successful, particularly in the complex condition. This implies that configural information helps performance with both meaningful and meaningless stimuli. In the meaningless stimuli, the task became almost impossible without this information.

The type of part, abstract or concrete, was taken as a measure of figural goodness of the part. Figural goodness influenced the performance in the tasks. In the whole probe task, figural goodness of the part had little effect on the meaningless stimuli. In the meaningful stimuli, however, when the concrete part was used, performance in the complex condition was as good as, if not better than, in the simple condition. In the part probe task, figural goodness of the part reduced the number of errors in the meaningless stimuli in the simple condition. The concrete parts had little effect on the meaningful stimuli. Abstract parts, however, reduced errors in the meaningful, complex stimuli so that performance was equivalent to the simple stimuli. This suggests that figural goodness was actually impairing performance in the complex, meaningful stimuli. Figural goodness, therefore, appears more important in the meaningful than in the meaningless stimuli. It also appears more important in the complex, meaningful stimuli than in the simple, meaningful stimuli.

The differences in performance across simple and complex stimuli can be assessed by a direct comparison, in which better performance with the meaningful objects is an object superiority effect (OSE) and a better performance with the meaningless objects is an object inferiority effect (OIE). In the whole probe task, only the concrete parts produced an OSE, and only in the complex condition. In the part probe task, only the abstract parts produced an OSE, and, again, only in the complex condition. In no conditions was an OIE produced.

The whole probe advantage across the stimuli was examined more formally. In all cases, the configural information aided performance in the task. In the meaningless stimuli, neither complexity nor figural goodness of the part had a great deal of effect on the WPA. In the simple condition, the advantage was very similar across meaningful and meaningless stimuli. The WPA appeared slightly higher in the meaningless stimuli, although this was probably a function of the very poor performance with meaningless stimuli in the part probe task. The WPA with the meaningless stimuli was the same in the complex stimuli as in the simple stimuli. In the meaningful stimuli, however, figural goodness had a role to play. The WPA with the concrete parts in the complex condition was considerably higher than with the abstract parts. This was partly due to a better performance in the whole probes with the concrete parts, but also due to a better performance in the part probes with the abstract parts.

Configural information helped in the performance of the task with both the meaningful and the meaningless stimuli. Without the configural information, the task was extremely difficult in the meaningless task. One explanation is that the short presentation time made it very difficult to encode the stimulus. Having a poor representation to search for a part would lead to a poor performance. When the whole stimulus was presented, however, this could provide a cue to the information needed to complete the task. If the representations formed were also well-structured, in terms of "good" parts and strong configural properties, this would enhance any cueing mechanism.

Dosher and Rosedale (1997) explain how this could work. First, one had to assume that the memory processes form a continuum. At one end was a configurally-based mechanism, leading, at the other end to a part-based mechanism. In the latter condition, where the parts are independent, the effect of several cues is additive. In the former condition, the effect is multiplicative; the relationship among the parts can also act to enhance the cue. However, the more holistic mechanism can only be used in certain conditions; that is, when there is a strong relationship among the parts due to the strong organisational properties.

Configural information appeared more important in the complex than in the simple meaningful stimuli. In the previous study of this thesis, one possible explanation put forward to explain the results was that the type of representation could change across complexities. These findings would support that idea. One way of representing more information at a time, under difficult conditions such as a short presentation time, is to "chunk" the material (Miller, 1956). It could be that configural information allows precisely this in meaningful stimuli, with a subsequent increase in importance for the configural information. This would be particularly noticeable in the complex stimuli, where more information has to be represented relative to the simple stimuli.

Figural goodness is less influential in the performance with the meaningless stimuli than with the meaningful stimuli. It appears particularly important in the meaningful, complex stimuli. The WPA supports this. This was equivalent across all conditions with the meaningless stimuli. However, figural goodness affects the WPA in the meaningful stimuli. In the simple condition, the pattern show by the WPA across featural goodness was similar to that in the meaningless stimuli. However, in contrast, with the complex, meaningful stimuli, there was an increase in WPA when the part was figurally good. This is consistent with the idea that complex, meaningful stimuli are "special" as far their representations go. One possible explanation for this is as follows. When the stimuli are simple, the tasks are relatively easy, and the same performance is produce for both meaningful and meaningless stimuli. When the stimuli are complex, the tasks become harder. The representation of the meaningful stimuli can compensate for the increase in difficulty through the use of configural information; this is mediated by the goodness of the parts. The representation of the meaningless stimuli allows no such compensation, and performance is poor. If either the configural information or the goodness of the parts is removed, then performance in the meaningful stimuli becomes the same as that of the meaningless stimuli. If both are removed, then the simple, meaningful stimuli also produce a poorer performance.

The holistic representation in the complex, meaningful stimuli, therefore, appears to rely on both configural information, and the goodness of parts. If this were so, then it would imply that holistic properties do not always work in the same direction. Some properties, such as those forming good parts, connectedness or closure, may work on the formation of parts, rather than towards the unification of the representation. Evidence is available to support this idea. It has been suggested that closure operating with local components could encourage the representation of the parts, whereas the grouping of components could encourage holistic representation (Han & Humphreys, 1999; Han, Humphreys, & Chan 1999a; Han, Humphreys, & Chan 1999b).

One constraint in the comparison of the whole probe task and the part probe task lies in the way the tasks were carried out. In the part probe task, the participant was given two probes; the representation of the whole stimulus had to be searched for the presence of one of these probes. In the whole probe task, the participant was given two whole stimuli, and the representation of the whole stimulus had to be searched for a match with one of these probes. However, there were two points of difference between the two whole probes, one unique to each probe. Therefore it could be argued that this task would automatically be easier than the part probe task, because either of the points could be used diagnostically. If this is so, then the WPA in all cases is inflated. Although an interpretation of the WPA in an absolute sense would therefore need to be made with caution, this study was more concerned with comparing the relative WPAs across conditions. Each condition was subject to the same difference between the two tasks.

In addition, it is not certain that the two tasks are performed in exactly the same way under any circumstances. They are designed to distinguish between the use of configural and component information, and the use of different information may entail different processes. Two different sorts of tasks were used in this research; one involved whole stimulus comparisons, and the other used a whole-part comparison. Ankrum and Palmer (1991) provide several accounts for better performance in the whole matching tasks. First, based on the similarity hypothesis (Palmer, 1978), performance in a whole matching task should be easier (and therefore better under the same conditions) than a whole-part comparison. This is because there are more points of similarity in the two whole stimuli than in the whole-part stimuli. Second, the identity hypothesis (Farell, 1985) posits that in a whole probe comparison a match is determined by a conjunctive criterion; all parts match and are identical. In contrast, in whole-part comparisons a match is determined by a disjunctive criterion; all parts cannot match because some are simply not there in the part. The former match is easier than the second match (Posner & Mitchell, 1967).

These differences could lead to completely different processes being used. The whole probe task may have been performed as a matching task, whereas the part probe task was performed as a part search task. Therefore, it would be worth considering the two tasks separately, to see if the findings of the two tasks converge on the same conclusions that were produced when the tasks were considered together.

The whole probe task could be described as a matching task. The meaningless stimuli and the meaningful stimuli with the abstract parts show an effect of complexity, where there are more errors in the complex condition. This suggests a part by part matching strategy. The meaningful stimuli with the concrete parts do not show this effect; in fact, there are fewer errors in the complex condition. This suggests that either the match is made through a holistic, template like process, or else that the figural goodness aided the part search in the complex condition.

The part probe task can be described as a part search task. When the parts that are being searched for are concrete, there is an effect of complexity, with the complex stimuli producing a worse performance than the simple stimuli. This is consistent with a serial search of the parts in the representation. This is found in both meaningful and meaningless stimuli. When the parts are abstract, there is no effect of complexity. This could be due to a parallel search of the parts, or to the formation of new parts specifically for the task. The high error rates would tend to support the latter rather than the former (see Chapter 4).

Another point that needs considering is that the concrete part, rather than enhancing figural goodness, enhances distinctiveness (see Fournier, Erikson & Bowd, 1998). Based on the work of Fournier, Erikson and Bowd, we could assume that the search of the stimulus representations begins with the most distinctive parts, and that the concrete part fulfills this criterion. In that case, the concrete probe will capture attention early in the task; the earlier the part is detected, the more error free the process is likely to be. The whole probe task may also aid discriminability by providing a location for the parts. However, this does not explain why the concrete part is less effective in the meaningless stimuli than the meaningful stimuli. It also does not explain the effect of complexity. If the concrete part captures attention, it should do so early in the processing of both simple and complex stimuli.

In summary, configural information, although useful for all the stimulus types, appears particularly important in the complex, meaningful objects. There is also interaction between the configural information and the goodness of parts in the complex, meaningful stimuli that is not found in the other conditions. One explanation could be that the representations of the meaningful objects have a better structure than the representations of the meaningless objects, allowing facilitation of certain tasks. This is particularly noticeable in the complex stimuli, where processing is likely to be more difficult than in simple stimuli.

CHAPTER 6

OBJECT REPRESENTATIONS AND THE EFFECT OF MEANING

ABSTRACT

The final chapter begins by summarising the results from the previous three chapters. The findings, when combined and reviewed, suggested that the meaningful stimuli are associated with more holistic representations than the meaningless stimuli. Furthermore, these holistic representations are associated with both configural properties and the figural goodness of the parts. The configural properties appear stronger in the complex, meaningful stimuli than in the simple, meaningful stimuli. These configural properties also appeared robust over different tasks, even if the format produced a detriment in performance in the task. These conclusions are then discussed, and set into the context of other research and theoretical approaches. The findings presented here allow some speculation about how and why meaning might encourage more configural processing.

CHAPTER 6

INTRODUCTION

The aim of this thesis was to access the underlying representations of objects, and examine the effect of meaning on the format of those representations. The aim of this chapter is to draw together the results of the three empirical studies performed within the thesis. The three studies were derived from different backgrounds, and have, therefore, the potential for different interpretations. The difference in tasks leads to a degree of caution in making a comparison; however, if the outcome from one task can converge with the outcome from a completely different task, then it adds weight to the conclusions that can be drawn from the results. The converging outcomes are more likely to apply to the stimuli, which are constant throughout the tasks, than to some factor of the tasks themselves, which are not constant. In addition, however, using different tasks may also tease apart the different kinds of holistic representations that may be found. The focus of the thesis is on the difference in representation between the meaningful stimuli and the meaningless stimuli. The common theme throughout the studies was the use of complexity to provide an insight into the format of the representations. The results of the studies are summarised below.

Summary of Results

Rotation Task

The first study looked at the effect of the complexity of the stimuli on the performance in a rotation task. The task used the whole stimulus for the initial presentation, which was left on screen for the participant to study for as long as required. The second presentation also consisted of a whole stimulus, and was presented for 750 ms. The orientation of the second presentation was varied across four angles: 0°, 50°, 100°, and 150°. The task was to say whether the two stimuli were the same or different. If they were the same, they were identical. If they were different, then one area of the stimulus had been transposed to an adjacent area. A deadline procedure was used. In this, the participant had to respond as soon as they were cued by the words "same or different" on the monitor, after the second presentation.

The logic used in this task exploited the different sources of error available in segment-based and holistic representations. In both types of representation, there would be an effect of angle, with an increase in errors as the angle increased. In both types of representation there would be an effect of complexity, where more complex objects would result in more errors than simple objects. However, in the segment-based representations there was an additional source of error, where the parts had to be re-aligned into their correct positions. This would lead to an interaction between the number of parts and the angle of rotation. This source of error would be present for both parallel and serial processing of segment-based representations. If a part-based representation were used, then there would be an interaction between the angle of orientation and the complexity. If a holistic representation were used, there would be no interaction, because the positions would be maintained through the configural information.

An interaction was found for the meaningful stimuli, but not for the meaningless stimuli, suggesting a segment-based representation for the meaningful objects and a holistic representation for the meaningless objects. However, the difference between the two types of stimuli appeared strongest at 0°. After this angle, there was little difference between the meaningful and meaningless objects in terms of the slope of the error rates. The results with the response times supported the results with the error rates, although the deadline procedure may have affected response times. The response times were flat from 50° onwards, although they increased from 0° to 50°.

Rotation had been used as a diagnostic tool in this task. However, rotation may actually have added a complication. The flat results from 50° onwards in both the error rates and the response times could also indicate that rotation was not taking place. This was a source for some concern.

The prediction was based on the occurrence of mental rotation of the representation, and if rotation was not present, then the predictions could not apply. Furthermore, although we predicted increases in error for greater angles of rotation, it is not clear from the literature that error rates are as diagnostic for rotation as response times. If rotation is not occurring, two explanations are possible. First, the task could be performed by examining only a small part of the stimulus, and it is not necessary to rotate the representation to access this information. Second, the deadline procedure could have been too short to allow the representation to be rotated, so the participants had to give up and guess.

Another possible explanation was that rotation itself took place in the same way for both meaningful and meaningless stimuli, and that the differences found were due to differences between performance in a comparison task (with no rotation involved, at 0°) and performance in a rotation task. Both representations may be holistic, but the meaningless representation consists of an undifferentiated representation, and the meaningful representation consists of parts and strong configural information. In the latter case, both components of the representation (parts and configuration) are available as the task demands require them.

The difference could lie mainly with the complex, meaningful objects. This is indirectly supported by the findings with the meaningless, familiar objects. When the meaningless objects were made familiar, the error rates dropped significantly, but the complex results were still higher than the simple results. This suggests that meaning had some particular facilitatory effect when associated with complexity that familiarity did not.

Part-Search Task

The second study looked at the effect of the complexity of the stimuli on the performance in a part search task, where the part was a small area of the whole stimulus. The task used the whole stimulus for the initial presentation. Again, the participants were allowed to study the stimulus for as long as they needed. Two parts were then presented on the screen. They remained on the screen until the participant responded. The task was to choose which of the two parts had been present in the initial stimulus. The task was exploiting the difference in independence among the parts in segment-based and holistic representations.

The logic behind the task suggested there would be a greater effect of complexity in a holistic representation than in a segment-based representation. In a segment-based representation, it would take longer to search a complex that a simple representation. This would also be true for a holistic representation, but in addition the part would have to be isolated from the other interdependent parts first. Not only would there be an effect of the amount of information present, but the level of interdependence would also affect performance. If the representation was completely undifferentiated then isolating a single part would be practically impossible, and performance would be at chance. The simple stimuli would also show a poor performance.

There was an effect of complexity for both meaningful and meaningless objects, but the effect was greater with the meaningful objects. This, in combination with the effects of complexity, suggests that the meaningful stimuli were represented holistically, and the meaningless objects produced segment-based representations. There was little difference in performance at the simple level, but, as in the last study, a difference occurred at the complex level. The error rates were below chance. This suggests that the holistic representation was due to configural information rather than either the parallel processing of several parts, or the use of undifferentiated representations.

One potential problem with this study was that the parts that were searched for might not have been the parts that were naturally parsed. If this is so, then a parsed representation could have been at a greater disadvantage than a non-parsed representation. In the parsed representation, the parts presented in the task may not have corresponded to the parts achieved by parsing. If this were so, the "natural" parts would have to be broken down into smaller components, then recombined in different groupings to form the task parts.

Whole and Part Probe Task

The aim of the third study was to examine the relative role of configural information. Two tasks were performed, one using a whole probe and one using a part probe. The reasoning behind this study was that if configural information were an important factor in the representation then there would be an advantage for the whole probe task. If the part-based information were an important factor in the representation, then there would be no advantage for the whole probe task, and a possible advantage for the part probe task. Rather than just looking for the presence or absence of a whole probe or part probe advantage, however, the aim of the experiment was to assess whether the findings were equivalent across the stimulus types and complexity. In addition, two types of part were used as the probe, differing in the level of figural goodness.

Both tasks used a whole stimulus as the initial presentation, and this remained on screen for 275 ms. The participant was then presented with two probes. One task used whole probes, in which two very similar, whole stimuli were presented. The other task used part probes, in which two parts were presented. The whole probes were formed by modifying a small area of the stimulus. Isolating this small area then formed the part probes. The probes remained on the screen until the participant responded. The response required was to indicate which of the two probes was either the same as the earlier stimulus, or was present in the earlier stimulus.

Both types of stimulus, meaningful and meaningless, and both levels of complexity, complex and simple, showed a better performance in the whole probe task. The improvement in performance with the whole probes may be partially accounted for by the fact that the whole probes differed in two areas, which would provide the participant with two opportunities to detect a difference when the whole probes were presented. However, there was still a differential effect of the tasks across type of stimulus and complexity that had to be accounted for. The increase in performance with the whole probes was the same for both complexities of the meaningless stimuli, but differed across complexities for the meaningful stimuli. This difference was related to the figural goodness of the part used as the diagnostic feature in the probe. The performance with abstract probes was compared with that of concrete probes. The biggest difference in performance was with the complex, meaningful stimuli, where performance with the concrete parts produced a higher whole probe advantage than with the abstract parts.

Review of the Findings

The studies presented in this thesis were independent pieces of research. However, they had a common theme in that they were looking at performance at different levels of complexity to access the underlying representation. The use of complexity to distinguish between representations was grounded in the different properties shown by the representational formats. Before we look at the findings, we need to review what these representational formats were.

There were three main understandings of the term "holistic" in relation to representations. First, a holistic representation could exist as a single unit, an undifferentiated whole. Here, the structure of the parts would be none existent. Different accounts have the undifferentiated representation arising through one of two routes. Tanaka & Farah (1993) suggested that this sort of representation was due to the very strong relationship among the parts. The configural information was as strong between the parts as it was within a part, and consequently the parts had lost their individual structure and independence. Palmer (1977) suggested that when an object was first represented, the early representation would be unstructured, and undifferentiated. This representation would then be parsed into parts, and the parts re-built to form a whole. The "whole" that resulted would also be a holistic representation, but a second type.

In this second type of holistic representation, there can be parts, probably well-defined and structured parts, but the predominant information in the representation is concerned with the relationship among those parts. The segment-based representation also consisted of the representation being parsed into good parts, but here the predominant information was found with the parts themselves, rather than their relations. The parts in the segment-based representation retained their independence.

Finally, holistic could refer to a segment-based representation where the parts were all being processed together. In a segment-based representation, the parts were relatively independent, and so each part could be isolated and processed apart from the others with relative ease. However, it has to be remembered that although the parts could be processed separately, they do not have to be processed separately. They could be processed in parallel. This would not be a holistic representation as such, but could be considered an example of holistic processing. It can be difficult to differentiate between holistic representations and holistic processing. Despite the lack of consensus of opinion on holistic representations, there was general agreement that: first, configural information was more important than part information in a holistic representation; and, second, the parts were not as independent in a holistic representation as in a segment-based representation.

