Haptic perception of multiple objects

Strategy, saliency and numerosity

Typesetting using LaTeX

Printed by Optima Grafische Comunicatie, Rotterdam

ISBN 987-90-393-5285-4

Haptic perception of multiple objects
Strategy, saliency and numerosity

Haptische waarneming van meerdere objecten
Strategie, saillantie en aantal

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof. dr. J.C. Stoof, in gevolge het besluit van het college voor promoties in het openbaar te verdedigen op maandag 1 maart 2010 des middags te 2.30 uur

door

Myrthe Amethyst Plaisier

geboren op 13 april 1982 te Rotterdam

Promotor: Prof. dr. A.M.L. Kappers
Co-promotor: Dr. W.M. Bergmann Tiest

Contents

Chapter 1

Introduction

On a daily basis, we take our keys out of our pocket and select the correct key to open our front door. The correct key can be easily recognised through touch by its shape or by its material properties. A car key, for instance, can be recognised because it has a part that is made out of plastic. If we were not able to extract such information through touch, we would have to constantly look at our hand to know what we were holding. Imagine how annoying this would be. Still, for a long time the general idea in the field of perception has been that the haptic system (touch) was not suitable for object recognition, or at least very inferior to vision. Only in 1985 did Klatzky, Lederman, and Metzger (1985) show that touch is in fact very accurate when it comes to object recognition. Later, Klatzky and Lederman showed that haptic object recognition can already be accomplished through a short static contact of only 200 ms, for which they coined the term 'haptic glance' (Klatzky & Lederman, 1995). They suggested that the long-standing misconception of touch not being suitable for object recognition was based on studies using raised line drawings (i.e. drawings that can be explored by touch) or nonsense shapes. In their 1985 study, on the other hand, common objects like a book, a sock or an umbrella were used. These objects contain three-dimensional shape information and are made out of different materials. Therefore, these objects are rich in properties that can be extracted using touch enabling fast and accurate recognition of these familiar objects.

In contrast to common objects, recognition of even very simple drawings of a house or a tree by touch is difficult and certainly much more difficult than recognition using vision (e.g. Loomis, Klatzky, & Lederman, 1991; Magee & Kennedy, 1980). This is probably due to high demands on spatiotemporal integration and memory as these drawings are explored using the index finger only. When moving your finger along the lines of the drawing you have to remember and integrate what you have felt over time to construct the complete image. In vision the field of view is much larger than the surface of a finger facilitating image recognition. When this field of view is limited by, for instance, looking through an aperture, image recognition using vision deteriorates and becomes comparable to recognition using touch. It has also been shown that when subjects draw what they have felt after haptically exploring a raised line drawing, they often recognise the object from their drawing (Wijntjes, Van Lienen, Verstijnen, & Kappers, 2008). Adding vision does not add information that was not available to the haptic system, but apparently allows the brain to process it in a different way. In such cases, where unlike vision haptic exploration is essentially serial, vision is usually much faster and more accurate at object recognition. Note, however, that haptic exploration does not have to be serial. A common object like a teacup can be recognised very fast by enclosing it with the hand, while recognition is much more difficult when the object has to be explored using one finger only.

It has been shown that different types of hand movements, or 'exploratory procedures' (EPs), are used for haptically extracting different types of information (Lederman & Klatzky, 1987). Examples are lateral motion for roughness, static contact for thermal properties, enclosure for global shape and contour following for local shape. When hand movements are restricted, object recognition can be impaired (Lederman & Klatzky, 2004). This indicates that allowing active exploration and leaving exploratory movements largely unconstrained can be important for investigating performance of the haptic system. Also the design of the stimulus is important. If the stimulus is poor in features that can be haptically extracted haptic perception may be slow or inaccurate. Note, however, that this does not mean that the haptic system is generally slow or inaccurate. When trying to compare haptic and visual perception one has to be aware that the experimental design can easily lead to an a priori advantage for vision at performing a certain task.

The studies discussed so far were concerned with recognition of a single object. This introduction, however, started with an example of selecting the correct key to open the front door. In that situation one has to recognise the correct key among the other keys. The difficulty of this task will depend on how similar the keys on the key chain are. If one key is made out of a material that is dissimilar in terms of, for instance, texture or thermal conductivity from that of the other keys, that particular key can be easy to find among the other ones. Tasks in which one has to search for a target object among distractor objects are generally referred to as search tasks. Besides finding the correct key, we can also perceive how many keys we have in our hand, i.e. judge their numerosity. Compared to vision, relatively little is known about haptic search or haptic numerosity judgement.

In the sparse haptic studies that are available on search and numerosity judgement, stimuli were often presented to the fingers only. An example of such a haptic search task is one in which materials were pressed onto the subject's separate fingers and the subject had to indicate whether a certain target material was present (Lederman & Klatzky, 1997). In the case of numerosity judgement, subjects had to indicate how many fingers were stimulated (Riggs et al., 2006). These studies are difficult to compare to daily-life haptic exploration because exploratory movements were restricted to small finger movements. Usually we explore a surface by moving our hand over it or three-dimensional objects by enclosing them in the hand. These situations are far more complex, because objects can be in contact with several parts of the hand. Furthermore, in the case of objects grasped in the hand, the objects can be freely rearranged in the hand.

1.1 Aim and structure of this thesis

As was pointed out at the beginning of this Introduction, we interact with all kinds of objects through touch on a daily basis. These objects vary in a wide range of physical properties. The haptic system is able to extract many of these physical properties like shape (e.g. Kappers, Koenderink, & Lichtenegger, 1994), weight (e.g. Jones, 1986), volume (Kahrimanovic, Bergmann Tiest, & Kappers, in press), roughness (e.g. Lederman, 1981), friction (e.g. Grierson & Carnahan, 2006) and thermal properties (e.g. Jones & Ho, 2008). Most previous studies were aimed at investigating how well we can haptically perceive a single property, or feature, in isolation. However, in daily life when we are interacting with multiple objects, we can recognise the object we are trying to find by a certain feature that makes that particular object stand out among the other objects (i.e. saliency). At the same time we also have to decide which parts belong to one object and which to another object (i.e. individuate the objects). Haptic feature saliency and object individuation both have received little attention in the literature until now. This thesis aims at providing insight into both of these processes.

To this end, a series of psychophysical studies was designed in which saliency and individuation were investigated by means of search tasks and numerosity judgement tasks. In all of these studies performance was measured in terms of response times. This means that the time needed for subjects to determine whether a certain object is present (search task) or how many items are present (numerosity judgement task) is measured. Using these tasks, it was investigated which physical object features play a role in recognition and individuation of objects through touch. Performance in these haptic tasks was compared to performance in similar visual tasks. Visual studies are used as a starting point to investigate similar questions in the haptic domain. Similarities in performance between the two modalities will be discussed in terms of modality independent processing of information.

In the remainder of this Introduction, search tasks and numerosity judgement tasks will be discussed in more detail and an overview of the different chapters will be given. Because there is a vast amount of visual research done on search as well as numerosity judgement, each section will start with a brief overview of what is known from vision.

1.2 Search tasks

1.2.1 Visual search

In visual search tasks, usually a number of items is shown on a screen. The observers are instructed to respond as fast as possible whether a certain target item is present among the other (distractor) items and response times are recorded. A target item is present in half of the trials and the total number of items is varied. This way the response times are measured as a function of the number of items. This function is usually linear and the slope represents the time needed per extra item. Therefore, the slope can be interpreted as the efficiency at which the search task was performed. Note, however, that there are two slopes; one for the target present trials and one for the target absent trials. These slopes do not necessarily have the same value.

If both the target present and absent response time slopes are near zero, the search is said to be performed in parallel because there is no extra time needed per extra item. This is interpreted as an indication that all items were processed simultaneously and the target item is said to 'pop-out' (Treisman & Gelade, 1980; Treisman & Souther, 1985). An example of such a parallel search task is searching for a red target dot among green distractor dots. However, when searching for an S among mirrored Ss response times increase with the number of items and search is said to be performed serially. This distinction between parallel and serial search is not as clear-cut as it seems. Most search tasks are performed at an efficiency somewhere in between parallel and serial. In fact, there exists a continuous range of response time slopes (Wolfe, 1998).

The absence of a clear-cut distinction between parallel and serial search does not prevent us from interpreting the slopes in terms of search efficiency. By choosing certain target and distractor item combinations, it is possible to investigate which features are salient and which are not. If a target item is distinguished from the distractor items by a salient feature, search efficiency will be high. When the distinguishing feature is not salient, search efficiency will be low. This way, search tasks can be used to provide insight into the role of certain object features in object recognition.

1.2.2 Haptic search

Compared to visual search, almost nothing is known about haptic search. In the few studies that are available, mostly items were presented to the fingers of an observer (Lederman, Browse, & Klatzky, 1988; Lederman & Klatzky, 1997; Overvliet, Smeets, & Brenner, 2007a, 2007b; Overvliet, Mayer, Smeets, & Brenner, 2008). In these cases items consisted of different materials or raised line drawings. What these previous studies have shown is that features that may be very salient in the visual domain, are not necessarily salient features for the haptic system. In vision, search for a line with a different orientation than the surrounding lines is very efficient (Treisman, 1985), while search for a ridge among among ridges with a different orientation is not efficient in touch (Lederman & Klatzky, 1997). Material properties like roughness, on the other hand, are very salient for the haptic system (Lederman & Klatzky, 1997).

Presenting items to the fingers only may be a very controlled way of stimulus presentation, but exploratory movements are very limited. As was mentioned before, restricting exploratory movement may impair haptic perception. Furthermore, in daily life we usually explore objects using our whole hand. The studies presented in this thesis were designed such that few restrictions were put on exploratory movements. Not only response times were measured, as was often done in previous studies, but also the exploratory strategies were analysed. To this end, two types of stimuli were designed. The first type is a haptic display consisting of a plane over which items could be distributed. In Chapter 2 this was a wooden display on which pieces of sandpaper could be placed (Figure 1.1a). This display could be explored using the whole hand. Hand movements were recorded and characterised. In this case the target item differed from the distractor items in roughness and subjects had to respond whether the target item was present. In Chapter 3, a virtual display rendered using a force-feedback device was used (Figure 1.1b). In that case, items consisted of areas with larger friction coefficients than the background of the display. Subjects had to determine whether the target item, which had a larger friction coefficient than the distractor items, was present. This type of display can be explored only with the index finger that was placed in a thimble-like holder connected to the force-feedback device. Note that in Chapter 2, movements using the

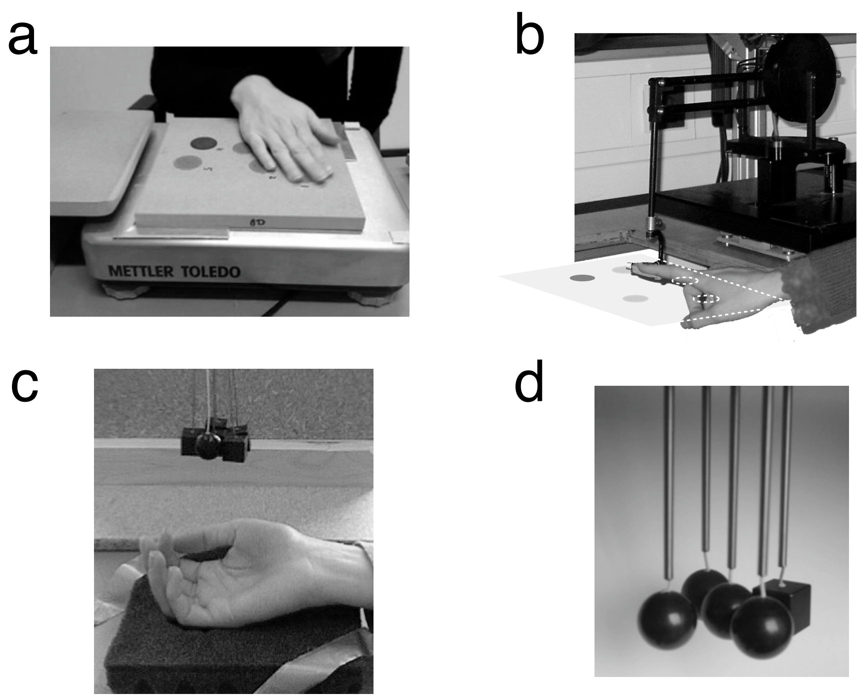

Figure 1.1: Overview of the different types of stimuli used in the search tasks in this thesis. a) A subject exploring a display from Chapter 2. In this case the subject was searching for a target item that was less rough than the distractor items. b) Exploration of a virtual display with the index finger placed into a thimble-like holder connected to a force-feedback device (Chapter 3). In this case the subject had to find a target item with a larger friction coefficient than the distractor items. c) A subject grasping upwards to enclose the three-dimensional shapes suspended above the hand (Chapter 4). In this case the subject had to determine whether a sphere was presented among the cubes. d) Shapes were suspended such that only rotation and small translations were possible (Chapter 5).

whole hand enabled a parallel search strategy. In the case of the virtual display, search could only be performed serially. To assist haptic exploration, in Chapter 3 also a visual stimulus spatially aligned with the haptic stimulus could be shown. This allowed us to investigate how visual spatial information about the display can be used to guide haptic exploration.

The second type of stimulus consisted of sets of three-dimensional shapes (spheres, cubes, tetrahedrons, ellipsoids and cylinders) that were suspended from flexible wires (Figure 1.1c). In this case, subjects grasped the shapes simultaneously and had to respond whether a certain target shape was present. Note that the observer has active control over the item positions in the hand and items can be released from the hand. This type of stimulus was used in Chapter 4 to investigate saliency of shape features by comparing several target-distractor shape combinations. To investigate if active control over the item position was important for performing such search tasks, the item positions were (partly) fixed in space in Chapter 5. The items were fixed in a way that allowed rotation and small translations (Figure 1.1d) or they were rigidly fixed. Search times were compared between both situations.

1.3 Numerosity judgement

1.3.1 Visual numerosity judgement

In visual numerosity judgement studies, usually a set of items is presented on a screen and observers are asked to report the number of items. Similar to search tasks, response times are measured as a function of the number of items. In numerosity judgement, performance is generally error-free and response times are small for numerosities up to 3 or 4 items. For larger numerosities, error rates and response times increase rapidly (e.g. Atkinson, Campbell, & Francis, 1976; Mandler & Shebo, 1982; Trick & Pylyshyn, 1993, 1994). The resulting behaviour of the response times as a function of the number of items can be described using a function consisting of two linear parts (bilinear function). The fast and accurate mechanism used for small numbers of items is known as ‘subitizing’, while the slower mechanism used for larger numerosities is usually referred to as ‘counting’ (Kaufman, Lord, Reese, & Volkmann, 1949). Consequently, the slope of the first part of the bilinear function can be interpreted as the subitizing slope and the second as the counting slope. Counting has a serial character and is less efficient than subitizing. It remains, however, unclear what kind of mechanism subitizing exactly is.

In the visual domain several suggestions as to the nature of subitizing have been made. It has been suggested that subitizing is based on pattern recognition (Mandler & Shebo, 1982). Pattern recognition could mediate fast recognition of certain numerosities. Others have suggested that it is caused by the fact that the relative differences between subsequent numerosities are large for small numerosities. Furthermore, there is a Weber fraction of 25% for numerosity discrimination (Ross, 2003). This means that subjects can reliably discriminate two numerosities without counting, when the difference between the stimuli is more than 25% percent. The transition to counting at 4 items could therefore be explained by the fact that the relative differences between subsequent numerosities becomes smaller than the Weber fraction for more than 4 items. This means that subjects would have to resort to counting to reliably determine the numerosity for more than 4 items. However, Revkin, Piazza, Izard, Cohen, and Dehaene (2008) showed that subitizing occurred when subjects had to judge 1, 2, 3, ..., or 8 dots, whereas it did not when they had to judge 10, 20, 30, ..., or 80 dots. Note that relative differences between subsequent numerosities are the same in both numerosity ranges. This indicates that subitizing is not accurate estimation enabled through the relative differences between subsequent numerosities being larger than the discrimination threshold. Finally, there is an explanation based on Visual Indexing theory (see Pylyshyn (2001) for an overview). This is a theory that follows from multiple object tracking experiments. Humans can visually track up to 5 items simultaneously (Pylyshyn & Storm, 1988). It has been proposed that this is mediated through mental pointers through which we can use to refer to a certain object without having to couple it to certain features or a certain location. Trick and Pylyshyn (1994) argued that this mechanism also enables subitizing.

So far, no consensus has been reached on which explanation for the accurate and efficient judgement of small (< 4) numerosities is true. What complicates dissociation of counting and subitizing is that humans seem to use a combination of the two. In vision, it has been found that numerosity judgement depends on item arrangement (Van Oeffelen & Vos, 1982a, 1984). Humans tend to group items together, judge each group and then add them all together to arrive at the total. A recent brain imaging study failed to find a brain circuit that is exclusively activated during judgment of small numerosities (Piazza, Mechelli, Butterworth, & Price, 2002). It is not clear whether this means that the same process underlies numerosity judgement in and outside the subitizing range, or whether this is due to subitizing being used in combination with counting during judgement of large numerosities.