The tasks used here can be divided into two broad groups. The rotation task and the whole probe task involve the comparison of two whole stimuli to detect the presence or absence of a difference between them. The search task and the part probe task involve the representation of the object being searched for the presence or absence of particular parts. The two types of task make different demands, leading to two sets of predictions (based on Farah, Wilson, Drain & Tanaka, 1998).

If the task requires the comparison of two whole representations, then two basic strategies can be used to complete the task. First, each part of the representation can be compared with the corresponding part of the other representation. Maintaining this information in memory is likely to involve a heavy processing workload. Alternatively, the whole representation can be compared to the other whole representation. The use of the second strategy will depend upon the ability to form a holistic representation, and, perhaps, on the nature of the difference between them. If the difference involves replacing one part with a completely different part, it may not be so readily detected as a change in the whole shape. For example, research has been carried out where facial features were replaced with general objects, and the face was still perceived as a face (Donnelly, Humphreys & Sawyer, 1986). Under controlled conditions, the participant might have had difficulty deciding whether the "mouth" was, in fact, a banana or a melon slice, because, although very different, both features are serving the same purpose in the whole. The tasks in this study used differences where one part was moved rather than replaced.

In the whole comparison tasks, there will be less effect of complexity in a holistic representation than in a segment-based representation. The holistic representation has one unit to compare, whereas the segment-based representation has several. There may still be an effect of complexity in the holistic representation, because the complex stimulus contains more information than the simple stimulus. In the segment-based stimulus there are more units to examine in the complex representation than the simple representation. Note, however, that this assumes a serial search of the segment-based representation. A parallel search may give less effect on response times than a serial search, duplicating the effects of a holistic representation, although the error rates may still be affected.

In the tasks where a part has to be searched out, then the part has to be separable from the rest of the stimulus. The alternative is that the part can make a difference to the whole object, and then a whole comparison can be used. However, if only the single part is shown, rather than the whole object, then to use the latter strategy, the relationship of the part to the rest of the object has to be clearly represented in memory.

Therefore, if a task requires a single part to be isolated before processing, then there will be less effect of complexity in a part-based representation than in a holistic representation. There will be an effect of complexity in the segment-based representation, because there are more parts, but these parts are already relatively separate. The effect of an increased number of parts will be additive. In a holistic representation, the individual part has to be isolated from the rest of the representation. Assuming the parts across complexities to be similar in size, the part in a complex, holistic representation will suffer more interference than an equivalent part in a simple representation. The effect of an increased number of parts will be multiplicative. Therefore, in a part-based representation, although there is likely to be an effect of complexity, this will be smaller than in a holistic representation of comparable complexity.

To summarise, in the whole stimulus matching tasks, successful performance can be achieved either by comparing the two representations, part by part, or by matching the two holistically. If there are no properties in the representation that will allow a holistic match, then the match will have to be part by part. A part by part match will entail an effect of complexity, with poorer performance in the complex stimuli than in the simple stimuli. When the results from all the matching tasks are combined and compared (see figure 6.01), it can be seen that this applies to the meaningless stimuli, and the meaningful stimuli when abstract parts were used. Therefore, the meaningless stimuli used a part by part match, and the meaningful stimuli used this approach when the difference consisted of a part with poor figural goodness. The meaningful stimuli with the "good" parts, however, showed a holistic match. An alternative explanation is that a part by part match also took place in these stimuli, but the figural goodness helped detection of the part. This would not seem correct, however, because it does not account for the better performance in complex stimuli compared with the simple stimuli.

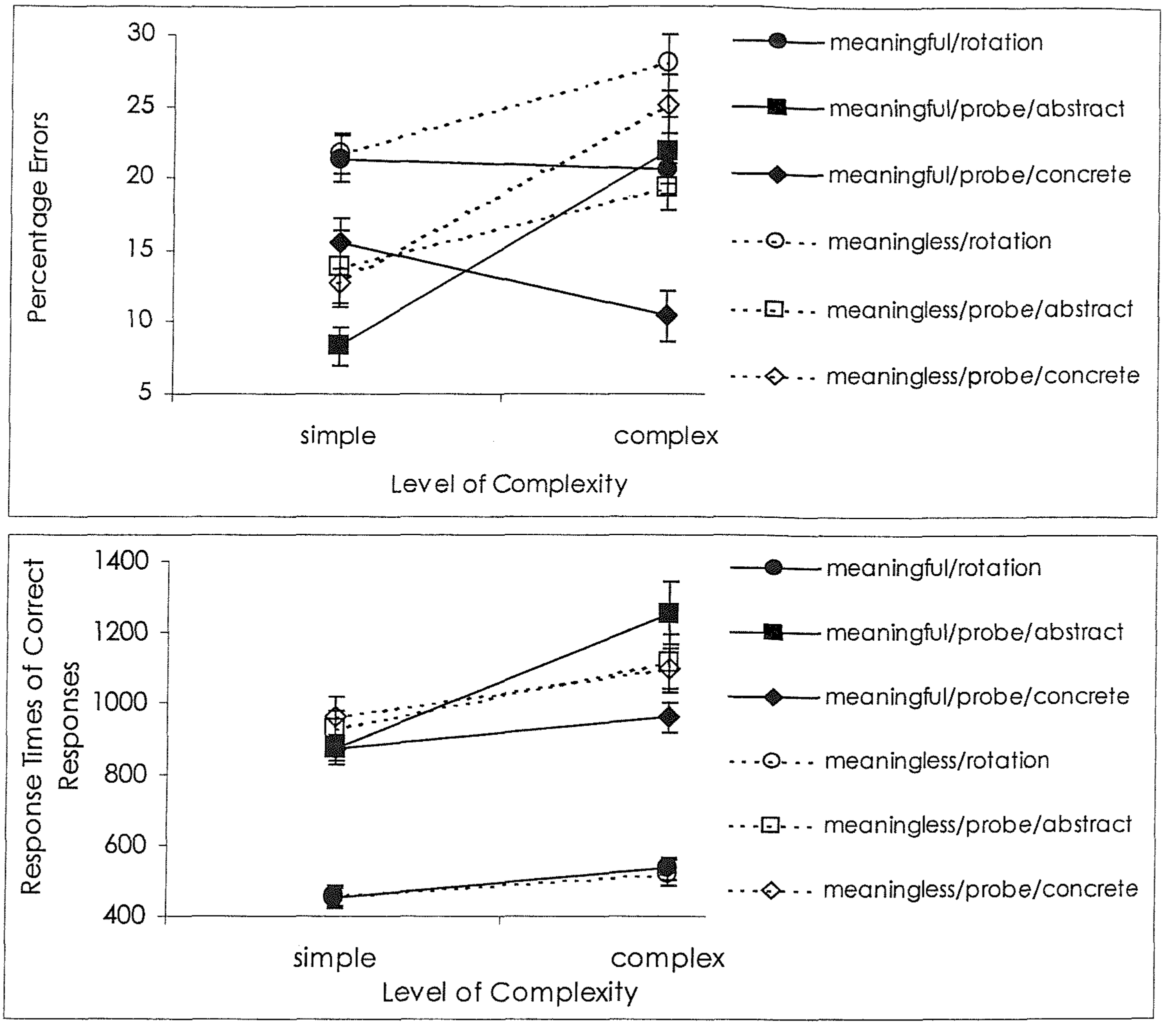

Figure 6.01: Performance across the Whole Stimulus Tasks.

Caution has to be used in comparing the results across tasks. The results from the rotation task show low response times. This is a function of the deadline procedure. The stimulus was presented for 750ms before the participant was prompted for a response; this time needs to be taken into consideration as part of the response time. The legend gives the types of stimulus, the type of task, and the type of part (if appropriate).

If the meaningful stimuli were represented holistically when the parts have figural goodness, and the other stimuli are represented segmentally, then there should also be an effect across representational formats at the simple level. There should be one part in the holistic stimuli, and more parts in the segmented stimuli, and this should give an effect of complexity. Although there was a difference between the error rates in the simple condition across the types of stimulus, this difference did not show an advantage for the representations deemed holistic above. The lack of such an effect suggests two possibilities. At the simple level, the meaningless stimuli (and the meaningful stimuli with poor parts) also had only one part. It is not clear, however, why this should be.

An alternative explanation is that perhaps the simple stimuli had more than one part in all cases. This implies that the complex, meaningful stimuli with good parts are in some sense more holistic than the simple equivalents. This is supported by the results, from both the whole probe task and the rotation task.

One important difference across the tasks was the presentation time of the original stimulus. The higher error rates could be caused by the short presentation time that was given in the whole probe task. If only a brief presentation time was given, then the participant may have had difficulty encoding the whole stimulus, leading to more errors. This effect would be more striking in the more complex stimuli. However, the error rates do not appear a function of the initial presentation time. In the rotation task, participants were allowed as long as they wanted to study the stimulus. However, in neither the complex nor the simple stimuli are there fewer errors in the rotation task than in the whole probe task.

A further possibility is that the short processing time in the rotation task (because of the deadline procedure) could increase the number of errors. However, although there were more errors in the rotation task than in the whole probe task, it would be expected that insufficient processing time would affect the complex stimuli more than the simple stimuli. This was not the case.

In the part search tasks, successful performance will involve matching the target part to an appropriate part of the stimulus. This will probably take place through a part by part search of the representation. If the parts are not already present in the representation, then they will have to be formed. When the target part is figurally good, then the representation may already contain that part. If the target part is figurally poor, then it will probably have to be formed during the task. If the representation is not parsed, then both figurally good and figurally poor parts will have to be formed.

In addition to the formation of the parts, the configural information will play a role in the performance of the task. If the configural information is strong, then the parts will be less independent, and it will be more difficult to isolate one part.

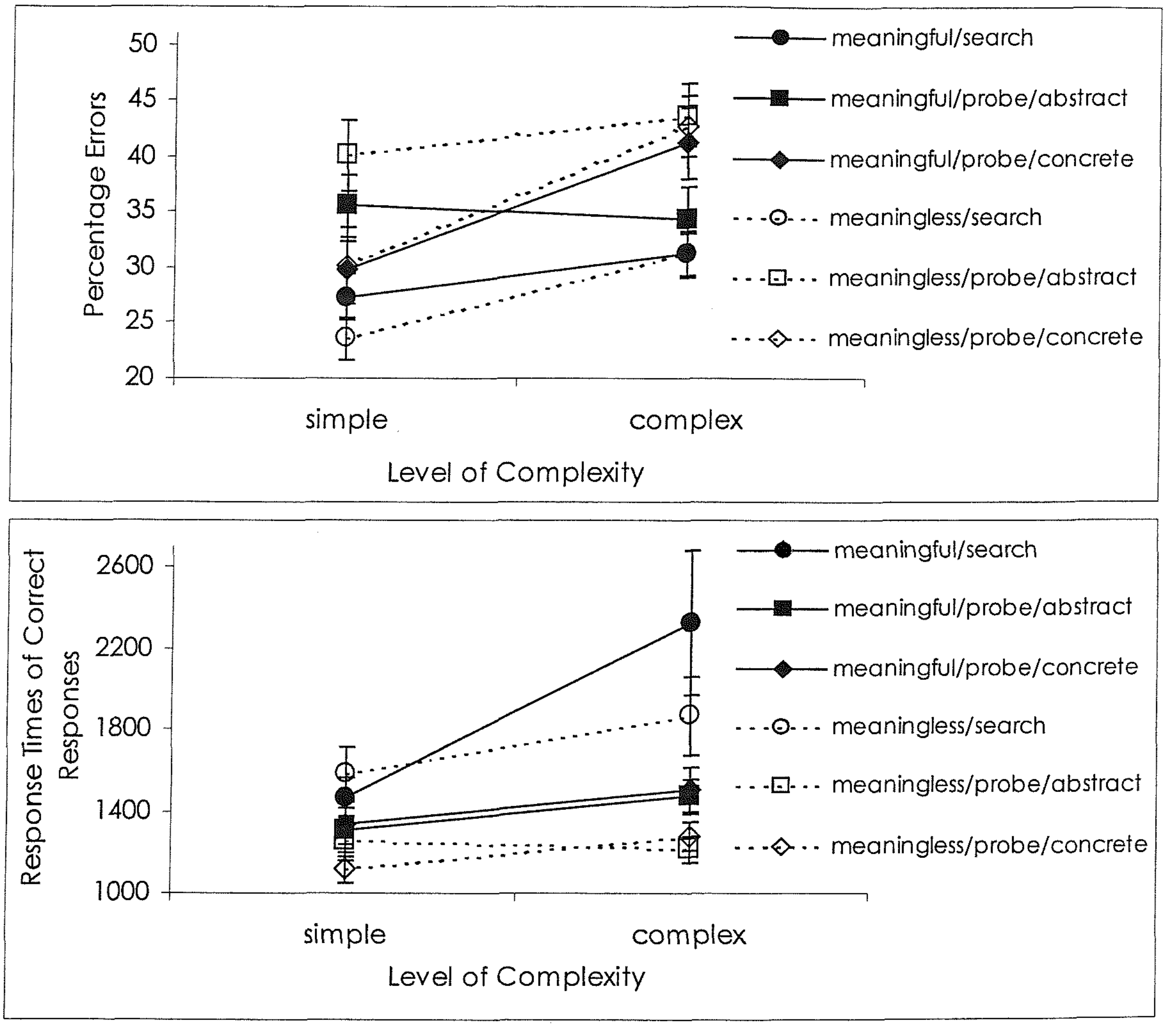

Figure 6.02: Performance across the Part Stimulus Tasks.

Caution has to be used in comparing the results. As with the summary of the whole matching tasks, the legend gives the types of stimulus, the type of task, and the type of part (if appropriate).

The results from the part search tasks were also combined and compared (see figure 6.02). When the parts that were being searched for are figurally good (the concrete probe task), then the poorer performance with the complex stimuli, both in response times and error rates, is consistent with a serial search of the parts in the representation. This was found in both meaningful and meaningless stimuli. When the parts were figurally poor, and there was only a brief presentation time (the abstract probe task), then there was no effect of complexity in the response times or error rates, and the error rates were high. This is consistent with the formation of new parts specifically for the task. When the parts are figurally poor, and there is a long presentation time (the part search task), then there is differential performance across meaningful and meaningless stimuli. In the meaningless stimuli, the results are consistent with a part by part search. In the meaningful stimuli, there is a higher difference in the performance across complexities. This could also indicate a serial search, but perhaps with configural information interfering.

To draw all these results together, the findings from Chapter Three indicated that there was a difference between meaningful and meaningless stimuli in terms of their representations. This difference was characterised by an improved performance in the complex stimuli, to a level better than that of the simple stimuli. The findings in Chapter Four showed that the meaningful stimuli were more holistic than the meaningless stimuli. Furthermore, the holistic representation that was produced was due to an increase in configural properties rather than either a lack of parsing, or parallel processing.

The findings presented in Chapter Five indicate that configural information appeared to aid performance in the task for all the stimulus types. However, the use configural information was more noticeable in the complex, meaningful objects. The configural information was associated with the goodness of parts in the complex, meaningful stimuli; this association was not found in the other conditions, even the simple, meaningful stimuli. One explanation put forward was that the meaning allowed a better structure of the representation, thus making performance in some tasks easier. This was most evident in the complex stimuli.

One area that needs to be considered is the "special" qualities of the meaningful, complex group of stimuli. The use of two complexities began as an analytical tool. However, the possibility arose that complexity may also affect format. In all the tasks, this group stood out from the other three. In the rotation study, the difference between the meaningful and meaningless pattern of results was due to the relatively low number of errors in the complex, meaningful stimuli. In the part search study, the difference between the meaningful and meaningless findings was again concentrated on the complex level, with little difference in performance at the simple level. In this case, the complex, meaningful stimuli showed a relatively poorer performance. Finally, in the probe task, the difference between the meaningful and meaningless stimuli was in the complex level, with the meaningful, complex stimuli showing a variation in the whole probe advantage, with a greater use of configural information combined with an effect of figural goodness.

These findings are consistent with the idea that the complex, meaningful stimuli are more holistic than the other types of stimuli in these studies. If this is so, then it suggests that the representations in the meaningful group of stimuli are more variable, in the sense that there appears to be some difference in representation across complexity that is not found in the meaningless stimuli. It also introduces the possibility that the complex, meaningful stimuli may be more inflexible in their representational format than the other groups of stimuli. The complex, meaningful stimuli are "fixed" in a holistic format; as such, this format has to be introduced into the process, even if using an alternative representation would be more suitable for the task.

A possible explanation for this is that the complex stimuli may have more configural properties. This may be because the simple stimuli have a stronger part-based component to the representation. It could also be a consequence of the greater number of parts in the complex stimulus leading automatically to a greater amount of relational information among the parts. When the complexity increases, the representations of the meaningful stimuli can become more holistic, thus reducing processing demands, and, consequently, the error rates. The representations of the meaningless stimuli cannot become holistic, and so performance becomes more difficult, and consequently poorer.

Another area for consideration is the role the task plays in determining the format. One explanation put forward to account for the findings in the rotation task was that the task itself might have contributed to the format. This does not seem the case throughout the other studies. If the task did constrain the representation to a large extent, then both meaningful and meaningless should produce the same type of representation within the same task. This was not found. Some degree of flexibility to produce a suitable format for completing the task may be present, but in the case of the complex, meaningful stimuli this seems to provide a lesser constraint than the properties of the stimulus itself in determining the format.

This leaves a question about the simple, meaningful stimuli. Are they less holistic than the complex stimuli, or are they more variable? The simple, meaningful and the simple, meaningless objects have generally produced performances similar to each other, and better than their complex counterparts. In some of the tasks, the complex, meaningful stimuli have equalled, or bettered, the performance with the simple stimuli. The complex, meaningless stimuli have generally produced the worst performance. This would seem to indicate a continuum of representation from very segment-based to very holistic. The simple stimuli from the two stimulus types lie relatively close together on the continuum. The complex, meaningful stimuli are more holistic, and the complex, meaningless stimuli are less holistic. This would mean that complexity interacts with the meaning of the stimulus to affect the format, rather than operating independently.

In each of the three tasks, the decision about whether a stimulus is holistic or segmented is not absolute. The claim is not that one stimulus is definitely holistic, and one is definitely segment-based; rather, the claim is that one is more holistic, and one is more segment-based. This, in turn, means that the definition of the representation of one type of stimulus is made in relation to the other stimulus. If a stimulus has access to more than one representation, then the best representation for the task can be selected from the options. If another type of stimulus has fewer options, then the choice will be more limited. So, the limited stimulus will keep to very much the same format whatever the task, and the more variable stimulus will be represented by whichever option is best. The variable stimulus will therefore produce fairly consistent results across tasks; that is, the optimum available given the other task constraints, whereas the more fixed representation will sometimes give better or worse performances depending upon how well the representational format interacts with the tasks demands. However, as the results are in relative terms, it is difficult to unpack how flexible the representations of the different stimuli are.