1.3.2 Haptic numerosity judgement

Although it is not clear what mediates the fast and accurate enumeration for small numbers of items, it has been shown to occur in touch as well (Riggs et al., 2006). Studying numerosity judgment in touch is of course important for understanding how the haptic system extracts numerosity information. However, it can also be valuable for the field of visual numerosity judgment. If enumeration is similar in both modalities, it is likely that it is mediated in a similar way. This would imply, for instance, that the underlying reason that subitizing can occur is not likely an essentially visual one.

As mentioned before, fast and accurate enumeration of small numbers of items has been shown to exist in touch. Riggs et al. (2006) showed that subitizing occurs when subjects have to judge the number of fingers that were stimulated using pins pressed onto the fingers. This may be a special situation, because we know how many fingers we have and where they are located. Judging a number of objects grasped together in the hand is a far more complex task. Before the number of items can be judged through any mechanism, they have to be individuated, i.e. it has to be determined what is one item and what is another item. This is more complicated for separate objects grasped in the hand than when the number of stimulated fingers has to be judged.

In this thesis, a series of experiments is presented in which subjects had to judge the number of three-dimensional shapes that are grasped in the hand. In Chapter 6 of this thesis, it was investigated whether subitizing occurs when subjects have to judge a number of spheres grasped together in the hand (Figure 1.2a). The results show that in this case numerosity judgement was more efficient up to three items than for larger items. We also investigated whether this fast performance relied on relative differences between subsequent numerosities. Furthermore, we investigated whether size or weight estimation played a role and whether it is crucial that numerosity information is present. In Chapter 7, we investigated whether these results are reproducible in vision, by carrying out a visual version of our haptic study (Figure 1.2a).

As pointed out earlier, for visual numerosity judgment humans tend to group items together. They judge each group and then add them all together to arrive at the total. In Chapter 8, we investigated whether such a group–and–add strategy can also be used in haptic numerosity judgement. To that end, we grouped items by presenting a set of spheres to both hands simultaneously (Figure 1.2b). If a group–and–add strategy is used, numerosity judgement should be facilitated in this case.

So far, numerosity judgment of objects that all have the same size and shape was studied. However, in the first part of this Introduction it was pointed out that in daily-life we often hold objects in our hand that do not all have the same size, weight or shape. Such features can be used to recognize a certain object. The question arises of whether they also play a

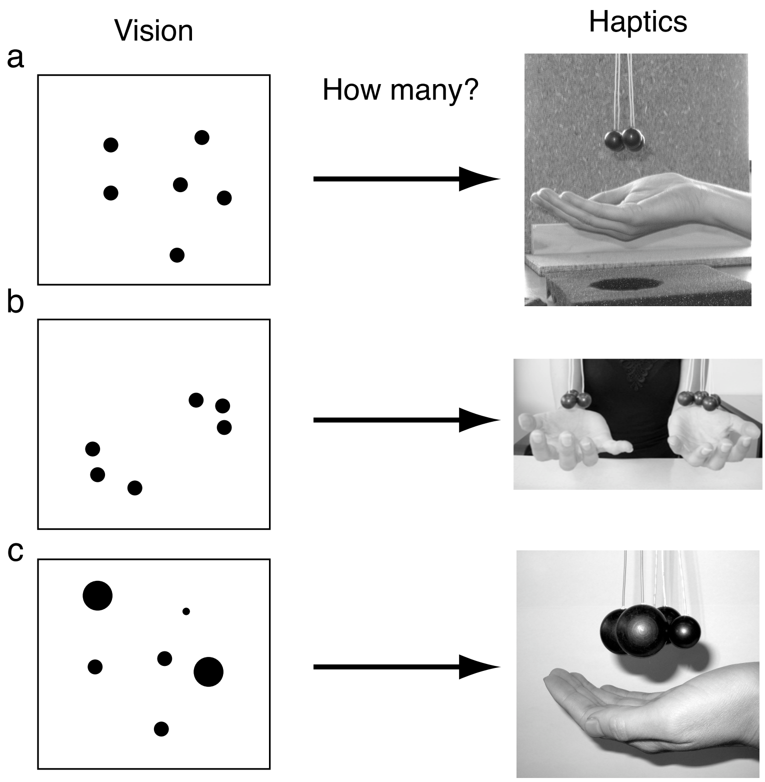

Figure 1.2: Overview of different visual numerosity judgement stimuli and their haptic counterparts. In this thesis only the haptic stimuli were used, except for the visual stimulus in (a) which was used in Chapter 7. In all cases, the subject had to respond the total number of items. a) Subjects grasped a set of spheres using the dominant hand (Chapter 6). b) Subjects grasped a set of spheres with each hand simultaneously (Chapter 8). c) Subjects grasped a set of spheres heterogeneous in size with the dominant hand (Chapter 9).

role in object individuation. To answer this question, subjects were asked to judge the number of items for sets of spheres of different sizes mixed together (Figure 1.2c) or a mixture of spheres and cubes in Chapter 9.

1.4 Summary

Using search tasks it was investigated which object properties are most salient and therefore enable fast object recognition. Furthermore, this paradigm was used to gain insight into haptic search strategies. The numerosity judgement tasks were aimed at investigating whether the fast mechanism for enumeration of small numbers of items in vision, known as subitizing, is also available in active touch. This knowledge was used to draw conclusions about object individuation. Together these studies provide insight into how information about multiple objects is extracted and perceived through active touch.

Chapter 2

Haptic pop-out in a hand sweep

Published as:
Plaisier, M.A., Bergmann Tiest, W.M. & Kappers, A.M.L. Haptic pop-out in a hand sweep. *Acta Psychologica*, 128: 368-377, 2008.

Abstract

Visually, a red item is easily detected among green items, whereas a mirrored S among normal Ss is not. In visual search, the former is known as the pop-out effect. In daily life, people often also conduct haptic *(tactual) searches, for instance, when trying to find keys in their pocket. The aim of the present research was to determine whether there is a haptic version of the pop-out effect. Blindfolded subjects had to search for a target item which differed in roughness from the surrounding distractor items. We report reaction time slopes as low as 20 ms/item. When target and distractor identities were interchanged the slopes increased indicating a search asymmetry. Furthermore, we show that differences in search slope were accompanied by search strategy differences. In some conditions a single hand sweep over the display was sufficient, while in others a more detailed search strategy was used. By relating haptic search slopes to parallel and serial search strategies we show, for the first time, that pop-out effects occur under free manual exploration.*

2.1 Introduction

Every day we reach into our pocket to take out our keys or we try to find a light switch in the dark. These are some common examples of the haptic searches humans conduct. Like in visual search, some haptic searches are much easier than others. Visual search tasks have been researched extensively over the years. Typically, the task is to find a certain target item among a varying number of distractor items. This can yield large differences in response times among tasks. Models of visual search try to explain these differences. However, relatively little is known about haptic search.

Treisman and Gelade (1980) proposed the Feature Integration Theory (FIT). This theory distinguishes between processing of visual information at the 'pre-attentive' stage and at the 'attentive' stage. They suggested that searches for basic features, so-called 'visual primitives' (e.g. colour) can be processed at the pre-attentive level. At the pre-attentive level information is processed in parallel, which means that response times are independent of the number of distractor items and the target item is said to 'pop out'. Searches at the attentive level (e.g. an 'S' among mirrored 'S's) are processed serially and the response time increases linearly with the number of items in the display. However, in practice this division between parallel and serial searches is not as rigid as suggested by this theory. Many conjunction searches, e.g. a red vertical bar among red horizontal and blue vertical bars, are processed more efficiently than the purely serial processing predicted from the Feature Integration Theory. Therefore, another theory of visual search, the 'guided search model', was proposed (Wolfe, Cave, & Franzel, 1989). This model suggests that the efficiency differences between visual search tasks can be explained from variations in the extent to which pre-attentive parallel processes can be used to guide attention in the attentive stage. One way of guidance is 'bottom-up' guidance, where attention is guided to a salient feature. In the case of a conjunction search there is 'top-down' guidance, which means that at the pre-attentive stage all red bars and all vertical bars, for instance, could be located and through feature binding this information could be used to make a single object representation and find the item that is both red and vertical. This could be an explanation for the fact that many conjunctions searches are performed more efficiently than predicted when the search would be performed serially.