We can now give a tentative answer to the question: What role does meaning play in determining the format of the representation of the object? Meaning appears to affect the representational format by encouraging a more holistic representation. However, this is not consistent over all levels of complexity. Specifically, more complex, meaningful objects are represented more holistically than simpler, meaningful objects. Furthermore, the holistic representation appears to contain both configural and part information, but be more dependent upon the configural information. The configural information also appears related to the goodness of the parts.

CHAPTER 6

DISCUSSION

Methodological Considerations

Before we consider the results of the present research in the light of current theories, we need to consider the stimuli themselves. From a theoretical perspective, we needed two sets of stimuli, one meaningful and one meaningless. From a methodological perspective we needed simple and complex examples of each type of stimulus.

The meaningfulness was assessed in terms of the ability to name the stimuli, and the confidence in the choice of name; the familiarity of the stimulus; and how representative the stimulus was of the object it depicted. It could be argued that these criteria do not fully define the term "meaningful". If this is so, then the results of this research apply to whatever term these criteria do define. The important point is that the criteria do not define familiarity alone. In the first study, training was used to increase familiarity, without confidence and representativeness being increased. This did not result in the same performance being produced as when all three factors were high.

Complexity was measured in terms of a compactness measurement that had been used in previous studies. The calculation (Attneave, 1957) involved taking the square root of the area, and dividing it by the number of lines in the perimeter. The calculation used here also incorporated the internal lines. Low compactness values signified high complexity.

The area of the object could affect the complexity calculation; a smaller area will give a higher complexity. This was not a problem in this series of studies because the areas were kept comparable across simple and complex stimuli. The use of the internal lines may have exaggerated the complexity of the stimuli if the cognitive system relied more heavily on the peripheral lines. Finally, the true complexity of the stimulus could lie in the amount of information conveyed, rather than the physical attributes (see Garner, 1970). It could be that some of the physical parts of the stimulus were redundant in terms of conveying information to the cognitive processes. All or any of these factors could contribute to an exaggeration of the complexity of the objects.

The important issue here is that the complex objects should be significantly more complex than the simple objects. If both simple and complex are actually simpler than the compactness values suggest, then the difference between them is maintained. A problem could occur if the complex objects were exaggerated more than the simple objects, resulting in a falsely high difference. As overall performance in every task was significantly better with the simple stimuli than with the complex stimuli, then this would seem evidence that both the information and physical properties of the complex objects were more complex than those of the simple objects.

The research in this thesis has been set within the framework of object recognition, with the purpose of considering the representations that might be formed during recognition processes. Recognition is about adding meaning to objects in the environment, and this thesis was interested in what happens to the representations of those objects when the objects become meaningful. The examination of the representations is achieved through various matching tasks rather than recognition *per se*. This might not generalise to the representations used in object recognition. Different processes may be used to match two representations from those used to recognise the object (e.g. see Tanaka & Sengco, 1997). However, the tasks used in this research do incorporate a memory element and a perceptual element, in that the matching task involves the comparison of a representation from memory with a representation coming from the perceptual system, which is what is posited as occurring during object recognition.

CHAPTER 6

Theoretical Considerations

One question that arose from the literature was how "holistic" should be defined. It has been used to describe undifferentiated material, differentiated material with a high reliance on configural information, and segmented material that is processed in parallel. However, before different terms are used to describe these very different situations, another question arises, and that is whether the different forms of holistic as described above can be distinguished empirically. This has been attempted in this research. The conclusions drawn here are that the representations of meaningful objects are parsed into parts, but then rely on configural information to a greater degree, particularly in complex objects.

To what extent does strong configural information rely on well-defined parts of the object? It would be worth reconsidering the two types of configural information that have been proposed (e.g. Carey & Diamond, 1986; Cooper & Wojan, 2000). If the parts are not well defined, then the categorical (or first order) information is still available - above, below and so forth. To take this to the next stage of measurement, be it co-ordinate or second-order, well-defined parts are needed. One cannot specify the distance from part A to part B unless one is quite clear about where part A ends and where part B begins. This would imply that to have precise configural information, one also needs clearly structured parts. Therefore, although experts have been found to rely on configural information to perform recognition tasks, perhaps the tasks used are not tapping in to the part-defining skills that may be present. If holistic processing is the processing of a single, undifferentiated unit, then configural processing, as discussed above, cannot fit into this definition.

With the exception of the complex, meaningful stimuli, the use of abstract parts did not impair performance relative to the use of concrete parts to any great extent. This has two possible implications. The first is that unless other factors act to prevent it, then the parts can be formed in response to the task. If the parts are formed in response to the task, then the relations among them must be formed at the same time. If configural information is an important, stable factor in a representation, then perhaps the parts generating the configural information also have to be stable. Alternatively, the "natural" parts may be smaller than the abstract parts, and so the abstract parts can be formed through a combination of natural parts, rather than them having to be broken. Saiki & Hummell (1998a) have proposed that the largest parts necessary will be formed to reduce the resource use in computing many relations. Perhaps such relational computations are also formed in response to the task, and otherwise not calculated. In experts, this information may be readily available because it has been calculated on many previous occasions. In novices, the calculations would have to be made from fresh.

We have suggested that meaning encouraged a holistic representation whereby figurally good parts and relations among the parts could be formed. We have also suggested that greater configural properties can arise as a result of the greater number of parts in a more complex stimulus. This would explain how meaningful stimuli could demonstrate a better performance, especially in the complex condition, than the meaningless stimuli.

However, it does not address what is happening in the meaningless stimuli. How can meaning have such an effect on the stimuli, and why does it not happen with the meaningless stimuli? One possibility is that the complex, meaningless stimuli are represented as an undifferentiated unit, and not easily parsed into natural parts. Although parts can be formed, these are not functional in a hierarchical sense, as meaningful parts are. If the parts are not stable, then the configurations will not be stable. If configural information has to be recalculated for every encounter, it might make it difficult to use. When it is used, it might deplete cognitive resources, leaving little processing power for other task demands, and, consequently, giving a poor performance. Computing the relations among many parts in a complex stimulus will be more demanding than computing relations among a few parts. This would explain why, even if parts and their spatial properties were calculated, the configural properties did not confer the same advantages to the complex, meaningless stimuli as they did in the meaningful condition.

The findings here can be seen as consistent with previous research. If we interpret these results in the framework of Palmer's (1977) series of representations, we can explain them by the meaningful stimuli having being parsed, but the parsing of the meaningless stimuli being less robust. Consequently, the meaningless stimuli are parsed in response to the task. The type of part will then have no effect on performance.

The meaningful stimuli have already been parsed. Here, if the part used in the task is "good", and so matches the parts already formed, then there will be a facilitation of performance. If the parts used are "poor", then performance will be comparable with that of the meaningless stimuli. In the part probe task, there is no difference between meaningful and meaningless stimuli. There is a difference in the whole probe task. If the meaningful stimuli have not only been parsed, but have also been combined back into a structured whole, the structure of the stimulus can also be used in the task. However, it cannot be used in the part probe task, because there is no structure present in the probe.

This is consistent with a movement from an undifferentiated representation, through a parsed representation, through to a re-structured representation, as put forward by Palmer (1977). This is supported by the findings of Gauthier, Williams, Tarr, and Tanaka (1998) who found a similar progression through representational formats when they trained participants to become "greeble" experts through familiarisation tasks. The undifferentiated stimulus may not be one single unit, but it has to function as a single unit in some contexts because it has no functional parts. It may, however, consist of many small features. This is consistent with Kimchi's idea (1998) that some of the levels of a hierarchy may be described more accurately as a texture, rather than as a single unit. The representation reaches the level of a structured whole through the use of well-defined, "good", functional parts, joined together through configural information that specifies their spatial arrangement. The parts and their configuration are closely related; one cannot exist without the other. Once this level has been reached, the representation can exist on a continuum, with the configural and part information playing smaller or greater roles, consistent with the ideas of Tanaka & Farah (1993). The demands of the task may specify an "ideal" representation. If this is not possible due to constraints from the stimulus, then performance in the task will be affected.

We can also speculate about what effect familiarity and meaning could have on this hierarchy of representation. One factor associated with meaning is that it could allow the representation to be parsed into functional parts. This would appear important in allowing a structured representation to exist, with stable parts and configurations. Ankrum and Palmer (1991) differentiated between scene-based and object-based hierarchies, although two levels in a larger hierarchy may be an equally apt concept. If the part could be interpreted by the cognitive system as an object in itself within a scene (perhaps the parts would be meaningful within themselves, such as the cat's tail), then any object advantage would move to the part. If the stimulus could not be treated in this way, then the parts would not be given such an advantage.

Familiarity could work in two places in the hierarchical model suggested earlier. First, it could work on the undifferentiated representation to allow parts to be formed. Familiarity might not be as efficient as meaning in parsing the parts, because the parts are not functional in the same way, but over time, stable parts could be produced. In the real world, familiarity would often be accompanied by meaning at this stage. Second, in the context of expertise, it could work through a practice effect to stabilise the computations of configural properties. As the configural properties become stronger, they will become more important within the representation. It is not clear whether or not meaning can also serve this purpose.

CHAPTER 6

Further Studies

The experiments within this thesis have raised far more questions than they have answered, and it is apparent that much more work is needed in this area before firm conclusions can be drawn. This section examines possible future experiments, and how these could help address some of the questions in this area.

There are three main approaches to further research that are considered here. The first approach considers experiments that could be added in to the present experiments to address some of the issues raised in the discussions. The second approach considers further experiments that could extend the studies reported here. Finally, the third approach looks at developing the research further.

Further experiments within these studies

Study 1: Rotation

One problem with this study is that it was not clear whether rotation was occurring or not. Although there was a clear difference in performance with meaningful and meaningless stimuli, the main effect of this was found at 0°. One way to address this concern would be to repeat the experiment, but remove the deadline procedure. This would allow an examination of the response times, and if rotation is occurring, a linear function increasing with the angle of rotation should be found.

Study 2: Part Search

The main criticism here was the type of part used in the task. It was predicted that a holistic representation would shower a greater effect of complexity than would a part based representation, because of the level of independence of the parts. However, if the parts formed in the representation interfered with the parts used in the task, then there could be a greater effect of complexity with part-based representations. To resolve this, the task should be repeated, with the abstract parts replaced with more concrete parts. If this reduced the effect of complexity in the meaningful stimuli, then it would imply that the effect of complexity was due to the parts rather than the configuration.

Study 3: Probe Task

The main two points to be addressed here are that the whole probes had two differences that could provide a response, and that the comparison of the parts was between concrete and abstract. One way to resolve this would be to re-design the experiment so that the same area of the whole object is used to produce three whole probes (resulting in only one difference between pairs of probes). Each of these three whole probes would then have one abstract part and one concrete part derived from the change. Concrete probes could then be compared with concrete probes, and abstract probes with abstract probes.

Extensions to the present studies

Converging evidence using different paradigms

This study used three approaches, but all relied upon the differences between complex and simple stimuli to distinguish between holistic and part-based representations. Having drawn conclusions based on the findings in these tasks, it would be valuable to see if the outcome can be extended to tasks not using complexity as a tool.

Several methods have been offered as useful in distinguishing between holistic and segment-based representations (see Farah et al, 1998). Two appear particularly promising. Farah and Tanaka (1993) made use of a mask. They reasoned that if the representation was a whole, then a mask of parts should not interfere to the same extent as a mask of the whole. Different masks may also give insight into the size of parts that might be used by the cognitive system.

The second method was originally used to distinguish between object-based and location-based attention (Lavie & Driver, 1996), but later used to look at the segmentation of objects and the effect of meaning (Ngohayon, Kawahara & Toshima, 1999). Two objects are presented next to each other (Ngohayon et al used Japanese kanji). Two probes were presented, either both black or white circles, or one black and one white circle, and the participant had to say whether they were the same or different. In half the trials, the probes were on the same object, and in the other half they were on different objects. A same-object advantage suggested that the two probes had been incorporated in the same, holistic representation. This method could be adapted, as with the last method, to look at general objects.

Other tasks that have been described in the literature exploit the disruption of configural information to assess how important this is to the overall representation. One procedure would be to invert the object, and time the recognition. When the object is represented holistically, inversion increases the recognition time (see, Diamond & Carey, 1986; and Yin 1969).

Stimuli

One limitation with the methods used in this series of research is that differences across complexities could not always be readily unpacked. In both the meaningful and meaningless stimuli there were implications that there may be a difference. One way to examine this more closely is to use more levels of complexity, and use the same methods to assess the representation at the higher and lower levels.

Different stimuli could also be used. One interesting stimulus to use could be faces. Faces are agreed by most researchers to be holistic; comparing them with the meaningful and meaningless stimuli could give a clearer picture of how the representations stand relative to each other. Three-dimensional stimuli could also be used.

Probes

In the whole and part probe study, there were effects of the type of part in some conditions and not others, and this would be worth exploring in more detail. In particular, it would be valuable to manipulate the parts to see if connectedness was the factor involved. Other factors to be considered include position of the part and the functionality.

The Role of Memory

It is not clear from the results of the present studies whether the differences in performance between meaningful and meaningless stimuli are focused at the perceptual level, during encoding and formation of the representation, or a function of memory processes. This would be worth unpacking, particularly in the search task. This could be addressed in two ways. One way is to constrain the time available for encoding (that is, the presentation of the first stimulus, the whole object). The second way would be to reduce the reliance on memory by changing the order of the procedure, so that the part is presented first, and has to be encoded by the participant. Then the whole stimulus is presented, and left on screen for the participant to make a decision whether the part was in the stimulus or not. This would also examine the encodability of concrete versus abstract parts.

Developing the studies

Meaning

There are several ways in which the present research can be developed. One approach of particular relevance is to take another look at what meaning is from a cognitive perspective.

In the rotation task, the main difference between meaningful and meaningless stimuli occurred at 0°, which was essentially a matching task. The rotation study also incorporated an experiment to assess the role of familiarity, by increasing familiarity through a training task. Training with familiarity did not produce the same pattern of performance as seen with the meaningful stimuli.

The training could be extended. Several types of training could be given to the participants. For example, a name for the meaningless stimuli could be introduced, or semantic information could be provided, or the function could be taught. These results could be compared to the original results to see what, if any, training can modify the performance with the meaningless stimuli to approximate that of the meaningful stimuli. If one of the tasks can achieve this, then it suggests that whatever property was learned is that being used by the cognitive system in the meaningful stimuli. Once this has been established, the property could be used to derive more meaningful and meaningless stimuli (with and without the particular property). It may also be possible to define complexity within the terms of the given property, giving a more functional operationalisation.

CONCLUSIONS

To conclude, the question asked in this research was whether meaning influenced the format of the representation of an object. The format that was under consideration was whether meaning encouraged a holistic or a segment-based representation relative to a comparable meaningless object. In addition, through the use of different approaches, the research was designed to differentiate between the three versions of holistic representation, and define exactly which, if any, were used in the meaningful or meaningless representations.

The findings presented here suggest that meaning acted by encouraging the use of configural information or processing. This effect is stronger in the complex meaningful stimuli than in the simple stimuli, and appears associated with the figural goodness of the parts. Furthermore, the use of configural information in the complex, meaningful stimuli appears robust, even at the detriment of performance in the task.

APPENDICES

APPENDIX 1: COMPLETE STIMULUS SET.

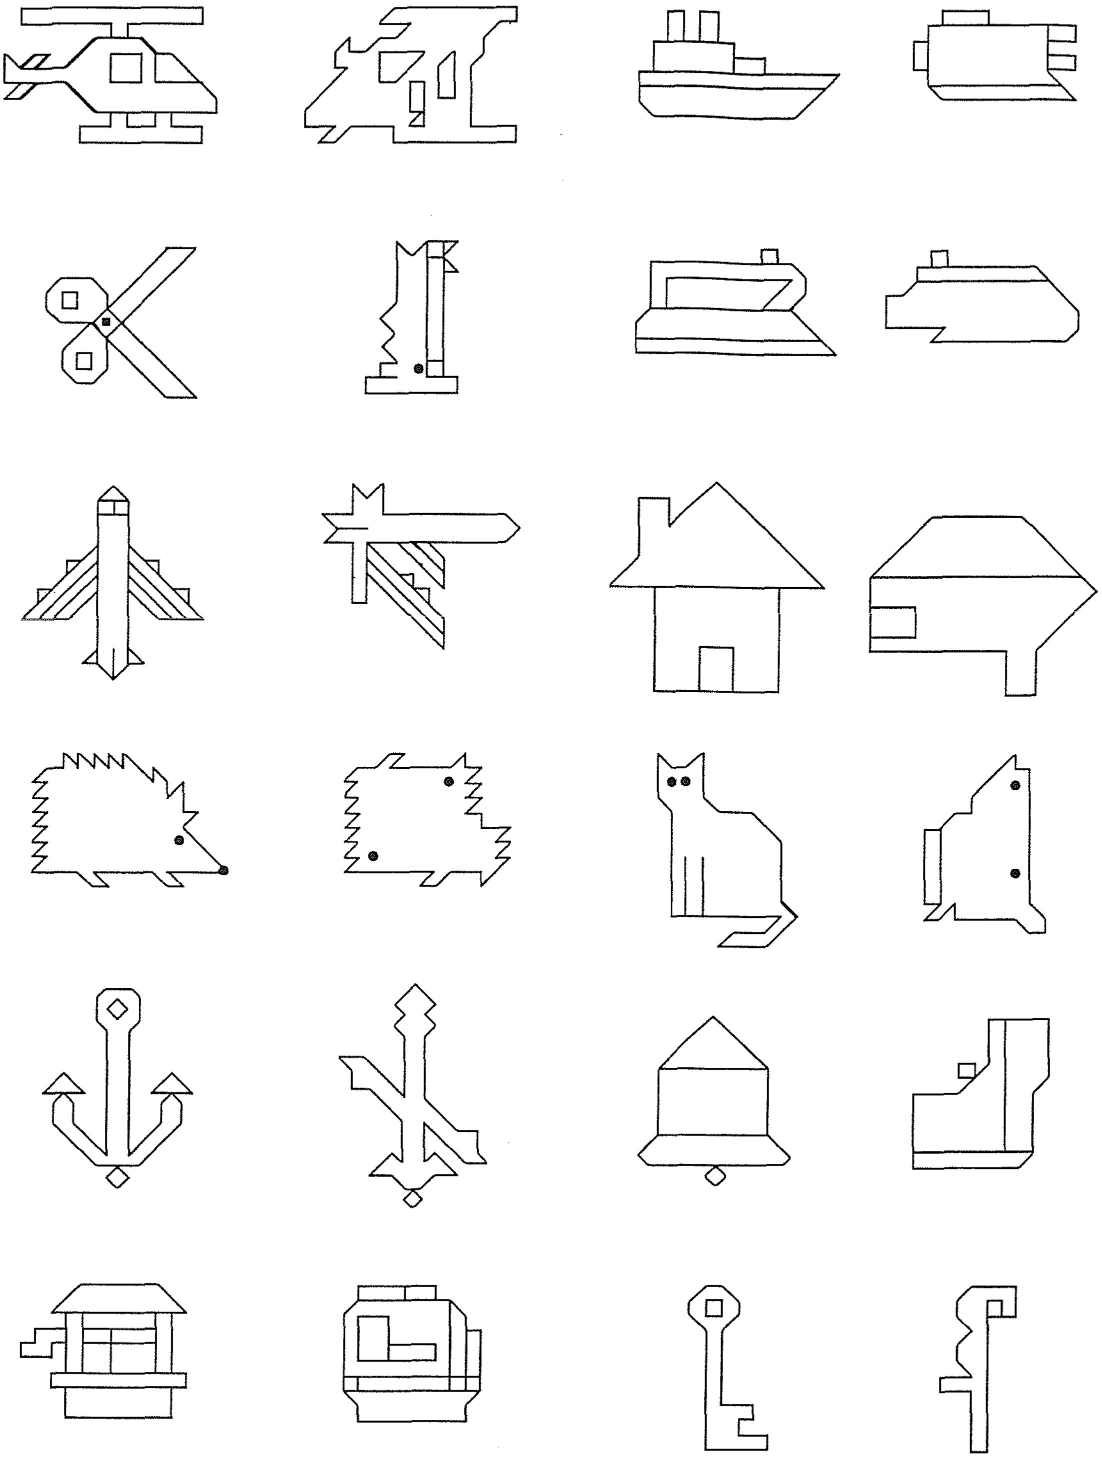

Figure 1: The Complete Stimulus Set
In columns from the left: meaningful complex shapes; meaningless complex shapes; meaningful simple shapes; meaningless simple shapes. The stimuli shown on the computer screen were larger than the ones shown here.