In previous research on haptic search tasks, target and distractor items were usually pressed onto the fingers of human test subjects (e.g. Lederman et al., 1988; Lederman & Klatzky, 1997; Purdy, Lederman, & Klatzky, 2004). Exploratory movements are then confined to small finger movements and the number of items that can be presented is limited to the number of fingers. The advantage of presenting haptic items to the fingers, on the other hand, is that all items are presented simultaneously. Since it can be expected that the information processing on a neurological level is similar to that in vision, visual search models may be easily extrapolated to haptic search tasks in which items are presented in this manner. However, these results cannot readily be generalised to haptic search under free exploration conditions. Although in case the items are randomly distributed on a display, the presentation of the items is similar to how this is generally done in vision, the way in which the information is extracted haptically can be considered quite different. In the haptic case subjects will always have to move their hand over the display, which introduces a serial component, but most importantly, they can adjust their exploratory strategy. Hand movements are not performed in the same way as eye movements which consist of saccades and fixations. Movement of the hand and probably the whole arm is relatively slow and this may have a large influence on haptic search times. If and how the haptic exploratory strategy of a display co-varies with difficulty of a search task has never been investigated. Note that while roughness perception is usually investigated in terms of cutaneous perception, under free exploration conditions, when items also vary in spatial location and hand and arm movements are made, proprioception also plays an important role. Hence such a task should be referred to as a haptic search task (a combination of cutaneous and proprioceptive perception) rather than a tactile search task.

In vision, the slope of the relationship between response times and the number of items in the display is used as a measure for the efficiency of a search. These slopes are referred to as search slopes. A serial self-terminating search is usually characterised by a $1:2$ ratio between the search slopes of the target present and target absent trials, while the intercept is the same. For serial search in target present trials, subjects only search on average half of the items before they find the target, while they always search the whole display in target absent trials (hence the ratio $1:2$). This might not be the case for haptic search, because it is difficult to determine whether the whole display was searched and subjects might search part of the display or possibly the whole display repeatedly. This might result in differences in intercept between target present and target absent trials. It also implies that the $1:2$ ratio between the slopes may not be a suitable indication for haptic serial self-terminating search.

When trying to find our keys or switching on the light, we make hand movements and the item we are trying to find can make contact with any part of our hand. The type of hand movements made, has been shown to depend on the type of haptic information that is to be extracted (Lederman & Klatzky, 1987). The natural exploratory movement for perceiving roughness, for instance, is a lateral motion. Perceiving thermal properties of a material, on the other hand, requires the skin to make contact with the material long enough to establish a certain amount of heat transfer. This is a relatively slow process which was also reflected in the results of Lederman and Klatzky (1997). Besides being relatively fast, roughness perception has been the subject of a considerable amount of research (e.g. Bergmann Tiest & Kappers, 2007; Goodwin & Wheat, 2004; Hollins & Risner, 2000; Johnson & Hsiao, 1992; Klatzky & Lederman, 1999; Lederman & Taylor, 1972). Lederman and Klatzky (1997) found that searches for material properties, like a rough target item among smooth distractor items, are relatively easy. In contrast, searches for relative orientation were shown to be more difficult and to depend strongly on the number of items. These results make some material properties, such as roughness, good candidates as 'haptic primitives'. Therefore, we decided to have subjects haptically explore surfaces covered with patches of differing roughnesses as target and distractor items.

We set out to find a haptic version of the pop-out effect under free exploration conditions by exploring search efficiency differences. We did this in terms of response times as a function of the number of items and in terms of exploratory strategy, i.e. movement track over a display. In analogy with visual search tasks, response times were measured while varying the number of items on the surfaces. We asked blindfolded subjects to freely explore the surfaces with their dominant hand. As there was no reason to expect otherwise, we assumed a linear relationship between response time and the number of items. In visual parallel search, all items on the display are perceived simultaneously and search times are independent of set size. In contrast with visual searches, subjects had to move their hand over the surface and therefore not all items were perceived simultaneously. All displays were of the same size and the target item could be placed anywhere on the display. Thus, set size by itself could not influence the response time. A search slope deviating from zero would therefore, like in the visual case, be caused by the influence of the distractor items. Slope differences between different conditions could be caused by differences in the haptic information processing mechanism, but also by the subjects' exploratory movements. Pilot experiments suggested that some haptic search tasks could be performed by a single hand sweep, while others required the subjects to visit each item with their fingers. The first method enables a more parallel intake of information than the second. Since the natural exploratory movement for perceiving roughness is a lateral motion, the most efficient way to explore the presented surfaces would be to sweep the hand over it. If the target item pops-out and distractor items have little or no influence it can be expected that subjects just sweep their hand across the surface once in order to detect a target item.

From visual experiments it is known that interchanging target and distractor identity can also cause differences in search slopes, an effect labelled 'search asymmetry'. These asymmetries can be caused by differences in processing of the items, but also by an asymmetry in the design of the stimulus (e.g. Rosenholtz, 2001). Search asymmetries in touch were already reported by Lederman and Klatzky (1997). In the present research we investigate whether they occur under active exploration and if an asymmetry in response times is accompanied by an asymmetry in exploration strategy.

Two experiments were conducted. Experiment 1 was a 'classic' search experiment in which subjects had to search for a single target item among a varying number of distractor items, while response times were measured as a function of the number of items. Two control experiments were conducted to assess that all items could be detected accurately. In Experiment 2 we partially repeated Experiment 1 while tracking the subjects' hand position on the display. Again a control experiment was conducted, this time to investigate whether the different types of items were detected using different exploratory strategies.

2.2 Experiment 1: Response times

In this experiment subjects actively searched a display with target and distractor items on it to investigate how efficient they can perform such a task. Furthermore, we compared different conditions to assess the effect of different types of target and distractor items. To investigate the effect of decreasing intensity contrast between target and distractor items, we compared a condition where the target item was rough while the distractor items had a finer texture with a condition where the target item was replaced by a somewhat less rough texture. We also included a condition in which the identities of the target and distractor items were interchanged, i.e. a fine textured target item among rougher distractor items. To investigate how well each of the types of sandpaper could be detected, we performed also two control experiments.

2.2.1 Method

Participants

Eight paid undergraduate students (3 females, 5 males; mean age = 20 ± 2 years) participated in each of the experimental conditions. All subjects were right-handed according to Coren's test (Coren, 1993). They gave their informed consent and were treated in accordance with the local guidelines.

Stimuli and apparatus

Figure 3.1 shows a schematic representation of one of the stimuli. The set consisted of 20 × 20 cm^2 displays, made out of Medium Density Fibre (MDF) board with a smooth surface into which 3 cm diameter holes had been drilled. There were 3, 5, 7, 9 or 11 holes and they were distributed randomly over the display at least 2 cm from the edges of the display. The rims of the holes were at least 1 cm apart. Two different displays were made for every number of holes. Plugs with sandpaper on them could be fitted into the holes, such that the surface of the sandpaper was at the same level

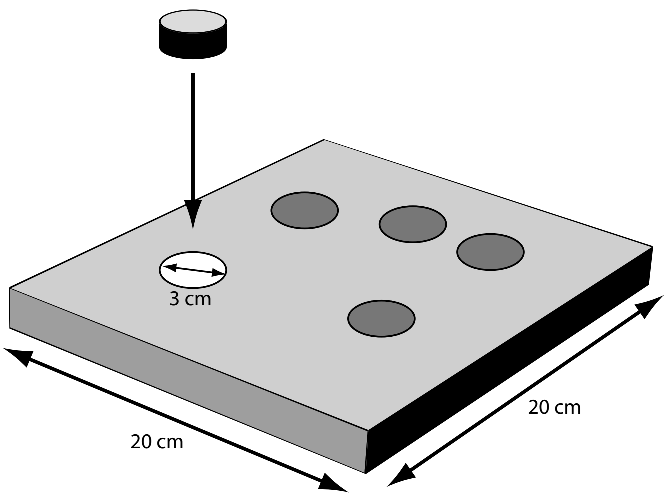

Figure 2.1: Schematic drawing of the display with five items. Plugs could be fitted into the holes to place the items on the display.

as the MDF surface. This allowed for items to be placed on the displays. Three types of sandpaper were used as items: fine (Siawat P360), medium rough (Sianor J P120) and rough (Sianor P60). These type codes indicate a mean particle diameter of 28.8 μm, 116 μm and 269 μm, respectively, according to the Federation of European Producers of Abrasives (FEPA) ‘P’ standard and in this case the particles were silica.

For response time measurements, the stimuli were placed on a computer-interfaced precision scale (Mettler Toledo SPI A6). Measurements were started when a weight change was detected due to a subject touching the stimulus. The scale had a time delay of 70 ms and this was added to the raw data. Measurements terminated with a verbal response registered using a headset microphone. The height of the scale remained stable upon pressure.

Task

The experiment consisted of three conditions in which subjects had to search for a target item among distractor items. Subjects had to say whether the target item was present or absent by calling out the Dutch equivalents of ‘yes’ and ‘no’, respectively. In the first condition, the target item was the rough sandpaper and the distractor items were fine sandpaper (condition 1). In the second condition, the target item was replaced by the medium rough sandpaper (condition 2) and in the third condition, the target item was fine sandpaper and the distractor items were made of medium rough sandpaper (condition 3).