APPENDIX 2: ASSESSMENT OF STIMULI FOR MEANINGFULNESS

One issue of concern was the values for the anchor, which showed lower values than the other stimuli for all three judgements. Questioning of the participants afterwards revealed that the problem here had been that they were not sure whether the anchor was an anchor or a candelabrum, and this had reduced their overall assessments. Both anchors and candelabra are meaningful, and the participants would not have to recognise the object by naming in any of the tasks. The anchor was kept in the set. It was considered that the properties it added to the overall set were useful, and the slightly low values it contributed to meaningfulness were in part due to confusion between two meaningful objects, and were compensated for by other members of the set.

APPENDIX 3: ASSESSMENT OF STIMULI FOR COMPLEXITY.

Table 1: Mean values for compactness values of the stimuli

	complex stimuli	compactness value	simple stimuli	compactness value
meaningful versions	helicopter	0.17	ship	0.42
	scissors	0.20	iron	0.42
	aeroplane	0.18	cat	0.73
	hedgehog	0.20	house	0.39
	anchor	0.15	bell	0.49
	water well	0.22	key	0.42
	mean	0.19	mean	0.45
	S.D.	0.03	S.D.	0.17
meaningless versions	helicopter	0.19	ship	0.37
	scissors	0.22	iron	0.46
	aeroplane	0.21	cat	0.86
	hedgehog	0.19	house	0.37
	anchor	0.17	bell	0.49
	water well	0.27	key	0.25
	mean	0.21	mean	0.47
	S.D.	0.03	S.D.	0.21
overall	mean	0.20	mean	0.46
overall	S.D.	0.03	S.D.	0.18

High figures signify a simpler shape, and lower figures signify a more complex shape. Both perimeter lines and internal lines were used in the calculation.

APPENDIX 4: COMPLETE SET OF MODIFICATIONS FOR ROTATION TASK

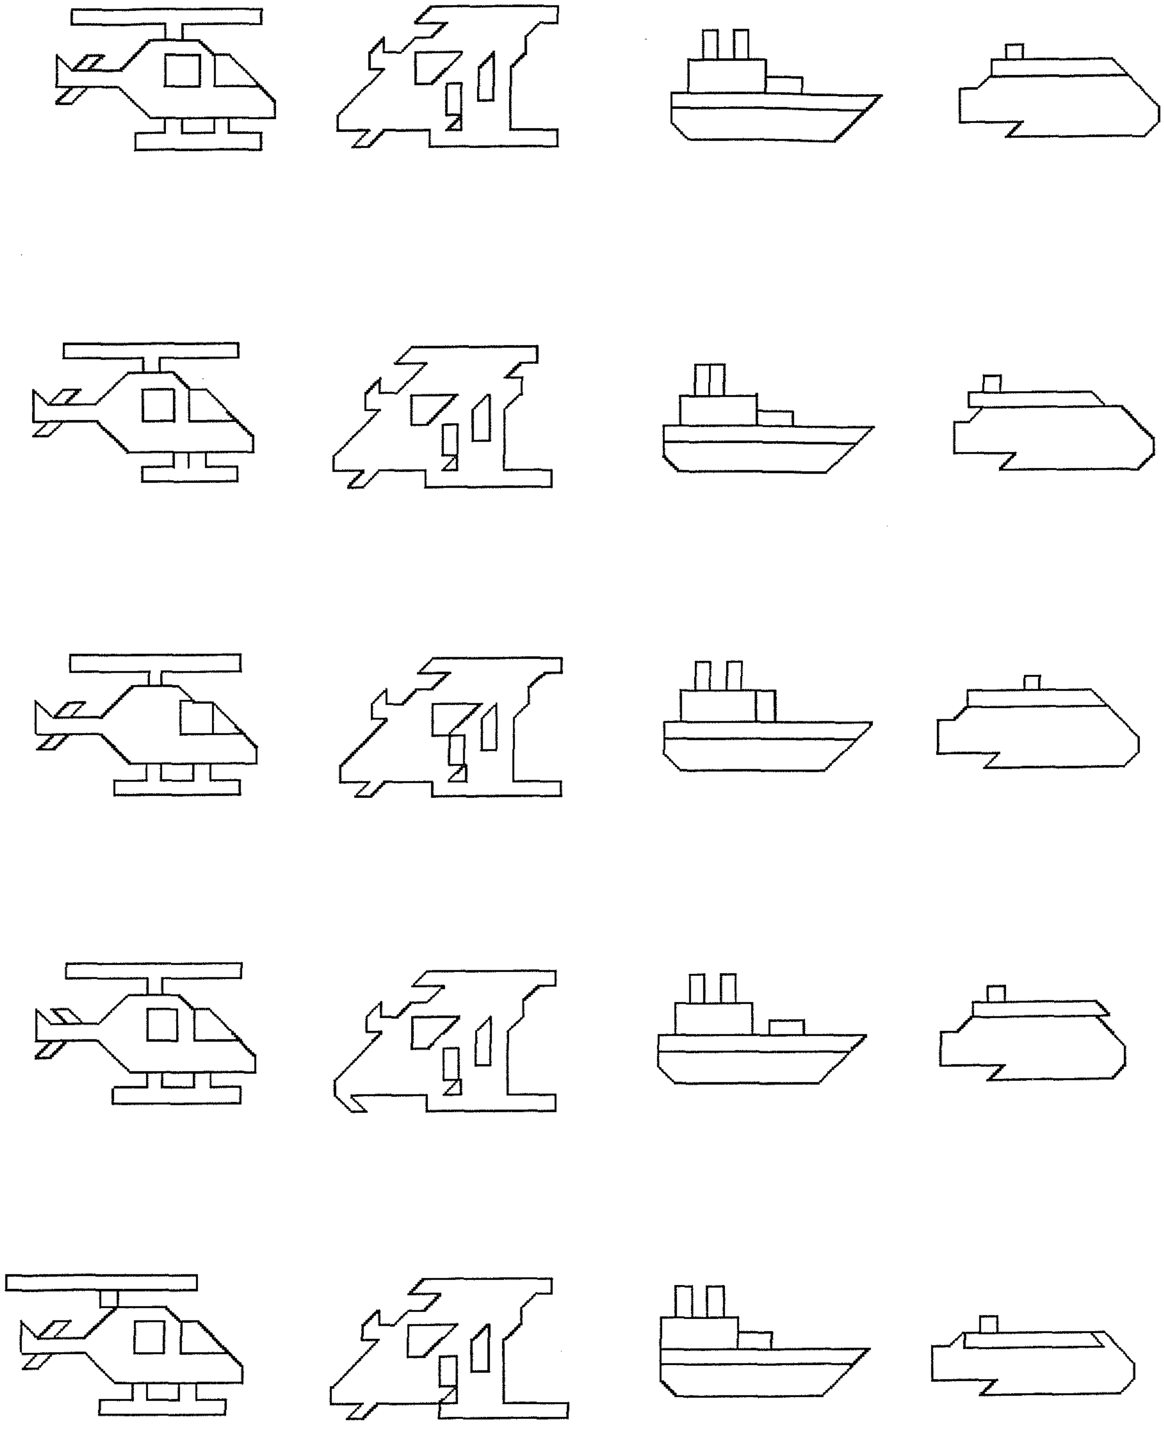

Figure 2a: Helicopter, ship, and meaningless equivalents.
A complete set of modifications is shown.
The original stimulus is shown at the top, and the four modifications are shown below the original.

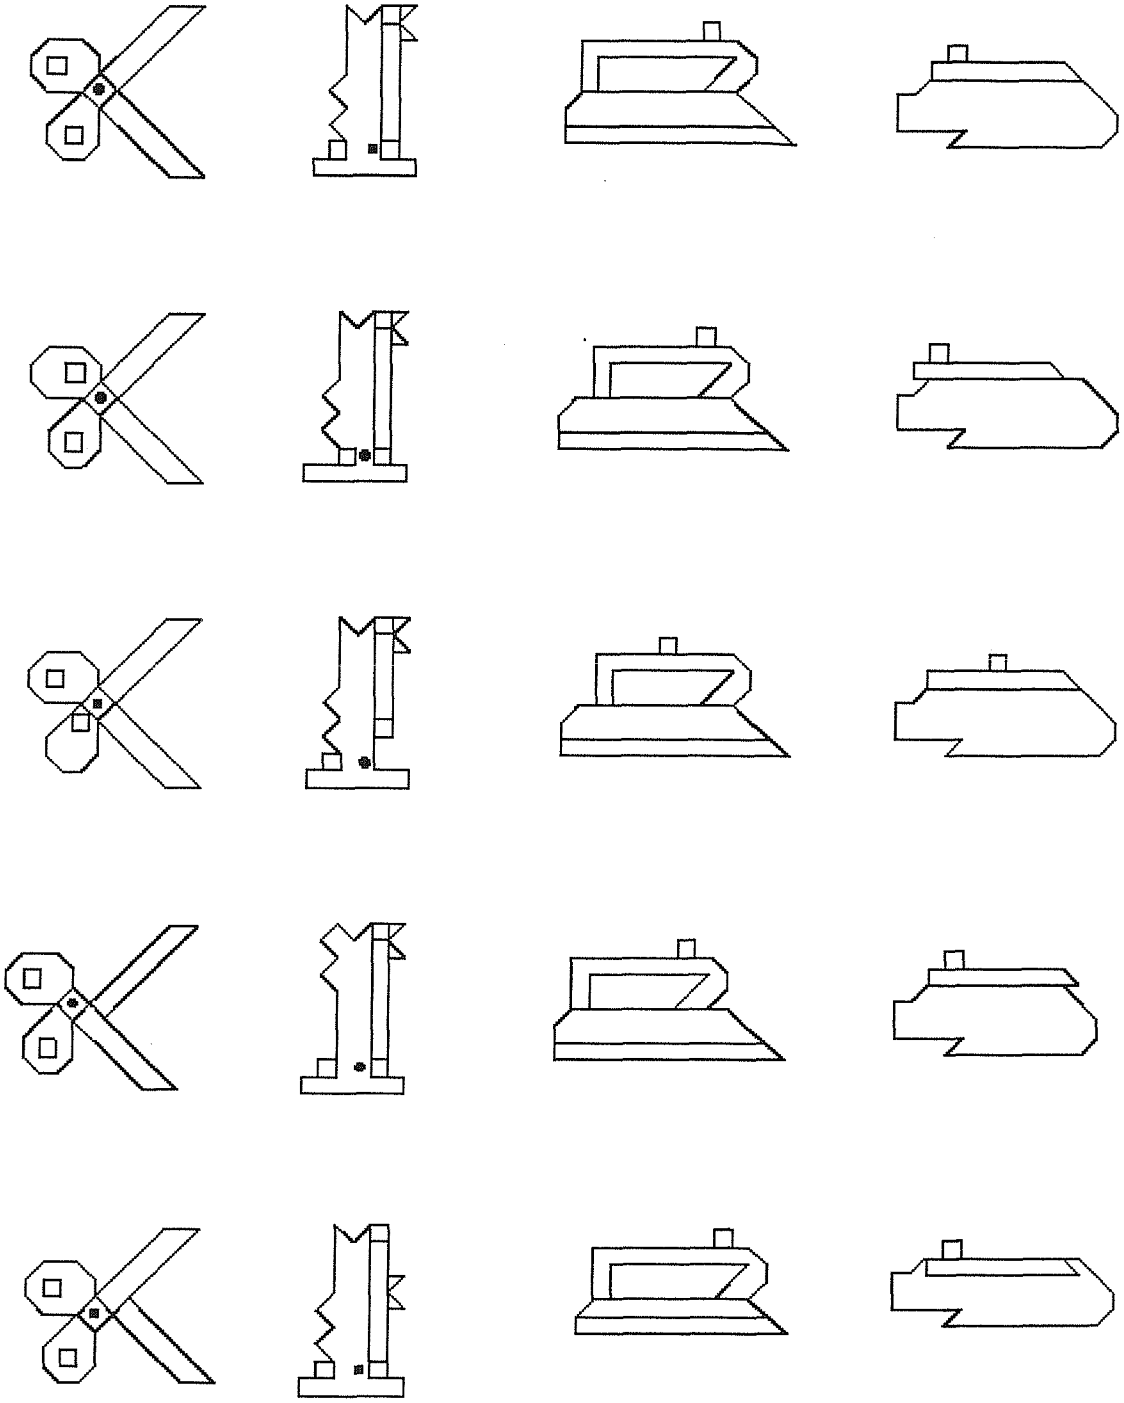

Figure 2b: Scissors, iron, and meaningless equivalents.
A complete set of modifications is shown.
The original stimulus is shown at the top, and the four modifications are shown below the original.

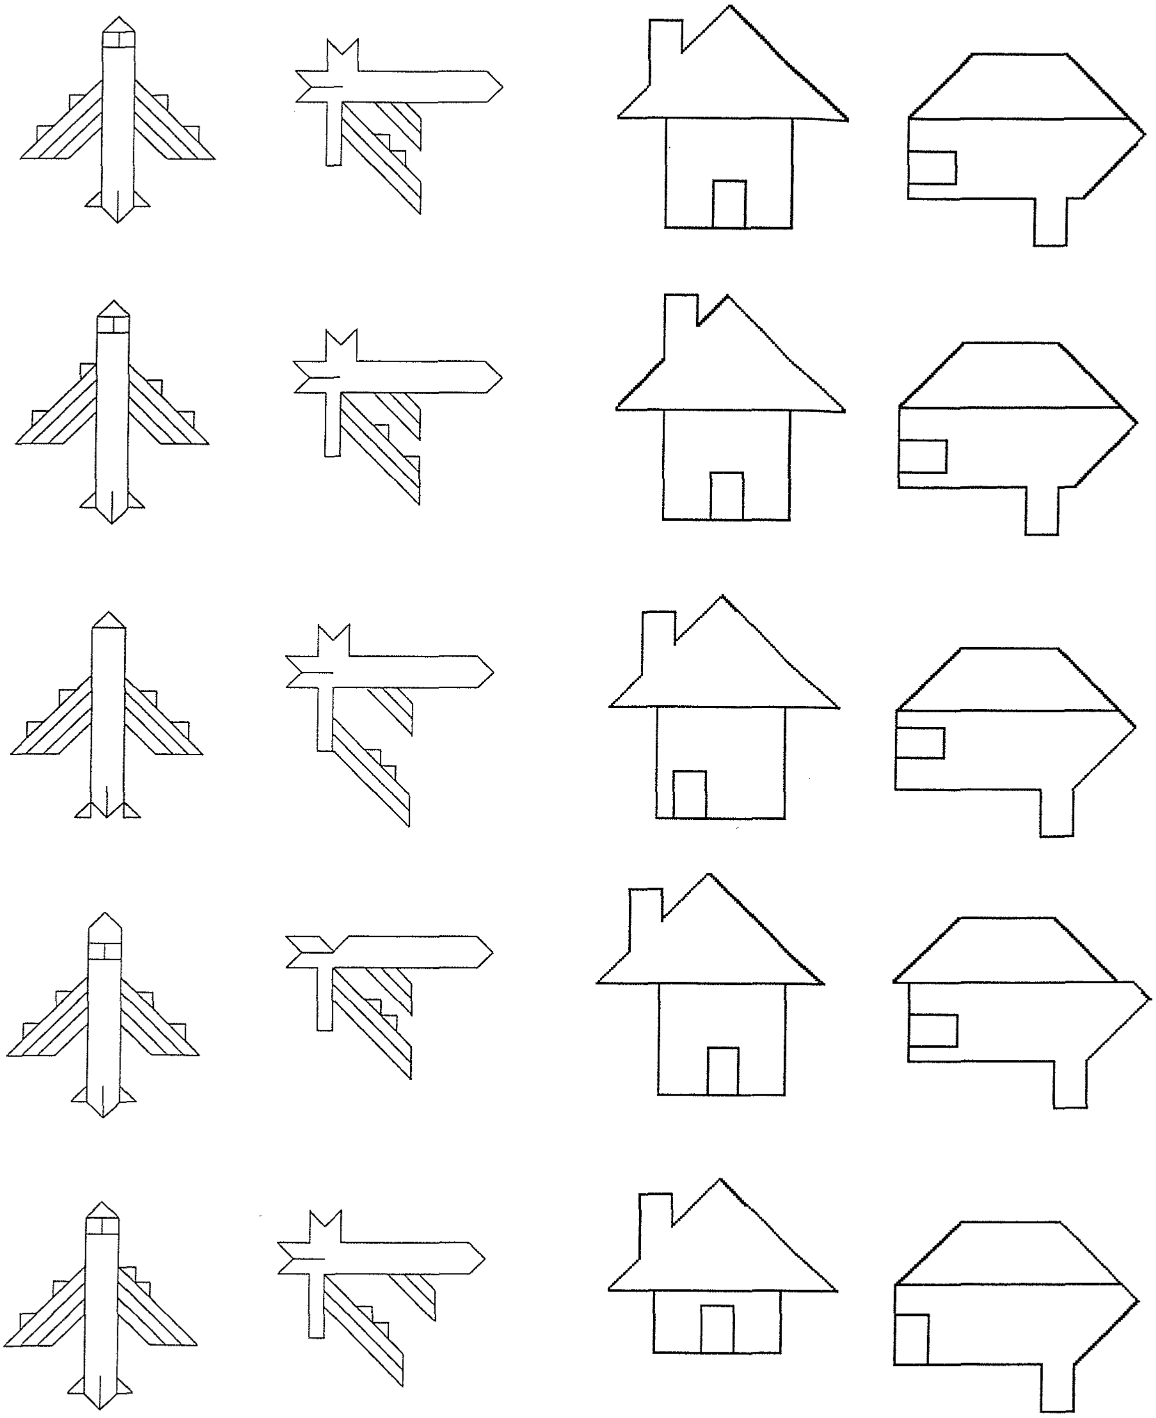

Figure 2c: Aeroplane, house, and meaningless equivalents.
A complete set of modifications is shown.
The original stimulus is shown at the top, and the four modifications are shown below the original.

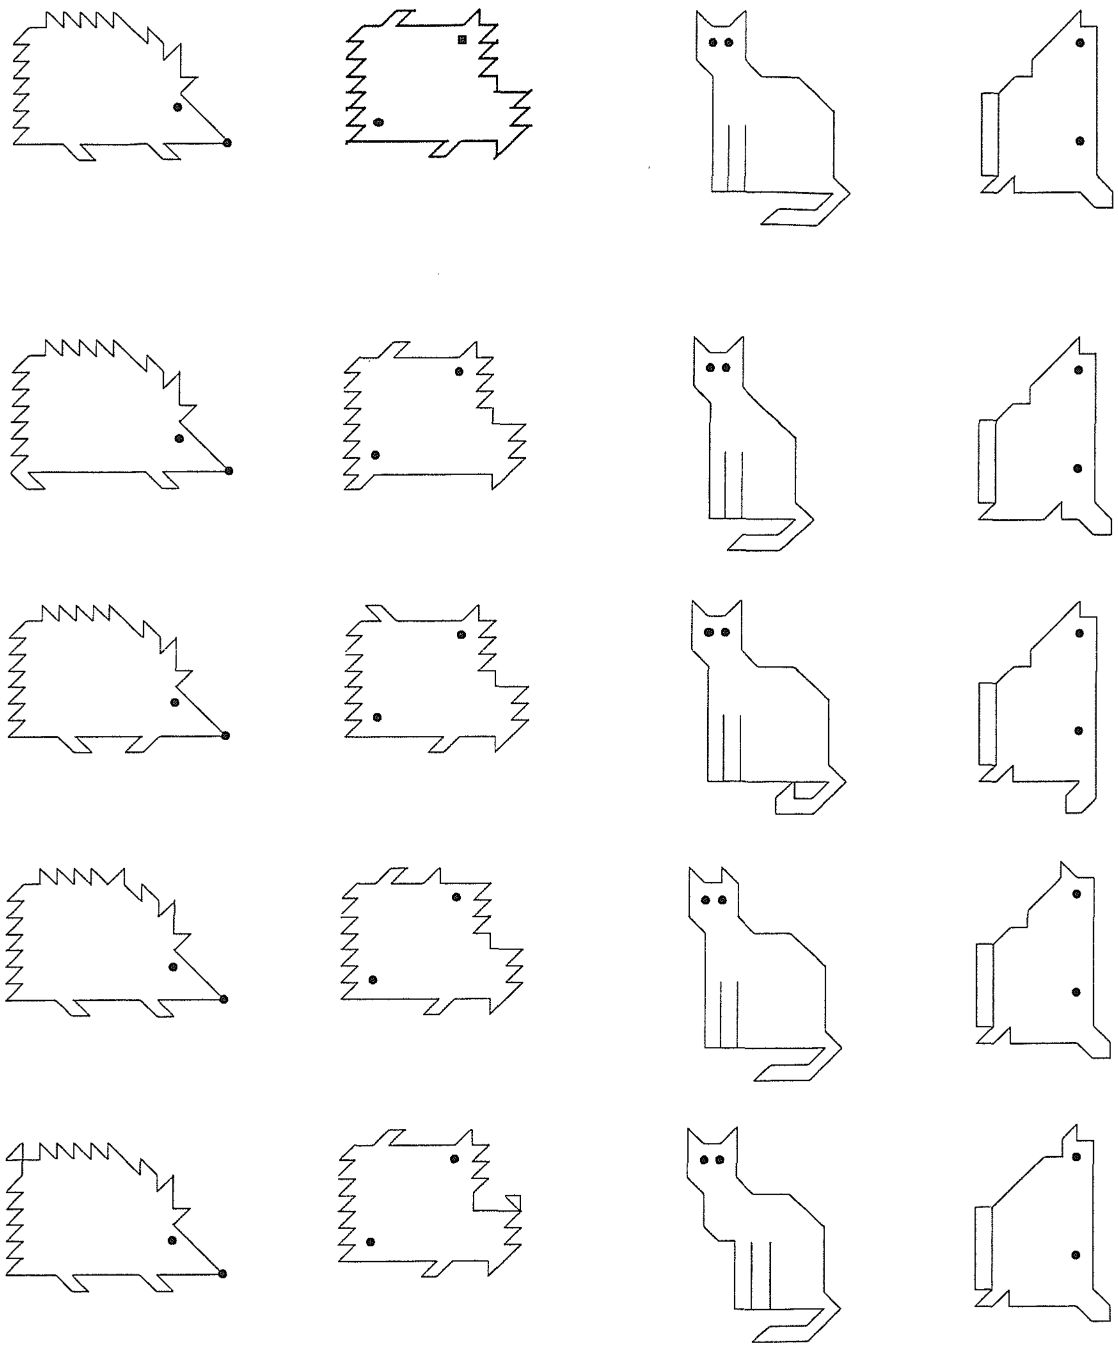

Figure 2d: Hedgehog, cat, and meaningless equivalents.
A complete set of modifications is shown.
The original stimulus is shown at the top, and the four modifications are shown below the original.

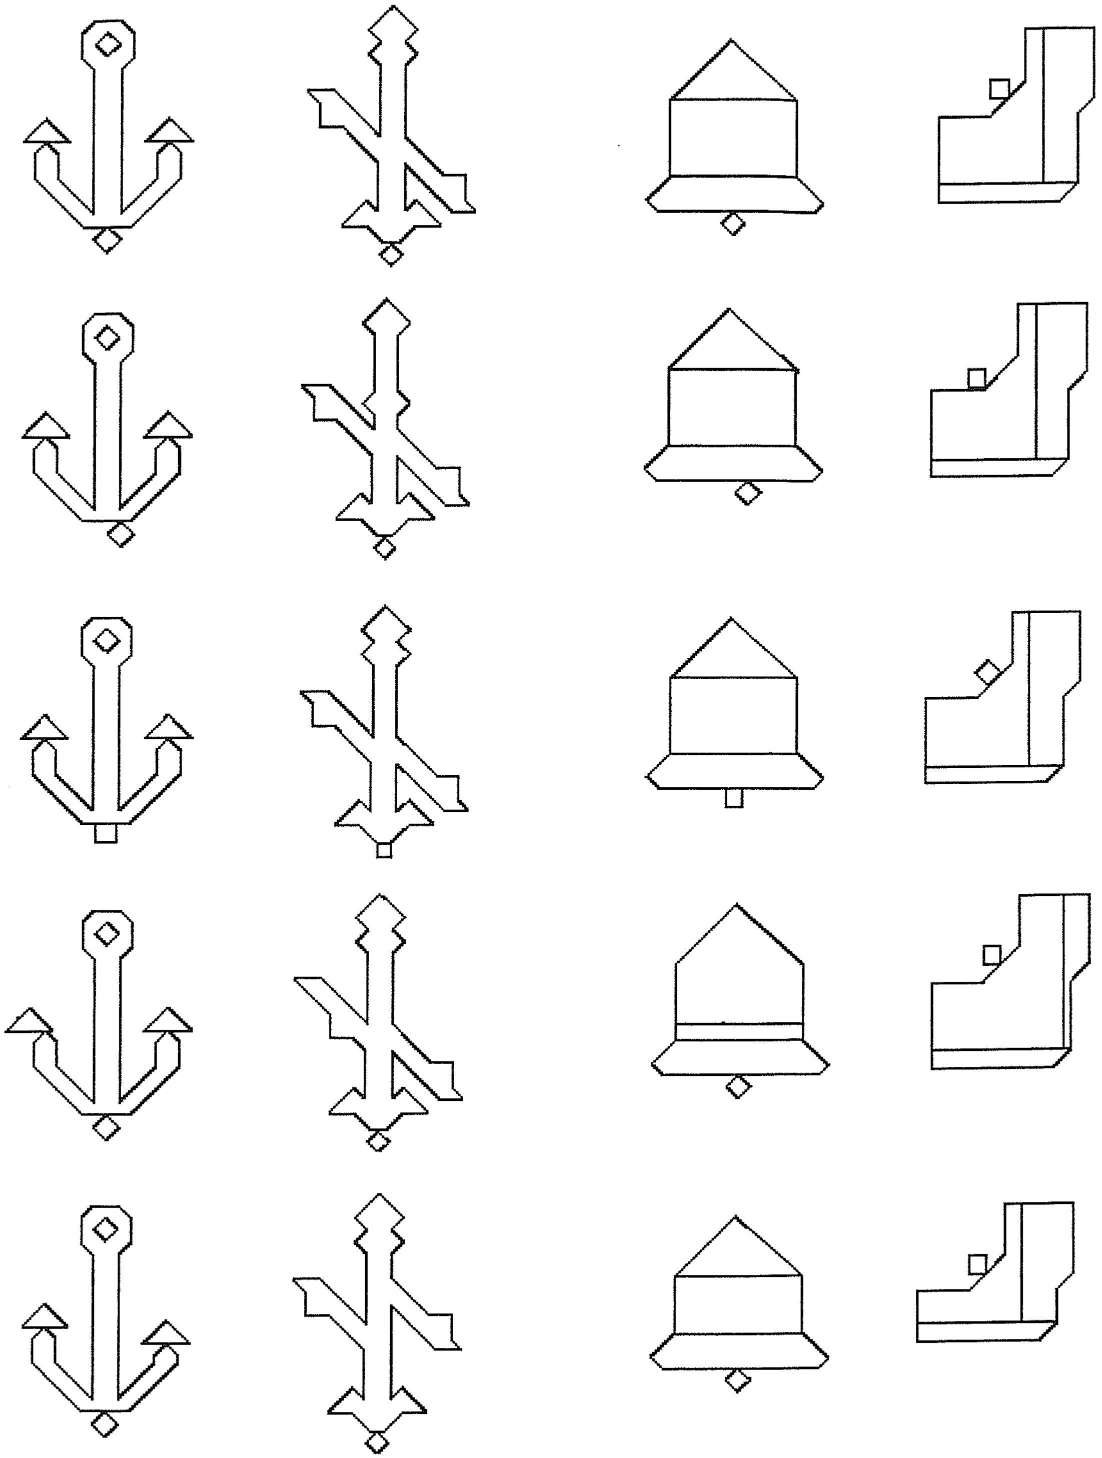

Figure 2e: Anchor, bell, and meaningless equivalents.
A complete set of modifications is shown.
The original stimulus is shown at the top, and the four modifications are shown below the original.

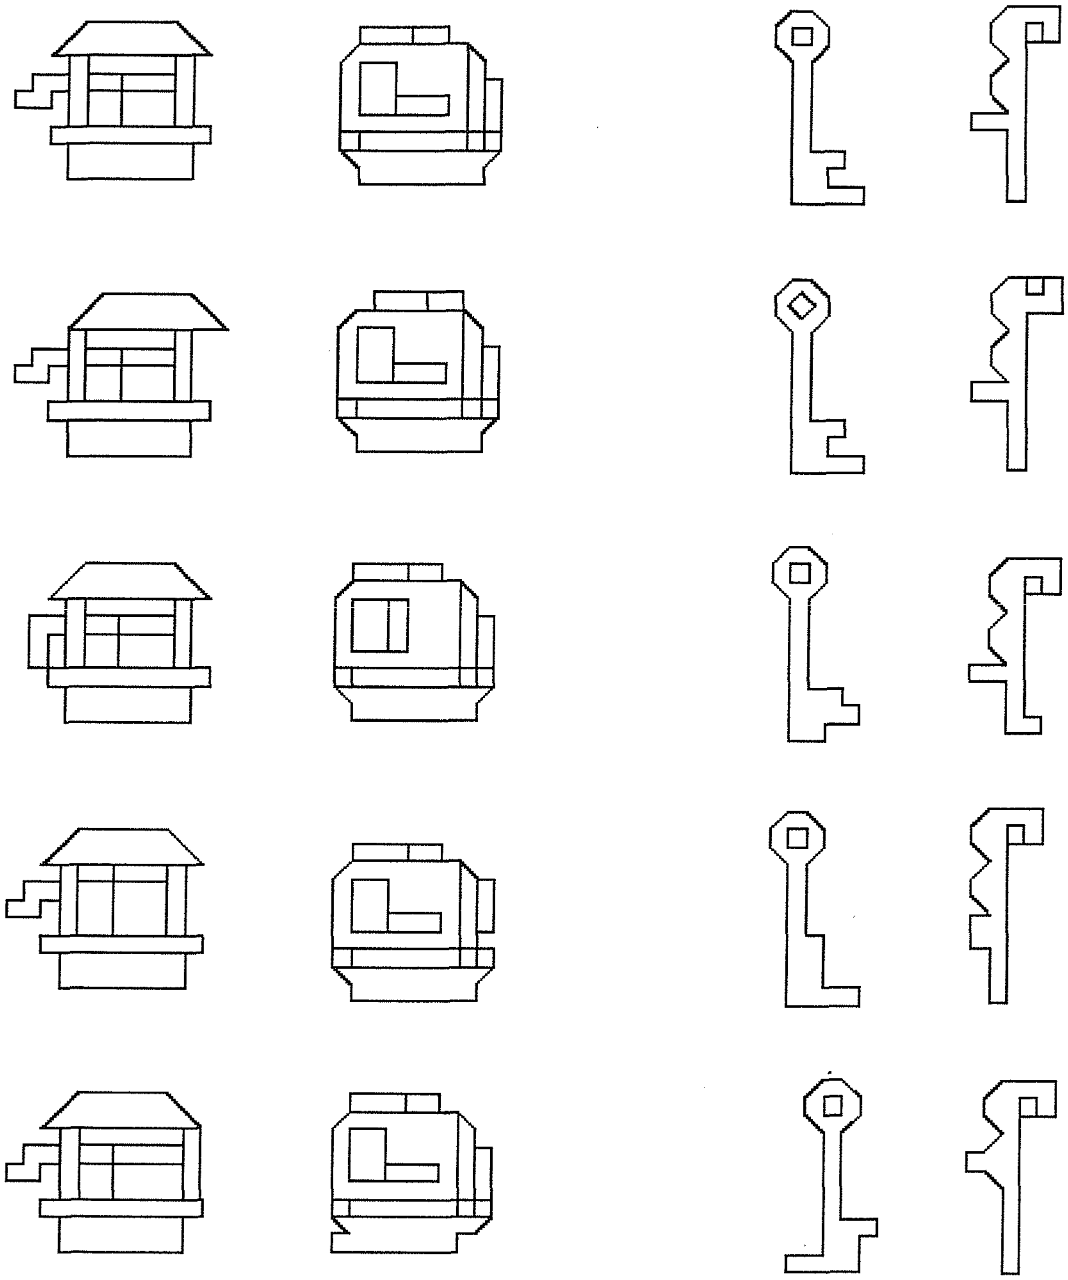

Figure 2f: Well, key, and meaningless equivalents.
A complete set of modifications is shown.
The original stimulus is shown at the top, and the four modifications are shown below the original.

APPENDIX 5: COMPLETE SET OF PARTS FOR PART SEARCH TASK

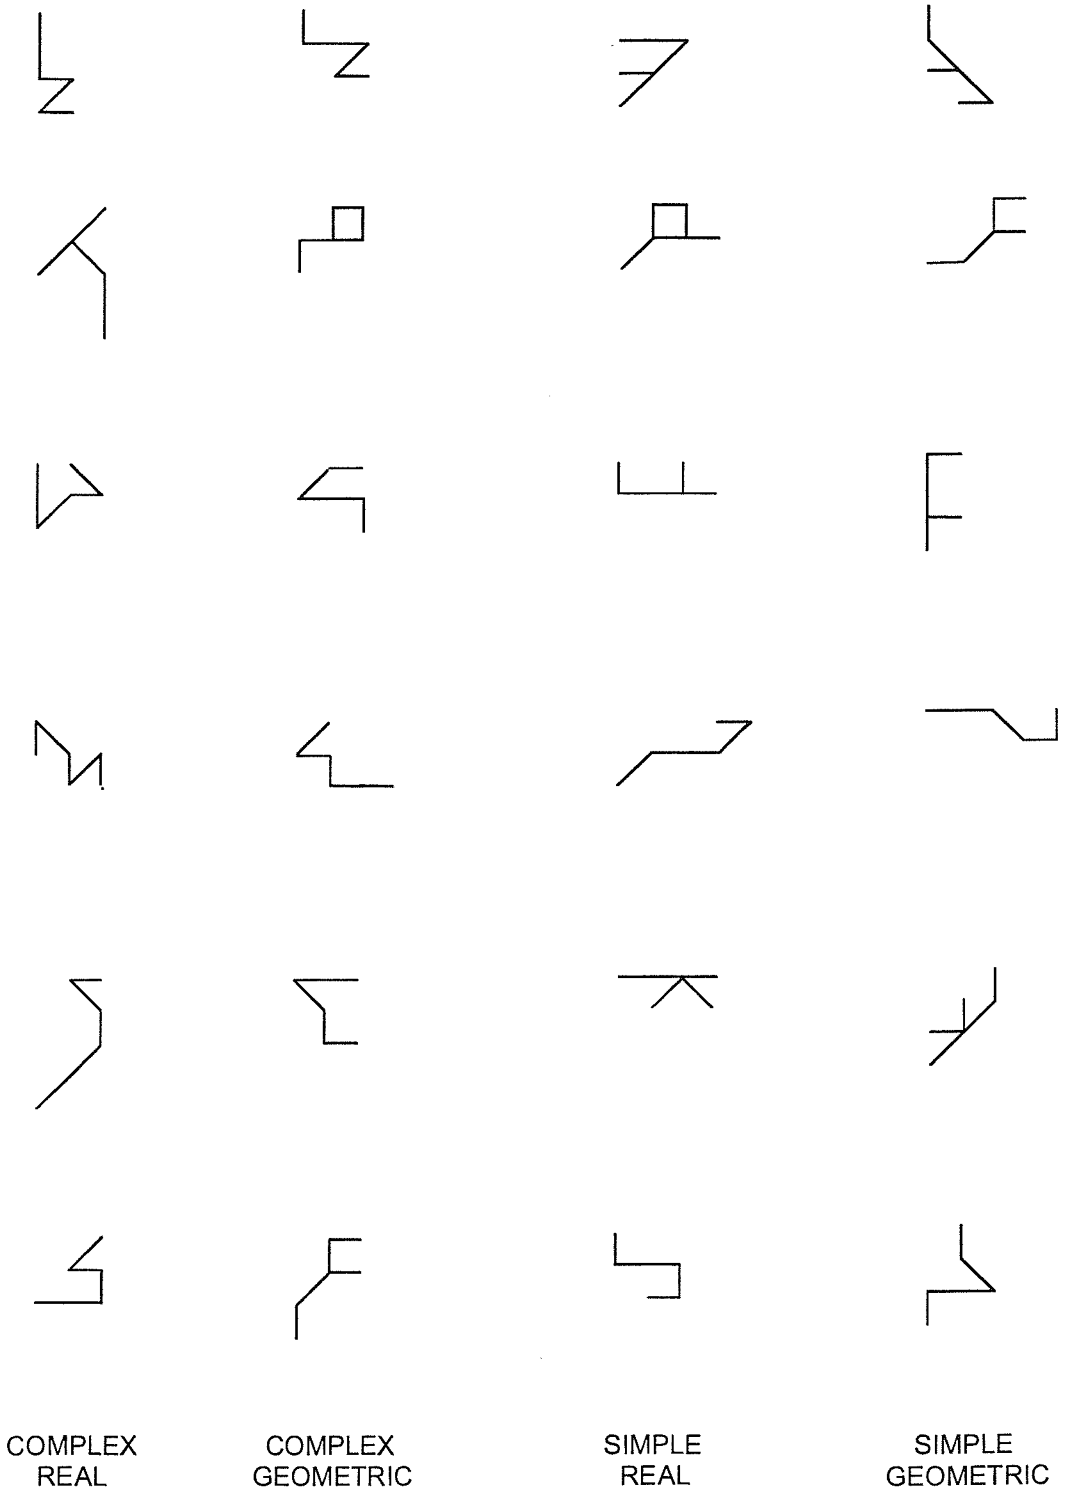

Figure 3: Complete set of parts for the part search task

The parts are shown here smaller than those presented on the computer screen.