Procedure

The blindfolded subjects were instructed to determine in the shortest possible time whether a target item was present, but it was also emphasised that they had to be correct. Incorrect trials were repeated at the end of the block so the average response time for each number of items was based on the same number of trials. Subjects used their dominant hand to explore the displays. Control experiments with and without earplugs did not reveal any difference in performance; therefore, to increase their comfort, subjects did not wear earplugs. Before a trial started, subjects placed their dominant hand on a hand rest. Since all subjects were right-handed the rest was always located on the right-hand side of the stimulus. The rest was levelled with the height of the stimulus so subjects could easily slide their hand from the rest onto the stimulus.

Each block of trials was preceded by a training session. During training, stimuli were presented until the subject was comfortable with the task and subjects were encouraged to find the fastest strategy. Then, trials were continued until ten in a row were correct before the actual experiment began. Throughout the training and the experiments, subjects received feedback as to whether their answers were correct.

The experiments were divided into blocks of approximately 45 minutes each and subjects performed no more than one block on the same day. The three conditions consisted of 150 trials each (30 trials for each number of items and in half of these trials the target was present). Each condition was divided into two blocks. The order in which the different conditions were performed was counterbalanced over the subjects. For each number of items there was one display with a target item and one without. The display that had a target item on it was interchanged between the two blocks and the position of the target item was randomised. After each trial, the display was rotated 90° to maximise the number of different displays available. Recorded response times that differed by more than three times the standard deviation from the mean were excluded from the raw data.

2.2.2 Results

For each number of items the response times averaged over all subjects are shown in Figure 2.2a for the target present trials and for the target absent trials in Figure 2.2b. The lines represent linear regression to the data. The values of the slopes and intercepts of the regression lines are indicated by s and y_0, respectively. Error rates did not exceed 5% in any of the conditions. Note that for all conditions the target absent trials yielded larger slopes and intercepts than the target present trials. The search slopes varied between the different conditions. For condition 1 it was rather shallow, while the search slope for condition 2 was somewhat steeper and for condition 3 the slope was quite steep.

For every subject in each condition, linear regression to the data from the target present and the target absent trials provided slopes and intercepts. Two separate 3 (condition)$\times$ 2 (target presence) repeated measures ANOVAs (Analysis Of Variance) with planned comparisons were performed on the slopes and the intercepts. For the slopes this showed significant main effects for condition ($F(2,14) = 28.40$, $p < 0.0005$) and target presence ($F(1,7) = 6.92$, $p = 0.034$). Also the interaction term was significant ($F(2,14) = 4.34$, $p = 0.033$). The main effects for the intercepts were also significant (condition $F(1.11, 7.77) = 16.61$, $p = 0.003$, target presence $F(1, 7) = 31.68$, $p = 0.001$, interaction term $F(1.03, 7.21) = 15.09$, $p = 0.006$). The effect of target presence was analysed further with paired samples t-tests. For each of the separate conditions the effect of target presence on the slopes was significant ($t(7) \leq -2.5$, $p \leq 0.040$) except in condition 2. For the intercepts the difference between target present and absent trials

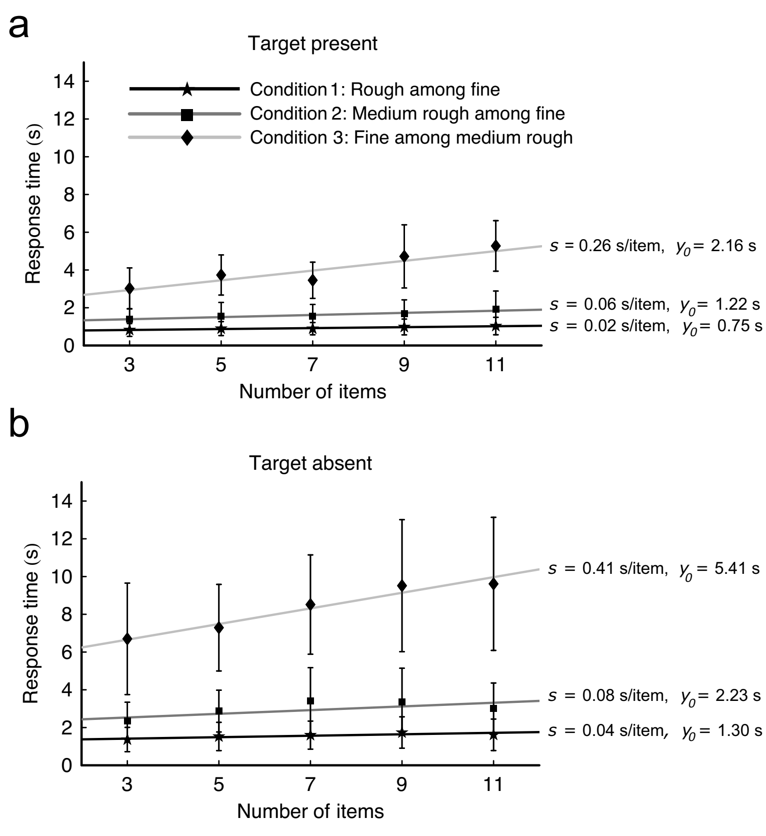

Figure 2.2: Experiment 1: subjects had to search for a target item among varying numbers of target items. The different conditions are indicated in the figure. Response times for each condition are shown, averaged over all subjects (N = 8) as a function of the number of items. a) represents target present trials and b) the target absent trials. The error bars represent the standard error of the mean and the lines represent linear regression to the data, where s represents the slope and y_0 the intercept.

was significant in all conditions ($t(7) \leq -4.2,\ p \leq 0.004$).

Planned comparisons between conditions 1 and 2 showed significant differences for the slopes ($F(1{,}7) = 6.36$, $p = 0.002$) as well as the intercepts ($F(1,\ 7) = 15.71$, $p = 0.005$). So decreasing the contrast between the target and distractor items increased both the intercept and the slope. The contrast between conditions 2 and 3 was also significant for both the slopes and the intercepts ($F(1,\ 7) = 25.23$, $p =0.002$ and $F(1,\ 7) = 12.71$, $p = 0.009$). This means that interchanging target and distractor identities caused an increase in both the slopes and intercepts.

2.2.3 Discussion

Compared to the other conditions, the search slope for condition 1 is rather shallow. Also, the intercept of less than a second is surprisingly low considering the fact that mechanical action is involved. If all items had to be found one by one to decide whether it was a target item, we would have expected much higher search times and slopes. Furthermore, in all conditions the intercept is larger for the target absent trials than for the target present trials. The intercept would be expected to be the same if the only difference between target present and absent trials is that subjects search on average only half of the display. The increase in intercept could be explained by subjects searching part of the display more than once in the target absent trials because they are uncertain of whether they did search the whole display.

The significant difference in slope between conditions 2 and 3 indicates a search asymmetry. A search asymmetry for rough and smooth items was also reported by Lederman and Klatzky (1997). They suggested that the search asymmetry is caused by the ends of a given continuum not being equally perceptually accessible. A rough patch would therefore be processed earlier than a fine patch. An alternative explanation could be that attention is guided by rough items more strongly than by less rough items, which would relate to Wolfe's guided search model (Wolfe et al., 1989). To investigate the origins of the differences in search slopes between the conditions we investigated detectability in Control experiments 1.1 and 1.2.

2.2.4 Control Experiment 1.1

In vision, search asymmetries are often caused by an asymmetrical design of the experiment (Rosenholtz, 2001; Rosenholtz, Nagy, & Bell, 2004). The search asymmetry reported in the previous experiment could thus be caused by an asymmetry in our experimental design. This might be due to detectability differences between the types of sandpaper. To investigate how accurate and how fast the three types of sandpaper were perceived a detection experiment was performed.

Method

The same subjects that participated in Experiment 1 also participated in this experiment. The set-up, procedure and also the stimulus design were the same, only this time there were just four displays: one blank display and three displays with a single item. The item could be any of the three types of sandpaper and subjects only had to say as fast as possible whether or not there was an item present. Each subject performed 60 trials; 15 trials for each type of sandpaper and for the blank display.

Results

Figure 2.3 shows the response times averaged over all subjects for the four conditions in the detection experiment. response times were below one second for displays with an item on them. The large standard error for the no item case was due to one subject having a much longer response time in this condition than the other subjects. For displays with the medium rough or rough sandpaper, no incorrect answers were given, while for the

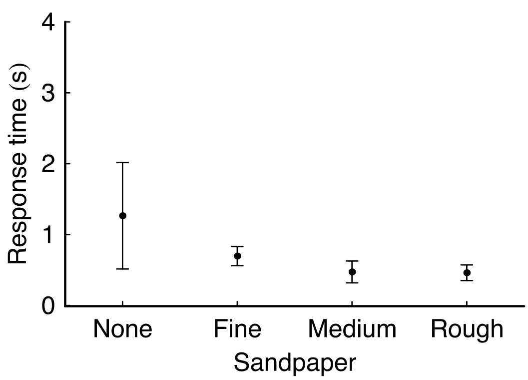

Figure 2.3: Control Experiment 1.1: subjects had to say whether an item was present on the display. The graph shows response times averaged over all subjects (N = 8) for the three types of sandpaper and an empty display. The error bars indicate the standard error of the mean.

displays without sandpaper and with the fine sandpaper error rates were 0.83% and 3.3%, respectively. This indicates that all types of sandpaper were detected accurately. response times for the rough and medium rough sandpaper were shorter than for the fine sandpaper and the no item case. A repeated measures ANOVA showed a significant main effect for the type of sandpaper, $F(1.057,\ 7.402) = 8.68,\ \ p = 0.019$. Pairwise comparisons with Bonferroni correction yielded significant differences in response time for the fine and the medium rough sandpaper ($p = 0.045$), as well as for the fine and the rough sandpaper ($p = 0.008$). These results show that all types of sandpaper were detected relatively fast, but the rough and medium rough sandpaper were detected significantly faster than the fine sandpaper.

Discussion

These findings indicate that the rough sandpaper had a higher contrast with the smooth background of the display than the fine sandpaper. This could be the reason for the slope difference between the search for a fine item among medium rough distractors and a medium rough item among fine distractors. However, it could also be that the differences between the different numbers of items were not perceived when the distractors consisted of the fine sandpaper. This would mean that the distractor items did not distract and therefore yielded the relatively shallow lines in the rough or medium rough target item among fine distractor items conditions. To be certain that differences between different numbers of items were perceived, Control experiment 2 was conducted.

2.2.5 Control Experiment 1.2

We conducted an experiment in which the subjects had to judge the number of items in the display to confirm that the differences between the varying numbers of distractor items in Experiment 1 were perceived. If subjects could judge the different numbers of items accurately then the differences between the varying numbers of items were perceived and their use as distractor items was justified.

Method

Again the subjects from Experiments 1 and 2 participated and the set-up and procedure of Experiment 1 were used. We took a subset of the displays from Experiments 1 and 2. Subjects were presented with displays having 0, 1, 3 or 5 items on them and they had to respond how many items were on the display. This experiment was done with both the fine and the medium rough sandpaper, which were used as distractor items in Experiment 1. Each subject performed 60 trials per type of sandpaper.

Results

Figure 2.4a shows the response times averaged over all subjects, as a function of the number of items for the fine sandpaper and Figure 2.4b for the medium rough sandpaper. It can be seen that response times did not vary systematically with the number of items and response times are in the same range for both types of sandpaper. A 4 (number of items) × 2 (type of sandpaper) repeated measures ANOVA on the response times did not show a significant main effect. On average the subjects were able to maintain an accuracy above 95% correct answers for both types of sandpaper. For the fine sandpaper the incorrect answers per number of items ranged from 0.83% to 10% and for the medium rough sandpaper from 0% to 7%. In both conditions the error rate increased with the number of items.

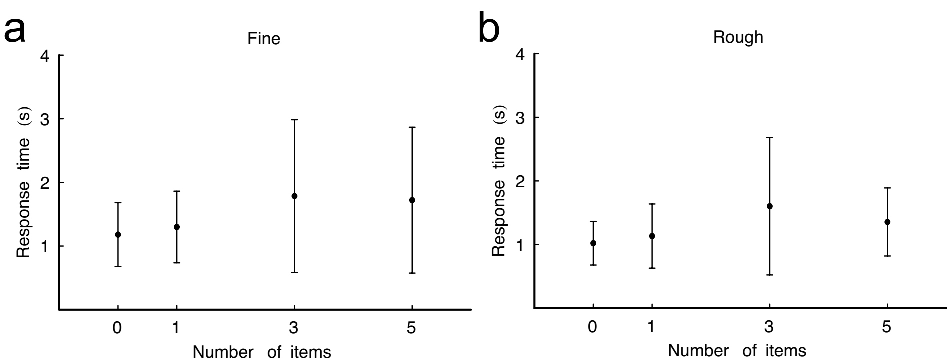

Figure 2.4: Control Experiment 1.2: subjects had to judge the number of items in the display. The graphs show the response times, averaged over all subjects (N = 8), as a function of the number of items. The error bars indicate the standard error of the mean. Items on the display were fine sandpaper (a) and medium rough sandpaper (b).

Discussion

The low error rates indicate that differences between the varying numbers of items could be perceived accurately. Numerosity judgements for the medium rough sandpaper, however, were not significantly faster than for the fine sandpaper. These results show that subjects could accurately estimate the number of items on a display of up to five items.

2.2.6 Conclusions

The search slopes show that when the target item was rough among fine distractor items, search slopes were relatively low. Since haptic exploration involves hand and arm movements, much higher slopes would be expected if the display would be scanned serially. When the roughness difference between target and distractor items was reduced both the slope and the intercept increased. The slope increase indicates that the influence of the distractors was larger when the roughness difference was smaller. The increase in intercept indicates that exploration speed decreased independently of the number of items on the display. Furthermore, there was an increase in slope and intercept comparing search for a rougher target item among finer distractor items with search for a finer target item among rougher distractor items, indicating a search asymmetry. The control experiments show that all types of sandpaper could be detected accurately (control 1.1) and also that the differences between different numbers of items could be detected (control 1.2). This means that all types of sandpaper used in this experiment could indeed act as target and distractor items. The differences in search slopes between the conditions were therefore caused by the differences in target and distractor identity and not merely detectability differences.

2.3 Experiment 2: Exploratory strategy

The differences in intercept and slope values between the conditions in Experiment 1 could be caused by subjects simply moving slower over the surfaces, or by a shift in search strategy. To investigate whether a strategy shift occurred we repeated Experiment 1 in part while tracking the subjects' hand movements. Also a control experiment was performed to investigate whether the different types of sandpaper were detected through different exploratory strategies.

2.3.1 Method

Eight new paid subjects (6 females, 2 males; mean age = 22 ± 2 years) participated in this experiment and all of them also performed the control experiment. All subjects were right-handed according to Coren's test (Coren, 1993) and gave their informed consent. None of them had any known hand deficits. The response time measuring set-up and stimuli from Experiment 1 were adopted. The subject's hand position was recorded using a movement tracking system (NDI Optotrak Certus). A marker consisting of an infra-red LED was placed on the nail of the index finger of the subjects' dominant hand and the marker position was recorded at a rate of 100 Hz. In Experiment 1 it was observed that subjects always moved over the surface with a flat hand and they did not spread their fingers and just moved their whole hand. They only rotated the hand with respect to the wrist when moving over the displays. Therefore, one marker was sufficient to detect strategy differences. As all subjects were right-handed they all entered the display from the right-hand side where they placed their hand on a rest before the trial started as they did in Experiment 1. This means they always entered the display from the same side, but the displays were rotated to randomise item positions as this was also done in Experiment 1.

The instructions were identical to those in Experiment 1. Each subject performed 2 target absent and 2 target present trials for each number of items in each of the three conditions from Experiment 1, totalling 60 trials. The order of the different conditions was counterbalanced over subjects.