APPENDIX 6: RELATIONSHIP OF PARTS AND WHOLES IN THE PART SEARCH TASK

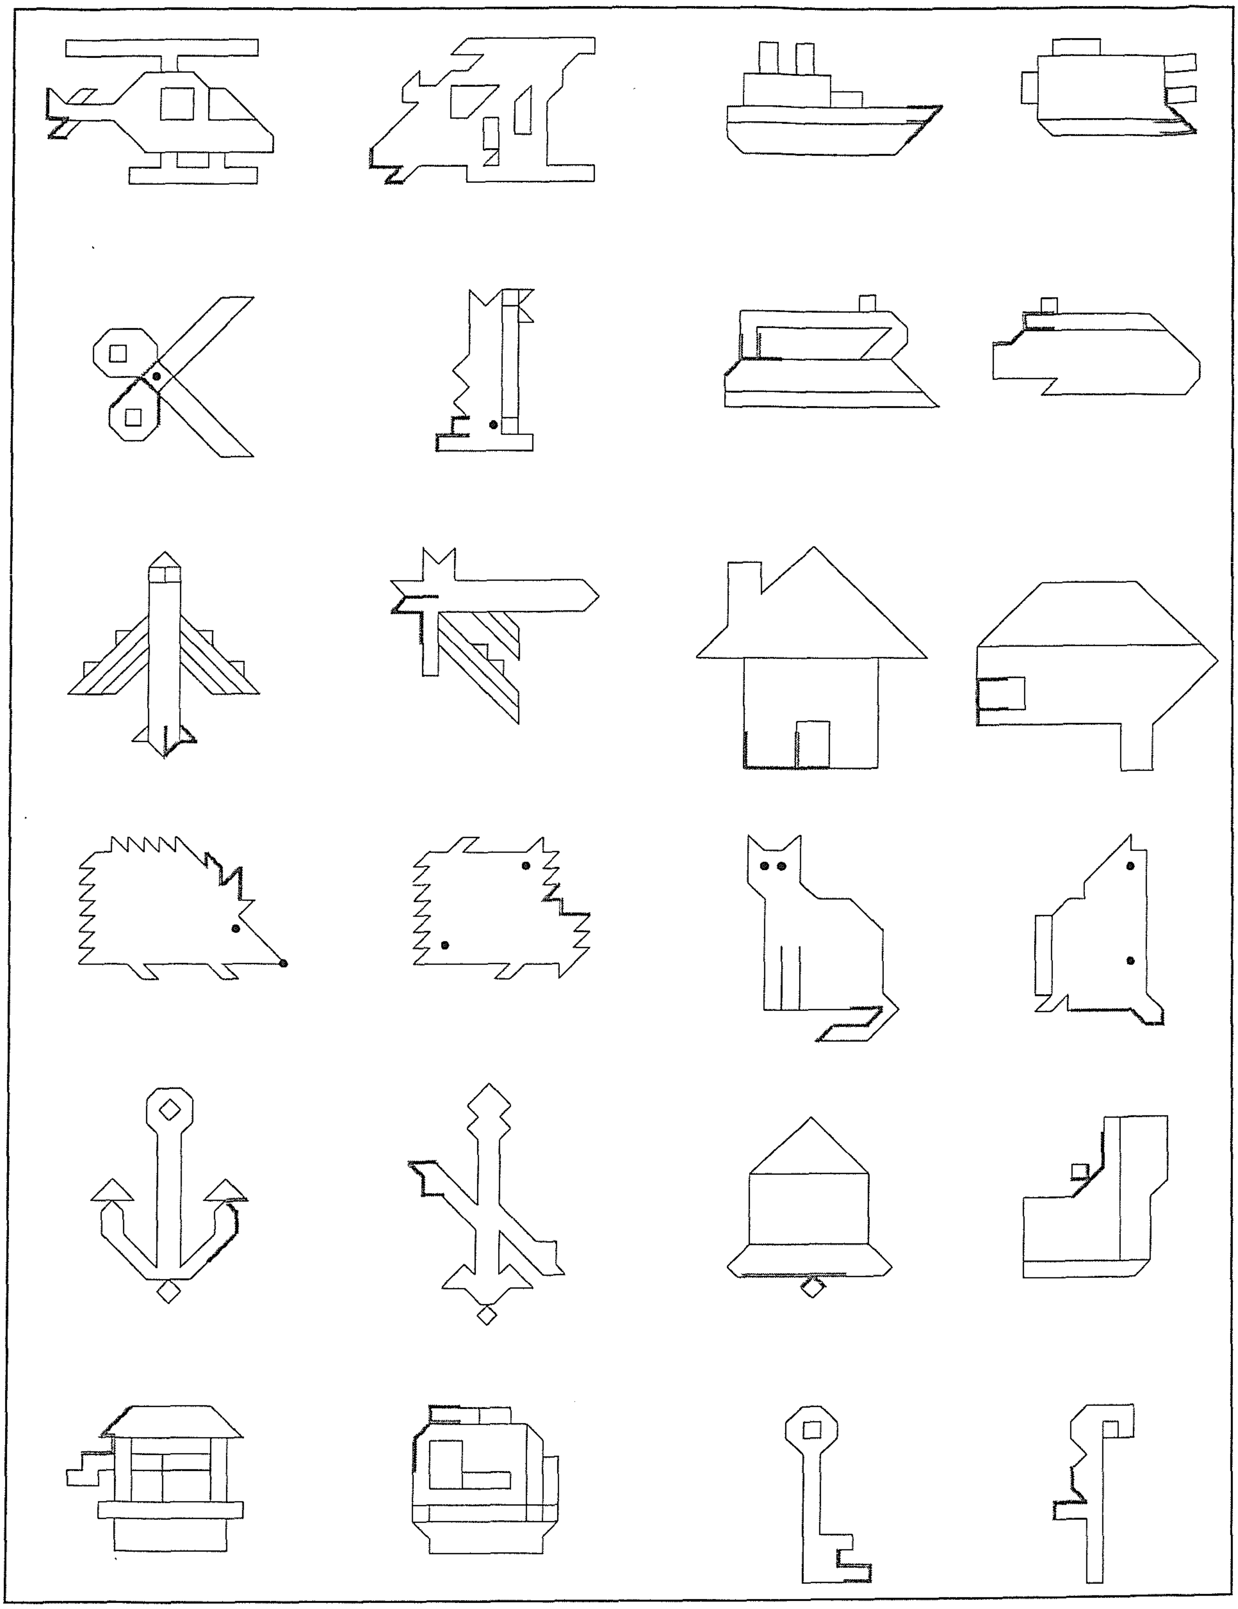

Figure 4: Relationship of parts and wholes for the part search task. This shows where in the whole stimulus the part was derived from.

APPENDIX 7: MEAN VALUES FOR THE RECOGNITION OF PARTS

Table 2: Recognition of Parts

	meaningful stimuli		meaningless stimuli	
stimulus	pre	post	pre	post
helicopter	0.00	0.00	0.00	0.00
scissors	0.04	0.14	0.00	0.00
aeroplane	0.11	0.18	0.00	0.00
hedgehog	0.04	0.54	0.00	0.00
anchor	0.00	0.07	0.00	0.00
water-well	0.00	0.14	0.04	0.07
ship	0.25	0.82	0.00	0.00
iron	0.00	0.11	0.00	0.04
house	0.00	0.54	0.04	0.04
cat	0.00	0.14	0.00	0.00
bell	0.00	0.29	0.00	0.00
key	0.00	0.25	0.00	0.00
mean	0.04	0.27	<0.01	0.01

The stimuli have been described in terms of the meaningful objects; the meaningless objects were derived from these objects. The recognition value is expressed as a proportion of the participants identifying the object that the part came from.

APPENDIX 8: PROBES FOR THE PART PROBE MATCHING TASK

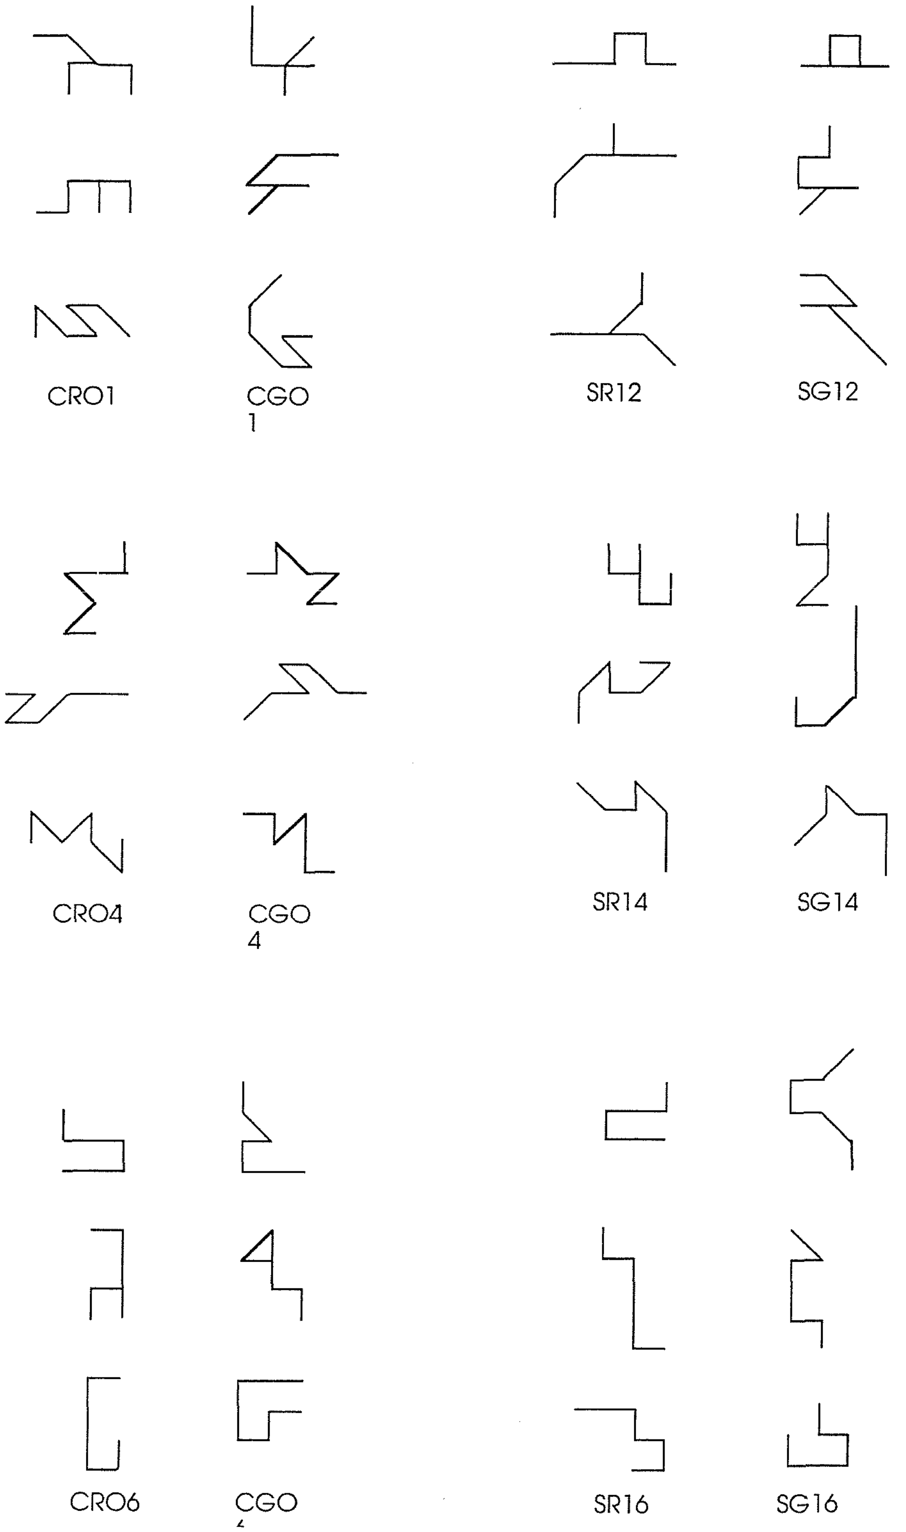

CRO1 CGO 1 SR12 SG12

CRO4 CGO 4 SR14 SG14

CRO6 CGO SR16 SG16

Figure 5: Part Probes for the Part Probe Matching Task.
The whole probes were formed by modifying a small area of the stimulus. The part probes were formed by isolating this area. Of the three probes per stimulus set, one was an abstract part, and one was a concrete part. These terms related to the level of figural goodness.