2.3.2 Results

The response times in these experiments were similar to those found in Experiment 1 and therefore, these results can be extrapolated to what we found in Experiment 1. A representative selection of the tracks over the stimuli from one subject is shown in Figure 2.5. The squares represent the display and the solid line marks the track of the subjects index finger over the display. The subject entered the display from the right hand side. It can be seen that in all experimental conditions the track tended to be longer in the target absent trials. Note that as the position marker was on the index finger sometimes tracks will extend across the display edges, but the subjects hand would then still be on the surface. Furthermore, between the different conditions the length of the track and the scale of the movements varied. In condition 1 the target present trials generally show only one sweep over the surface, whereas the tracks over the displays in condition 3

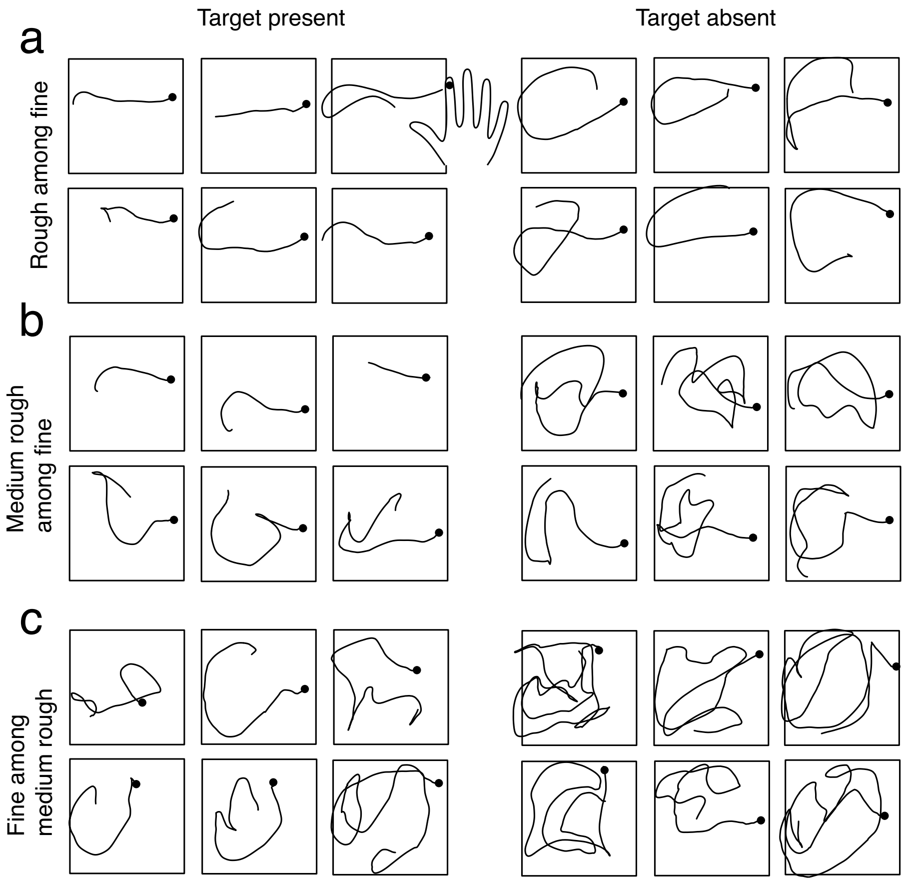

Figure 2.5: Experiment 2: subjects had to search for a target item among varying numbers of distractor items while the movement over the surface was recorded. A selection of tracks across the displays from a single subject is shown. The squares indicate the edges of the displays and the hand indicates the subjects' hand entering the display from the right side. The dot on the index finger depicts the position marker. The starting point of each track is marked with a dot. For each display the total number of items is indicated in the lower right corner. a) Rough target item among fine distractor items (condition 1). b) Medium rough target item among fine distractor items (condition 2) and c) Fine target items among medium rough distractor items (condition 3).

show a far more complicated movement profile.

In Figure 2.6 a selection of tracks from one subject in conditions 1 and 3 is given, now with the position of the items indicated. A grey filled disk marks the position of the target item. Note that in condition 1 the subjects did not necessarily have to move their fingers over the target item, they also used other parts of the hand. In condition 3 the movements concentrated on the areas with items present, while this is not apparent in condition 1. Furthermore, the length of the track of the target present trials varied markedly in condition 3, because it was highly dependent of the location of the target. If a subject happened to start searching near the target item the track was much shorter than when it was further away.

For a more quantitative analysis the length of the tracks and movement speed were analysed. First the length of the track was calculated from the position data. The track lengths were averaged over all numbers of items tested (3, 5, 7, 9, 12) for the target present and absent trials in each of the three conditions. Figure 2.7a shows the distance travelled across the display averaged over subjects for the target present and absent trials in the three conditions. In the target present trials from condition 1 the length of the track was approximately 20 cm, which equals the width of the displays

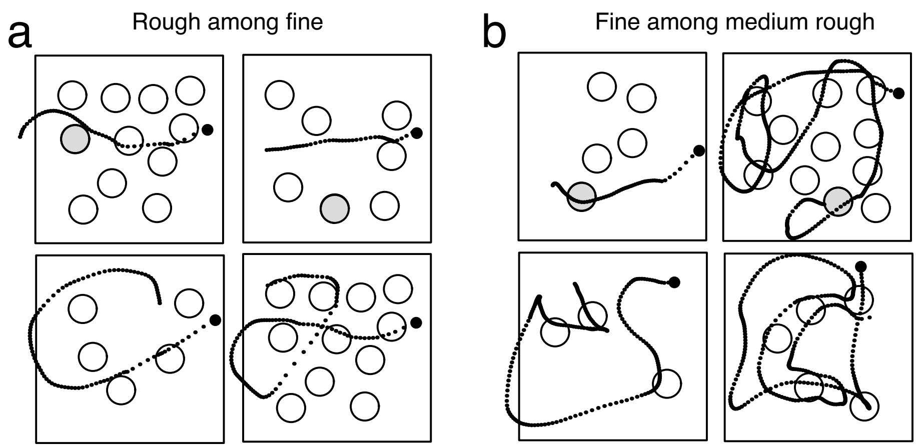

Figure 2.6: Experiment 2: Tracks across the displays from the same subject as in Fig 5 are shown with the position of the items indicated. A filled circle indicates a target item. The dots indicate the subsequent positions of the marker which was sampled at a rate of 100 Hz. A larger black dot indicates the starting position of a trial. a) Condition 1 and b) condition 3.

and suggests a single hand sweep was performed. Also the average speed at which subjects moved over the displays in the different conditions was calculated. The averaged speed is represented in Figure 2.7b. From this figure it can be seen that in each condition the average speed in the target

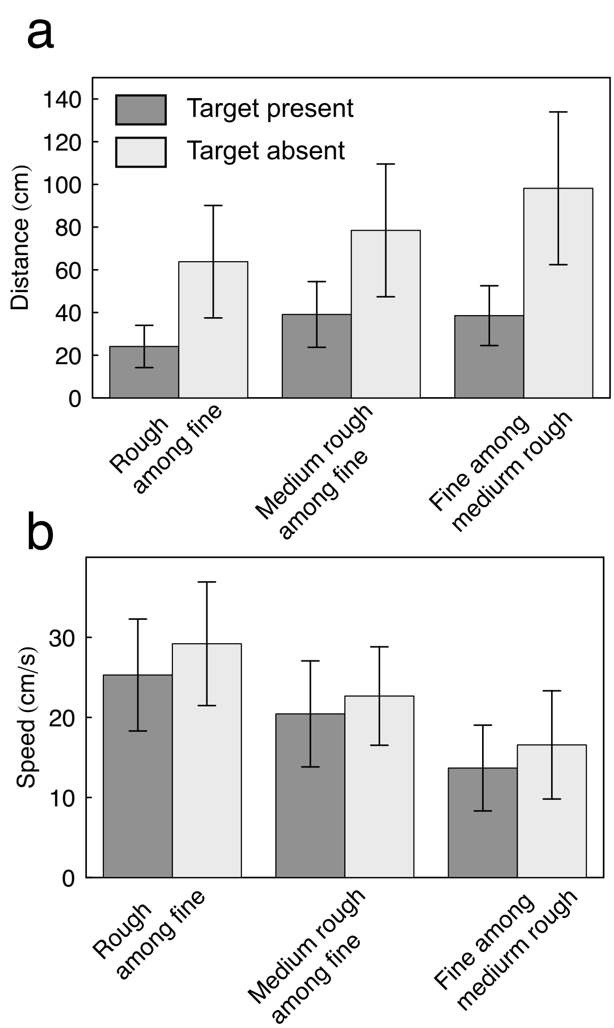

Figure 2.7: Experiment 2. a) The distance that subjects moved over the display and b) the speed at which they did this averaged over all subjects (N = 8) for the three conditions. The dark bars indicate target present trials and the light grey bars the target absent trials. Error bars indicate the standard error of the mean.

present trials was slightly smaller than in the target absent trials.

A 3 (condition)× 2 (target presence) repeated measures ANOVA with planned comparisons on the track length showed significant main effects for condition and target presence ($F(2,14) = 18.2$, $p < 0.0005$ and $F(1,14) = 21.3$, $p < 0.002$). Planned comparisons showed that the difference between condition 1 and 2 was significant ($F(1,7) = 18.2$, $p = 0.004$), as well as the difference between condition 2 and 3 ($F(1,7) = 5.9$, $p = 0.045$). In each of the conditions the averaged total track over the display was significantly longer in the target absent trials than in the target present trials (paired samples t-test, $t(7) \leq -3.2, p < 0.0151$).