APPENDIX 9: PROBES FOR THE WHOLE PROBE MATCHING TASK

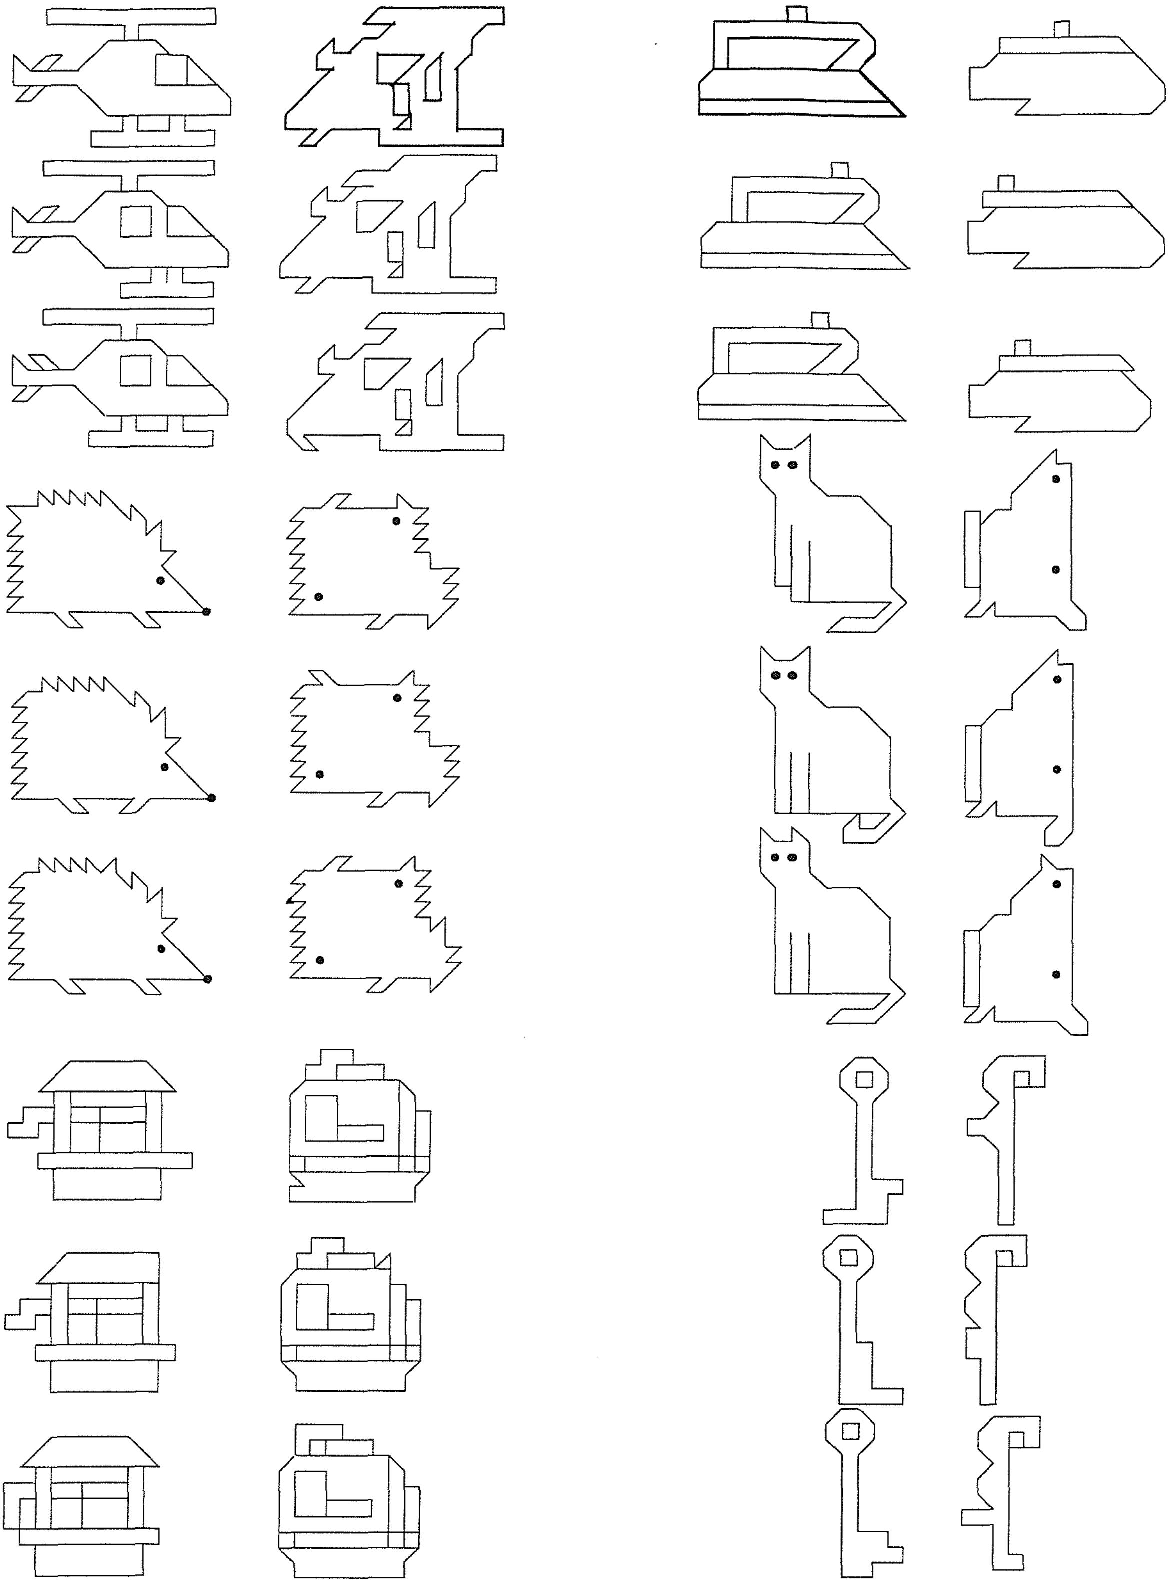

Figure 6: Whole Probes for the Whole Probe Matching Task

APPENDIX 10: RELATIONSHIP OF PARTS AND WHOLES IN THE PROBE MATCHING TASK

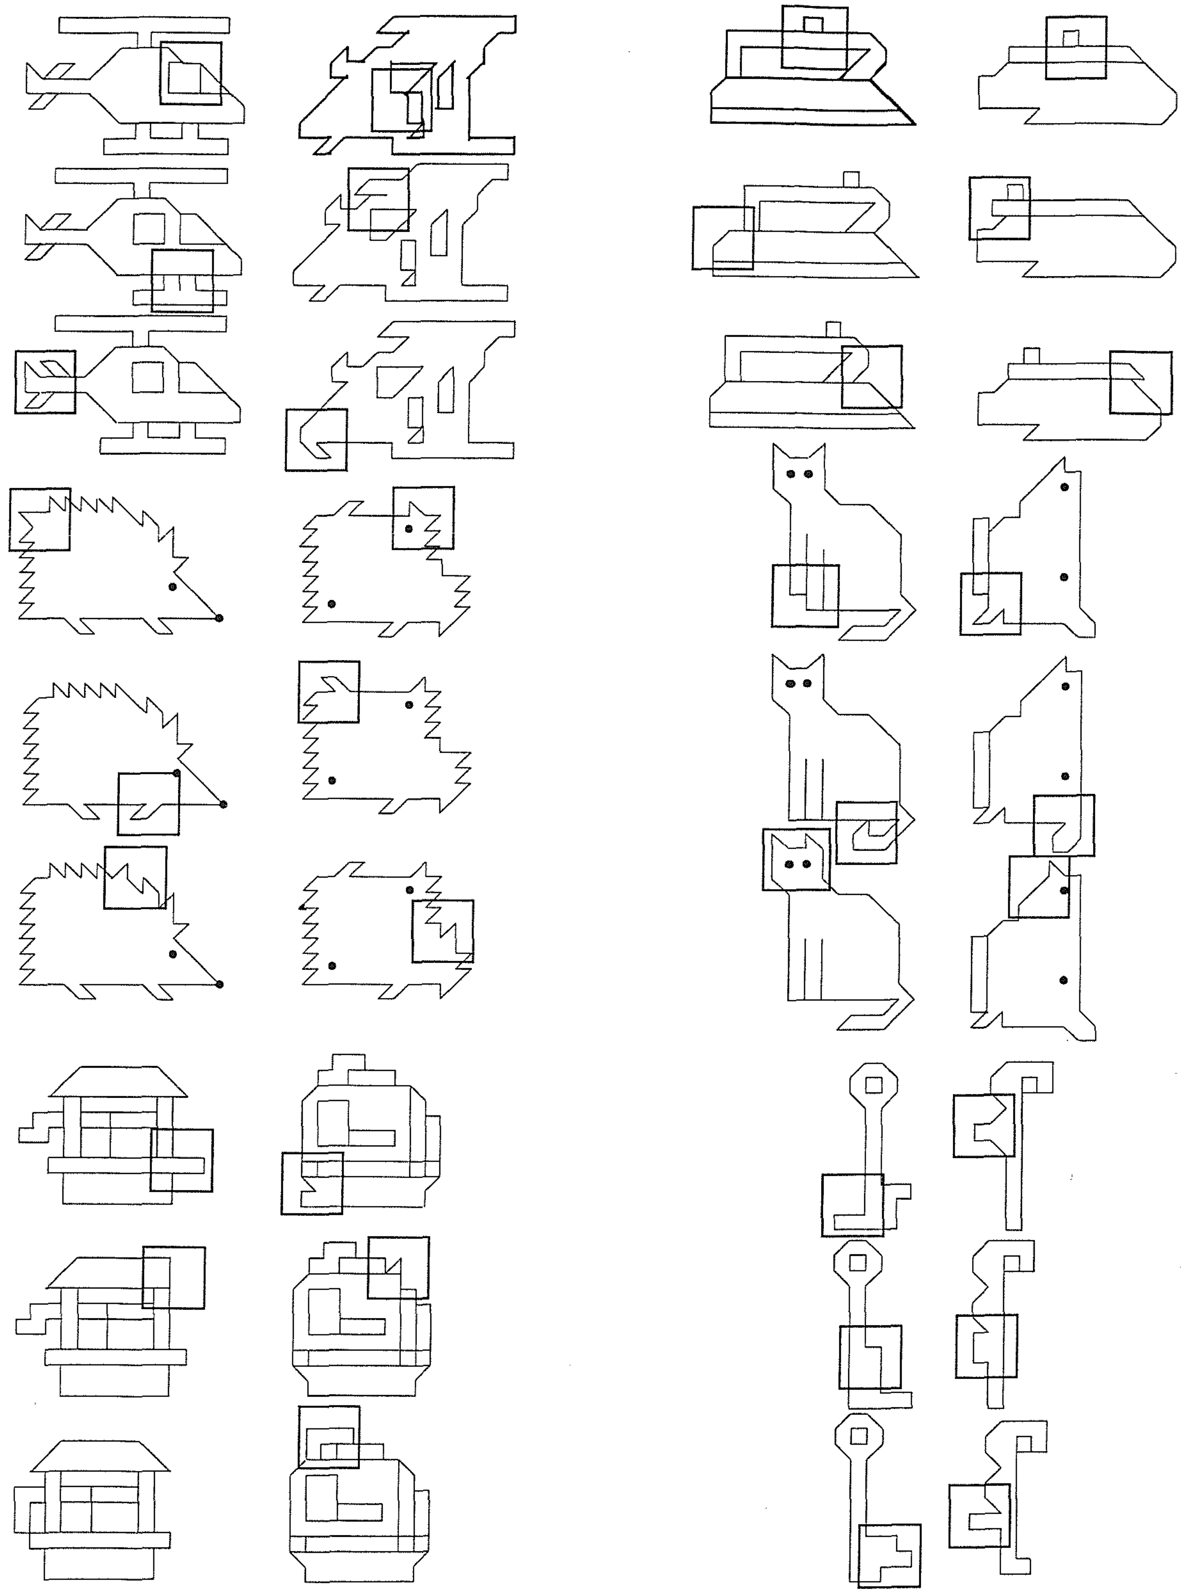

Figure 7: Relationship of Wholes and Parts in Probe Matching Task This shows the location of the part probe within the whole probe.

REFERENCES

Ankrum, C. & Palmer, J. (1991). Memory for objects and parts. *Perception and Psychophysics* **50** 141-156.

Ashby, F. G., Prinzmetal, W., Ivry, R., & Maddox, W. T. (1996). A formal theory of feature binding in object perception. *Psychological Review* **103** 165-192.

Attneave, F. (1957). Physical determinants of the judged complexity of shapes. *Journal of Experimental Psychology* **53** 221-227.

Bajo, M-T. (1988). Semantic facilitation with pictures and words. *Journal of Experimental Psychology: Learning Memory and Cognition* **14** 579-589.

Banks, W. P. & Prinzmetal, W. (1976). Configurational effects in visual information processing. *Perception and Psychophysics* **19** 361-367.

Behrmann, M., Moscovitch, M., & Winocur, G. (1994). Intact visual imagery and impaired visual perception in a patient with visual agnosia. *Journal of Experimental Psychology: Human Perception and Performance* **20** 1068-1087.

Bethell-Fox, C. E. & Shepard, R. N. (1988). Mental Rotation: effects of stimulus complexity and familiarity. *Journal of Experimental Psychology: Human Perception and Performance* **14** 12-23.

Biederman, I. & Gerhardstein, P. C. (1995). Viewpoint-dependent mechanisms in visual object recognition: reply to Tarr and Bülthoff (1995). *Journal of Experimental Psychology: Human Perception and Performance* **21** 1506-1514.

Biederman, I. & Shiffran, M. (1987). Sexing day old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, Memory and Cognition* **13** 640-645.

Biederman, I. (1972). Perceiving real-world scenes. *Science* **177** 77-80.

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review* **94** 115-147.

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1162-1182.

Blanz, V., Tarr, M. J., & Bülthoff, H. H. (1999). What object attributes determine canonical views? *Perception* **28** 575-599.

Boer, L. C. & Keuss, P. J. G. (1982). Global precedence as a postperceptual effect: an analysis of speed-accuracy tradeoff functions. *Perception and Psychophysics* **31** 358-366.

Boucart, M. & Humphreys, G. W. (1997). Integration of physical and semantic information in object processing. *Perception* **26** 1197-1209.

Brennan, T., David, D., Fluchaine, I., & Dellet, J. (1996). Naming faces and objects without comprehension: A case study. *Cognitive Neuropsychology* **13** 93-110.

Bülthoff, H. H., Edelman, S. Y. & Tarr, M. J. (1995). How are three-dimensional objects represented in the brain? *Cerebral Cortex* **5** 247-260.

Burton, A. M. & Bruce, V. (1992). I recognise your face, but I can't remember your name: A simple explanation. *British Journal of Psychology* **83** 45-61.

Carpenter, P. A. & Just, M. A. (1978). Eye fixations during mental rotation. In J. W. Senders, D. F. Fisher, & R. A. Monty (Eds.). *Eye Movements and the Psychological Functions* Hillsdale, NJ: Erlbaum.

Cave, C. B. & Kosslyn, S. M. (1993). The role of parts and spatial relations in object identification. *Perception* **22** 229-248.

Cedrus Corporation (1991). *Superlab: General Purpose Testing Package.* (Version 1.4 Macintosh Edition). .

Ceraso, J., Kourtzi, Z. & Ray, S. (1998). The integration of object properties. *Journal of Experimental Psychology: Learning, Memory and Cognition* **29** 1152-1161.

Chambers, D. & Reisberg, D. (1992). What an image depicts depends on what an image means. *Cognitive Psychology* **24** 145-174.

Chatterjee, A. & Southwood, M. H. (1995). Cortical blindness and visual imagery. *Neurology* **45** 2189-2195.

Chun, M. M. & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology* **36** 28-71.

Cochran, E. L., Pick, A. D. & Pick, Jr., H. L. (1983). Task-specific strategies of mental "rotation" of facial representations. *Memory and Cognition* 11 41-48.

Cohen, D. & Kubovy, M. (1993). Mental rotation, mental representation, and flat slopes. *Cognitive Psychology* **25** 351-382.

Cooper, E. E. & Wojan, T. J. (2000). Differences in the coding of spatial relations in face identification and basic-level object recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition* **26** 470-488.

Cooper, L. A. & Podgorny, P. (1976). Mental transformations and visual comparison processes: effects of complexity and similarity. *Journal of Experimental Psychology: Human Perception and Performance* **2** 503-514.

Corballis, M. C. (1976). Recognition of disoriented shapes. *Psychological Review* **92** 115-123.

Corballis, M. C., Macadie, L., Crotty, A. & Beale, A. (1985). The naming of disoriented letters by normal and reading-disabled children. *Journal of Child Psychology* **26** 929-938.

Corballis, P. M. & Corballis, M. C. (1988). How apparent motion affects mental rotation: Push or Pull? *Memory and Cognition* **21** 458-466.

Craver-Lemley, C. & Reeves, A. (1993). How visual imagery interferes with vision. *Psychological Review* **99** 633-649.

Cutzu, F. & Edelman, S. (1994). Canonical views in object representation and recognition. *Vision Research* **34** 3037-3056.

Davidoff, J. & Donnelly, N. (1990). Object superiority: a comparison of complete and part probes. *Acta Psychologica* **73** 225-243 .

Dewey, J. (1925). *Experience and Nature*. Chicago: Open Court.

Diamond, R. & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General* **115** 107-117.

Diamond, R. & Carey, S. (1990). On the acquisition of pattern encoding skills *Cognitive Development* **5** 345-368.

Donnelly, N. & Davidoff, J. (1999). The mental representations of faces and houses: Issues concerning parts and wholes. *Visual Cognition* **6** 319-343.

Donnelly, N., Humphreys, G. W., & Sawyer, J. (1986). Stimulus factors affecting the categorisation of faces and scrambled faces. *Acta Psychologica* **85** 219-234.

Dosher, B. A. & Rosedale, G. S. (1997). Configural processing in memory retrieval: multiple cues and ensemble representations. *Cognitive psychology* **33** 209-265.

Dror, I. E., Ivey, C., & Rogus, C. (1997). Visual mental rotation of possible and impossible objects. *Psychonomic Bulletin and Review* 4 242-247.

Earhard, B. & Walker, H. (1985). An "outside-in" processing strategy in the perception of form. *Perception and Psychophysics* **38** 249-260.

Edelman, S. & Bülthoff, H. H. (1992). Orientation dependence in the recognition of familiar and novel views of 3d objects. *Vision Research* **32** 2385-2400.

Edelman, S. (1995). Representation of similarity in three-dimensional object discrimination. *Neural Computation* **7** 408-423.

Eley, M. G. (1983). Symbol complexity and symbol identification with rotated symbols. *Acta Psychologica* **53** 27-35.

Fagot, J., & Deruelle, C. (1997). Processing of global and local visual information and hemispheric specialization in humans (Homo sapiens). and baboons (Papio papio). *Journal of Experimental Psychology: Human Perception and Performance* **23** 429-442.

Farah, M. J. (1988). Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review* **95** 307-317.

Farah, M. J., Hammond, K. M., Levine, D. N., & Calvanio, R. (1988). Visual and spatial mental imagery: dissociable systems of representation. *Cognitive Psychology* **20** 439-462.

Farah, M. J., Péronnet, F., Gonon, M. A., & Giard, M. H. (1988). Electrophysiological evidence for a shared representational medium for visual images and visual percepts. *Journal of Experimental Psychology: General* **117** 248-257.

Farah, M. J., Wilson, K. D., Drain, M, & Tanaka, J. N. (1998). What is “special” about face perception? *Psychological Review* **105** 482-498.

Farell, B. (1984). Attention in the processing of complex visual displays: detecting features and their combinations. *Journal of Experimental Psychology: Human Perception and Performance* **10** 40-64.

Farell, B. (1985). "Same" – "Different" judgements: A review of current controversies in perceptual comparisons. *Psychological Bulletin* **98** 419-456.

Faw, B. (1997). Outlining a brain model of mental imaging abilities. *Neuroscience and Behavioral Reviews* **21** 283-288.

Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind.* Cambridge, MA: MIT Press.

Folk, M.D. & Luce, R.D(1987). Effects of stimulus complexity on mental rotation rate of polygons. *Journal of Experimental Psychology: Human Perception and Performance* **13** 395-404.

Fournier, L. R., Eriksen, C. W., Bowd, C. (1998). Multiple-feature discrimination faster than single-feature discrimination within the same object? *Perception and Psychophysics* **60** 1384-1405.

Garner, W. R. (1970). Good patterns have few alternatives. Information theory's concpet of redundancy helps in understanding the Gestalt concept of goodness. *American Scientist* **58** 34-42.

Gauthier, I., Williams, P., Tarr, M. J., Tanaka, J. (1998). Training 'greeble' experts: a framework for studying expert object recognition processes. *Vision Research* **38** 2401-2428.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: Exploring the face recognition mechanism. *Vision Research* **37** 1673-1682.

Goff, L. M. & Roediger III, H. L. (1998). Imagination inflation for action events: Repeated imaginings lead to illusory recollections. *Memory and Cognition* **26** 20-33.

Hamm, J. P. & McMullen, P. A. (1998). Effects of orientation on the identification of rotated objects depend on the level of identity. *Journal of Experimental Psychology: Human Perception and Performance* 24 413-426.

Han, S. & Humphreys, G. W. (1999). Interactions between perceptual organization based on Gestalt laws and those based on hierarchical processing. *Perception and Psychophysics* **61** 1287-1298.

Han, S., Humphreys, G. W., & Chen, L. (1999a). Parallel and competitive processes in hierarchical analysis: Perceptual grouping and encoding of closure. *Journal of Experimental Psychology: Human Perception and Performance* **25** 1411-1432.

Han, S., Humphreys, G. W., & Chen, L. (1999b). Uniform connectedness and classical Gestalt principles of perceptual grouping. *Perception and Psychophysics* **61** 661-674.

Hassebrock, F., Johnson, P. E., Bullemer, P., Fox, P. W., & Moller, J. H. (1993). When less is more: representation and selective memory in expert problem solving *American Journal of Psychology* **106** 155-189.

Hodges, J. R, Greene, J. D. W. (1998). Knowing about people and naming them: Can Alzheimer's disease patients do one without the other? *Quarterly Journal of Experimental Psychology Section A-Human Experimental Psychology* **51** 121-134.

Hoffman, D. D. & Singh, M. (1997). Salience of visual parts *Cognition* **63** 29-78.

Hoffman, J. E. (1980). Interaction between global and local levels of a form. *Journal of Experimental Psychology: Human Perception and Performance* **6** 222-234.

Hole, G. J. (1994). Configurational factors in the perception of unfamiliar faces *Perception* **23** 65-74.

Hollingworth, A., Henderson, J. M. (1998). Does consistent scene context facilitate object perception? *Journal of Experimental Psychology-General* **127** 398-415.

Homa, D., Haver, B., & Schwartz, T. (1976). Perceptibility of schematic face stimuli: evidence for a perceptual Gestalt. *Memory and Cognition* 4 176-185.

Hummel, J. E. & Biederman, I. (1992). Dynamic Binding in a neural network for shape recognition. *Psychological Review* **99** 480-517.

Hummel, J. E., Stankiewicz, B. J. (1998). Two roles for attention in shape perception: A structural description model of visual scrutiny *Visual Cognition* **5** 49-79.

Humphrey, G. K. & Khan, S. C. (1992). Recognizing novel views of three-dimensional objects. *Canadian Journal of Psychology* **46** 170-190.

Ishai, A. & Sagi, D. (1997). Visual imagery: effects of short- and long-term memory *Journal of Cognitive Neuroscience* **9** 734-742.

Jolicoeur, P. (1985). The time to name disoriented natural objects. *Memory and Cognition* **13** 289-303.

Jolicoeur, P., & Milliken, B. (1989). Identification of disoriented objects: Effects of context of prior presentations. *Journal of Experimental Psychology: Learning, Memory and Cognition* **15** 200-210.

Jolicoeur, P., Corballis, M. C., & Lawson, R. (1998). The influence of perceived rotary motion on the recognition of rotated objects. *Psychonomic Bulletin and Review* **5** 140-146.

Kimchi, R. & Bloch, B. (1998). Dominance of configural properties in visual form perception. *Psychonomic Bulletin and Review* **5** 135-139.

Kimchi, R. & Palmer, S. E. (1985). Separability and intergrality of global and local levels of hierarchical patterns. *Journal of Experimental Psychology: Human Perception and Performance* **11** 673-688.

Kimchi, R. (1992). Primacy of wholistic processing and global/local paradigm: a critical review. *Psychological Bulletin* **112** 24-38.

Kimchi, R. (1994). The role of wholistic/configural properties versus global properties in visual form perception. *Perception* **23** 489-504.

Kimchi, R. (1998). Uniform connectedness and grouping in the perceptual organization of hierarchical patterns. *Journal of Experimental Psychology: Human Perception and Performance* **24** 1105-1118.

Kinchla, R. A. & Wolfe, J. M. (1979). The order of visual processing: "Top-down", "bottom-up", or "middle-out". *Perception and Psychophysics* **24** 237-242.

Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological Bulletin* **112** 500-526.

Kosslyn, S. M. (1994). *Image and Brain.* Cambridge, MA: MIT Press.

Kurbat, M. A. (1994). Structural description theories: Is RBC/JIM a general-purpose theory of human entry-level object recognition? *Perception* **23** 1339-1368.

Lamb, M. R. & Roberston, L. C. (1990). The effect of cisual angle on global and local reaction times depends on the set of visual angles presented. *Perception and Psychophysics* **47** 489-496.

Landau, B, Smith, L, Jones, S (1998). Object shape, object function, and object name. *Journal of Memory and Language* **38**.1-27.

Lanze, M., Maguire, W., & Weisstein, N. (1985). Emergent features: A new factor in the object-superiority effect? *Perception & Psychophysics* **38** 423-442.

Larsen, A. (1985). Pattern matching: effects of size ratio, angular difference in orientation, and familiarity. *Perception and Psychophysics* **38** 63-68.

Lawson, R, Jolicoeur, P (1998). The effects of plane rotation on the recognition of brief masked pictures of familiar. *Memory and Cognition* **26** 791-803.

Lawson, R. & Humphreys, G. W. (1998). View-specific effects of depth rotation and foreshortening on the initial recognition and priming of familiar objects. *Perception and Psychophysics* **60** 1052-1066.

Lawson, R. (1999). Achieving visual object constancy across plane rotation and depth rotation. *Acta Psychologia* **102** 221-245.

Leder, H. (1996). Line drawings of faces reduce configural processing .*Perception* **25** 355-366.

Leeuwenberg, E., van der Helm, P., & van Lier, R. (1994). From geons to structure. a note on object representation. *Perception* **23** 505-515.

Liter, J. C. (1998). The contribution of qualitative and quantitative shape features to object recognition across changes of view. *Memory and Cognition* **26** 1056-1067.

Liu, Z. (1996). Viewpoint dependency in object representation and recognition. *Spatial Vision* **9** 491-521.

Love, B. C., Rouder, J. N., & Wisniewski, E. J. (1999). A structural account of global and local processing. *Cognitive Psychology* **38** 291-316.

Macho, S. & Leder, H. (1998). Your eyes only? A test of interactive influence in the processing of facial features. *Journal of Experimental Psychology: Human Perception and Performance* **24** 1486-1500.

Markman, A. B. & Dietrich, E. (2000). In defence of representation. *Cognitive Psychology* **40** 138-171.

Marr, D. & Nishihara, K. (1978). Representation and recognition of the spatial organisation of three-dimensional shapes. *Philosophical Transactions of the Royal Society (London)., B200,* 269-294.

Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.

Massaro, D. W. (1973). Perception of letters, words and non-words. *Journal of Experimental Psychology* **100** 349-353.

McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception. Part 1. An account of basic findings. *Psychological Review* **88** 375-407.

McDermott, K. & Roediger III, H. L. (1994). Effects of imagery on perceptual implicit memory tests. *Journal of Experimental Psychology: Learning Memory and Cognition* **20** 1379-1390.

Mermelstein, R., Banks, W., & Prinzmetal, W. (1979). Figural goodness effects in perception and memory. *Perception & Psychophysics* **26** 472-480.

Miller, G. A. (1956). The magic number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* **63** 81-93.

Minsky, M. (1975). A framework for representing knowledge. In P. H. Winston (ed.). *The Psychology of Computer Vision* New York: McGraw-Hill.

Modigliani, V, Loverock, DS, Kirson, SR (1998). Encoding features of complex and unfamiliar objects. *American Journal of Psychology* **111** 215-239.

Morgan, M. J. (1983). Mental rotation: a computationally plausible account of transformation through intermediate steps. *Perception* **12** 203-211.

Mou, L-C., anderson, N. S., Vaughan Jr, W. S., & Rouse Jr, R. O. (1989). Recognition memory for non-object drawings. *Bulletin of the Psychonomic Society* **27** 399-401.

Murray, J. E. (1995). Imagining and naming rotated natural objects. *Psychonomic Bulletin and Review* **2** 239-243.

Murray, J. E. (1997). Flipping and Spinning: Spatial Transformation Procedures in the Identification of Rotated Natural Objects. *Memory and Cognition* **25** 96-105.

Navon, D. & Norman, J. (1983). Does global precedence really depend on visual angle? *Journal of Experimental Psychology: Human Perception and Performance* **9** 955-965.

Navon, D. (1977). Forest before trees: the precedence of global features in visual perception. *Cognitive Psychology* **9** 353-383.

Nelson, K. (1974). Concept, word, and sentence: Intercorrelations in acquisition and development. *Psychological Review* **81** 267-285.

Ngohayon, S. L., Kawahara, J., & Toshima, T. (1999). The effect of meaning on visual image segmentation. *Psychologia* **42** 170-182.

Nickerson, R. S. (1965). Short-term memory for complex meaningful visual configurations: A demonstration of capacity. *Canadian Journal of Psychology* **19** 155-160.

Norman, G. R., Brooks, L. R., Coblentz, C. L., & Babcock, C. J. (1992). The correlation of feature identification and category judgements in diagnostic radiology. *Memory and Cognition* **20** 344-355.

Paivio, A., Walsh, M., & Bons, T. (1994). Concreteness effects on memory: when and why? *Journal of Experimental Psychology: Learning Memory and Cognition* **20** 1196-1204.

Palmer, S. E. & Rock, I. (1994). Rethinking perceptual organization: the role of uniform connectedness. *Psychonomic Bulletin and Review* **1** 29-55.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology* **9** 441-474.

Palmer, S. E. (1978). Structural aspects of visual similarity *Memory and Cognition* **6** 91-97.

Paquet, L. (1991). Mental rotation of compound stimuli: the effects of task demands, practice and figural goodness. *Memory and Cognition***19** 558-567.

Perrett, D. I., Harris, M. H., & Looker, S. (1992). Use of preferential inspection to define the viewing sphere and characteristic views of an arbitrary machined tool part. *Perception* **21** 497-515.

Pezdek, K., Maki, R., Valencia-Laver, D., Whetstone, T., Stoeckert, J., & Dougherty, T. (1988). Picture memory: recognizing added and deleted details. *Journal of Experimental Psychology: Learning Memory and Cognition* 14 468-476.

Pezdek, K., Whetstone, T., Reynolds, K., Askari, N., & Dougherty, T. (1989). Memory for real world scenes: the role of consistency with schema expectation. *Journal of Experimental Psychology: Learning Memory and Cognition* **15** 587-595.

Pilon, D. J., Friedman, A. (1998). Grouping and detecting vertices in 2-D, 3-D, and quasi-3-D objects. *Canadian Journal of Experimental Psychology* **52** 114-127.

Podgorny, P. & Shepard, R. N. (1983). Distribution of visual attention over space. *Journal of Experimental Psychology: Human Perception and Performance* **9** 380-393.

Policardi, E., Perani, D., Zago, S., Grassi, F. & Làdavas, E. (1996). Failure to evoke visual images in a case of long-lasting cortical blindness *Neurocase* **2** 381-394.

Pomerantz, J. R. & Pristach, E. A. (1989). Emergent features, attention and perceptual glue in visual form perception. *Journal of Experimental Psychology: Human Perception and Performance* **15** 635 -49.

Pomerantz, J. R. & Sager, I. C. (1975). Asymmetric integrality with dimensions of visual pattern. *Perception and Psychophysics* **18** 460-466.

Pomerantz, J. R. (1981).Perceptual organization in information processing. In M. Kubovy & J. R. Pomerantz,(Eds.). *Perceptual Organization* Hillsdale, NJ: Erlbaum.

Pomerantz, J. R. (1983). Global and local precedence: Selective attention in form and motion perception. *Journal of Experimental Psychology: General* **112** 511-535.

Posner, M. I. & Mitchell, R. F. (1967). Chronometric analysis of classification. *Psychological Review* 74 29-336.

Prinzmetal, W. & Banks, W. P. (1977). Good continuation affects visual detection. *Perception and Psychophysics* **21** 389-395.

Purcell, D. G. & Stewart, A. L. (1986). The face-detection effect. *Bulletin of The Psychonomic Societ y* **24** 118-120.

Purcell, D. G. & Stewart, A. L. (1988). The face-detection effect: Configuration enhances detection. *Perception & Psychophysics* **43** 355-366.

Purcell, D. G. & Stewart, A. L. (1991). The object-detection effect: configuration enhances perception. *Perception and Psychophysics* **50** 215-224.

Reed, S. K. & Johnson, J. A. (1975). Detection of aprts in patterns and images. *Memory and Cognition* **3** 569-575.

Reinitz, M. T., Lammers, W. J., & Cochran, B. P. (1992). Memory-conjunction errors: Miscombination of stored stimulus features can produce illusions of memory. *Memory and Cognition* **20** 1-11.

Reisberg, D. & Morris, A. (1985). Images contain what the imager put there: a nonreplication of illusions in imagery. *Bulletin of The Psychonomic Society* **23** 493-49.

Rhodes, G. & McLean, I. G. (1990). Distinctiveness and expertise effects with homogeneous stimuli: Towards a model of configural coding. *Perception* **19** 773-794.

Rhodes, G. (1995). Face recognition and configural coding. In T. Valentine (Ed.). *Cognitive and Computational Aspects of Face Recognition: Exploration in Face Space*. London: Routledge.

Rhodes, G., Brake, S., & Atkinson, A. P. (1993). What's lost in inverted faces? *Cognition* **47** 25-27.

Rhodes, G., Carey, S., Byatt, G., & Proffitt, F. (1998). Coding spatial variations in faces and simple shapes: A test of two models. *Vision Research* **38** 2307-2321.

Richards, D. D., Goldfarb, J., Richards, A. L., & Hansen, P. (1989). The role of the functionality rule in the categorisation of well-defined concepts. *Journal of Experimental Child Psychology* **47** 97-115.

Riddoch, M.J. & Humphreys, G. W. (1987). Picture naming in G. W. Humphreys & W. M. J. Riddoch (Eds.). *Visual Object Processing: A Cognitive Neuropsucholgocial Approach* Hillsdale, NJ: Erlbaum.

Robertson, L. C. & Lamb, M. R. (1991). Neuropsychological contributions to theories of part/whole organization. *Cognitive Psychology* **23** 299-330.

Robertson, L. C. & Palmer, S. E. (1983). Holistic processes in the perception and transformation of disoriented figures. *Journal of Experimental Psychology: Human Perception and Performance* **9** 203-214.

Rock, I. (1974). The perception of disoriented figures. *Scientific American* **230** 78-85.

Roediger III, H. L. (1996). Memory illusions. *Journal of Memory and Language* **35** 76-100.

Rouw, R., Kosslyn, S. M., & Hamel, R. (1997). Detecting high-level and low-level properties in visual images and visual percepts. *Cognition* **63** 209-226.

Rugg, M. D., Schloerscheidt, A. M., & Mark, R. E. (1998). An electrophysiological comparison of two indices of recollection. *Journal of Memory and Language* **39** 47-69.

Saariluoma, P. & Kalakoski, V. (1998). Apperception and imagery in blindfold chess *Memory* **6** 67-90.

Saariluoma, P. & Kalakoski, V. (1997). Skilled imagery and long-term working memory. *American Journal of Psychology* **110** 177-201.

Saariluoma, P. (1992). Do visual images have Gestalt properties? *The Quarterly Journal of Experimental Psychology* **45A** 399-420.

Saiki, J. & Hummel, J. E. (1998). Connectedness and part-relation integration in shape learning. *Memory and Cognition* **26** 1138-1156.

Saiki, J. & Hummel, J. E. (1998). Connectedness and the integration of parts with relations in shape perception *Journal of Experimental Psychology: Human Perception and Performance* 24 227-251.

Sanocki, T. (1993). Time course of object identification: evidence for a global-to-local contingency. *Journal of Experimental Psychology: Human Perception and Performance* **19** 878-898.

Schacter, D.L., Cooper, L. A., & Delaney, S.M(1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: General* 119 5-24.

Schacter, D.L., Cooper, L. A., Delaney, S.M., Peterson, M.A., & Tharan, M. (1991). Implicit memory for possible and impossible objects: Constraints on the construction of structural descriptions. *Journal of Experimental Psychology: Learning Memory and Cognition* **17** 3-19.

Schyns, P. G. & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning Memory and Cognition* **23** 681-696.

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

Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology* **75** 221-242.

Shepard, R. N. & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science* **171** 701-703.

Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: psychophysical aspects. *Perception* **25** 399-424.

Smith, E. E. & Haviland, S. E. (1972). Why words are perceived more accurately than non-words. *Journal of Experimental Pscychology* **92** 59-64.

Srinivas, K. & Schwoebel, J. (1998). Generalization to novel views from view combination. *Memory and Cognition* **26** 768-779.

Stevenage, S. V. (1995). Expertise and caricature advantage. In T. Valentine (Ed.). *Cognitive and Computational Aspects of Face Recognition: Exploration in Face Space*. London: Routledge.

Suzuki, S. & Cavanagh, P. (1995). Facial organization blocks access to low-level features: an object inferiority effect. *Journal of Experimental Psychology: Human Perception and Performance* **21** 901-913.

Takano, Y. (1989). Perception of rotated forms: a theory of information types. *Cognitive Psychology* **21** 1-59.

Tanaka, J. W. & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology* **46A** 225-245.

Tanaka, J. W. & Sengco, J. A. (1997). Features and their configuration in face recognition. *Memory and Cognition* **25** 583-592.

Tarr, M. J. & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance* **21** 1494-1505.

Tarr, M. J. & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology* **21** 233-282.

Tarr, M. J. & Bülthoff, H. H. (1998). Image-based object recognition in man, monkey and machine. *Cognition* **67** 1-20.

Tarr, M. J. & Gauthier, I. (1998). Do viewpoint-dependent mechanisms generalize across members a class? *Cognition* **67** 73-110.

Tarr, M. J., Kersten, D., & Bülthoff, H. H. (1998). Why the visual recognition system might encode the effects of illumination. *Vision Research* **38** 2259-2275.

Tarr, M. J,, Williams, P., Hayward, W. G., & Gauthier, I. (1998). Three-dimensional object recognition is viewpoint dependent ..*Nature Neuroscience* **1** 275-277.

Tlauka, M. & McKenna, F. P. (1998). Mental imagery yields stimulus-response compatibility. *Acta Psychologica* **98** 67-79.

Ullman, S (1998). Three-dimensional object recognition based on the combination of views. *Cognition* **67** 21-44.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition* **32** 193-254.

Valentine, T., Hollis, J., & Moore, V. (1998). On the relationship between reading, listening, and speakingIt's different for people's names. *Memory and Cognition* **26** 740-753.

Vecera, S. P. & Farah, M. J. (1997). Is visual image segmentation a bottom-up or an interactive process? *Perception & Psychophysics* **59** 1280-1296.

Verso, J. (1988). *Picture it in Cross Stitch.* Devon : David & Charles.

Voyer, D. (1995). Effect of practice on laterality in a mental rotation task. *Brain and Cognition* **29** 326-335.

Wallace, B. (1984). Apparent equivalence between perception and imagery in the production of various visual illusions. *Memory and Cognition* **12** 156-162.

Warrington, E. K. & Taylor, A. M. (1978). Two categorical stages of object recognition. *Perception* **7** 695-705.

Watson, J. B. (1920). Is thinking merely the action of language mechanisms? *British Journal of Psychology* **11** 87-104.

Weisstein, N. & Harris, C. S. (1974). Visual detection of line segments: An object-superiority effect. *Science* **186** 752-755.

Wheeler, D. D. (1970). Processes in word recognition. *Cognitive Psychology* **1** 59-85.

Yin, R. K. (1969). Looking at upside down faces. *Journal of Experimental Psychology* **81** 141-145.

Yuille, J.C. & Steiger, J.H. (1982). Nonholistic processing in mental rotation: Some suggestive evidence. *Perception and Psychophysics* **21** 201-209.