In Figure 2.7b the average speed across the displays is shown. The average speed was highest in condition 1 and lowest in condition 3. Repeated measures ANOVA showed significant main effects for condition and target presence ($F(2,14) = 33$, $p < 0.0005$ and $F(1,14) = 26.8$, $p = 0.001$). For each of the conditions the difference between target present and absent trials was significant (paired samples t-test, $t(7) \leq -2.4, p < 0.049$). Planned comparisons revealed significant differences between condition 1 and 2 ($F(1,7) = 14$, $p = 0.007$) and between conditions 2 and 3 ($F(1,7) = 18.2$, $p = 0.004$). This shows that interchanging target and distractor identity caused subjects to switch to exploration movements at a lower average speed, but also to make longer exploratory tracks over the display surfaces.

2.3.3 Discussion

These results show that subjects performed a hand sweep across the displays in condition 1, while in condition 3 they switched to searching the displays at a smaller scale and at a lower speed. The search strategy in condition 2 was an intermediate between a hand sweep and the strategy in condition 3. This suggests that in condition 1 the search had a parallel character in which the target could be found through a hand sweep. In condition 3 the search had a far more serial character in which items were examined sequentially. Summarising, subjects adjusted their search strategy to a more parallel or a more serial model depending on the contrast between target, distractor and background.

2.3.4 Control Experiment 2.1

In Control Experiment 1.1 it was already found that all types of sandpaper were detected accurately. In the present control experiment we investigated whether there is a strategy difference between detecting different types of sandpaper. If it is found that all types of sandpaper are detected through the same exploratory strategy, we can conclude that the strategy differences found between the conditions in the search task must have been caused by the presence of distractor items. To investigate this we repeated Control Experiment 1.1 while tracking the subjects' hand position.

Method

The stimuli and instructions were the same as in Control Experiment 1.1. The subjects performed 4 trials for each type of sandpaper and the empty

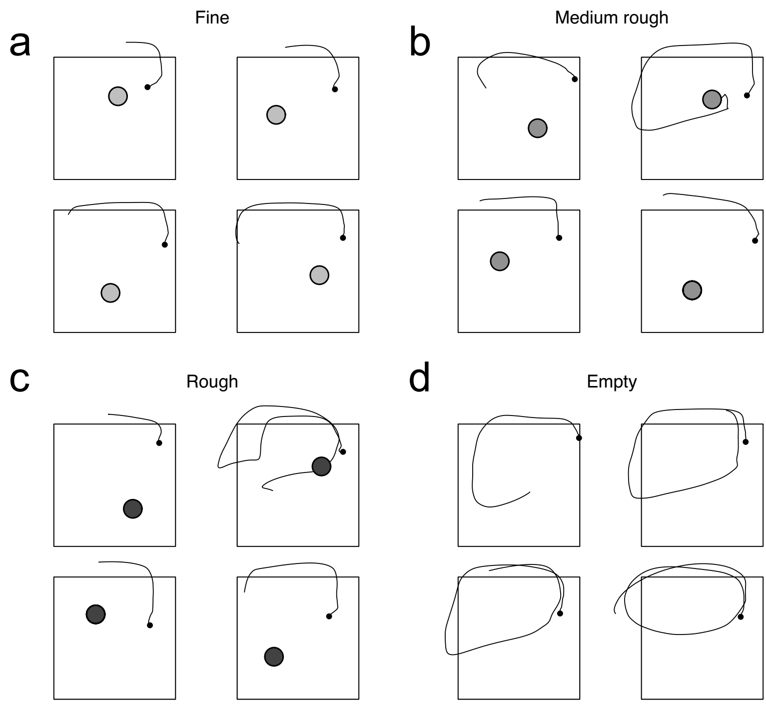

Figure 2.8: Control Experiment 2.1: subjects had to say whether an item was present on the display while the movement over the display was recorded. A selection of tracked movements over the displays from a single subject is shown. The grey disks indicate the item position and the items could fine sandpaper (light grey disk), medium rough sandpaper (intermediate grey disk) or the rough sandpaper (dark grey disk). The start of a track is marked with a black dot. The subjects responded whether there was an item present for fine sandpaper (a), medium rough sandpaper (b), rough sandpaper (c) and an empty display (d).

display. Since the displays were rotated 90° for each trial, the location of the sandpaper was roughly homogeneously distributed over the four quadrants of the display.

Results

A selection of tracks for the different types of sandpaper and the empty display is shown in Figure 2.8. It can be seen that the empty display was searched more extensively than the other displays. All displays with sandpaper were searched with one sweep over the surface and subjects did

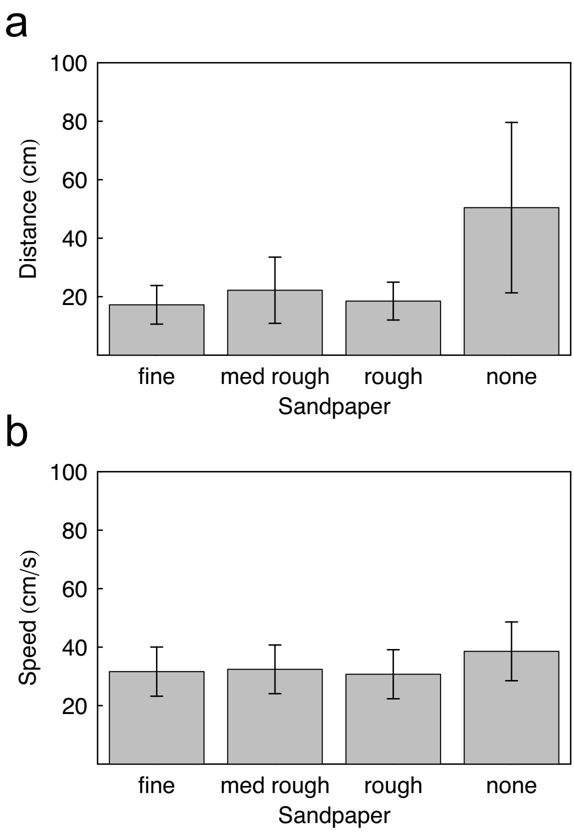

Figure 2.9: Control Experiment 2.1. a) The distance that subjects moved over the display and b) the speed at which they did this averaged over all subjects (N = 8) for each of the types of sandpaper and the empty display. Error bars indicate the standard error of the mean.

not have to move their fingers over the item but could use any part of the hand to detect it. The average distance travelled over the displays and average speed are shown in Figure 2.9 and it can be seen that for all sandpaper present trials the length of the track roughly equals the width of the display. Repeated measures ANOVA on the distance data showed a significant main effect of the four possible displays ($F(1.1, 7.9) = 9.4$, $p = 0.014$), but pairwise comparisons did not show any significant differences between the different displays. A planned comparison of the length of the tracks over the sandpaper present displays against track length over the empty display showed that the track length was significantly longer when no sandpaper was present ($F(1, 7) = 10$, $p = 0.016$). Analysis of the speed data did not show a significant main effect.

Discussion

These results show that all types of sandpaper were detected using a similar exploration strategy and that subjects could use any part of the hand. There were no significant differences in average speed of track length between the types of sandpaper.

2.3.5 Conclusions

The results show that different exploration strategies were used in the three conditions. They ranged from a parallel strategy (hand sweep) to a more elaborate serial strategy. There were no strategy differences in detection of the types of sandpaper and therefore we can conclude the strategy differences between the conditions were caused by the identity of target and distractor items.

2.4 General Discussion

The results from Experiment 1 show that there were large differences in search slopes between the three conditions. The difference between conditions 2 and 3 indicated a search asymmetry. In Experiment 2 it was shown that different exploratory strategies were used in the different conditions. When searching for a rough target item among fine distractor items (condition 1) subjects generally performed a hand sweep over the surface, while search for a fine target item among medium rough distractor items (condition 3) was performed through a more complex track of movements over the surface. Not only was the exploratory trajectory over the display longer in condition 3, also the speed was lower. Search for a medium rough target item among distractor items (condition 2) was performed through a strategy in between that of conditions 1 and 3. The control experiments showed that both the differences in search slopes and exploratory strategies were not caused by detectability differences. Therefore, these were truly effects of the target and distractor identities.

The difference between conditions 2 and 3 in search slopes (Experiment 1) accompanied by differences in search strategies (Experiment 2) indicate a search asymmetry. Finding a patch of rough sandpaper among fine sandpaper was easier than the reversed case. If the hand is moved along a textured surface there is cutaneous texture information, but there is also a frictional force. Note that the frictional forces are directly related to the roughness of the items. When moving the hand over a rough patch on a surface there will be local stretch of the skin because of higher friction and this friction is also likely to exert strain on the wrist. These cues can be used to efficiently determine whether a rough item is present among less rough items by just sweeping the hand over the display. In the reversed situation, on the other hand, subjects are searching for an item that is less rough and in that case the target item will not exert higher friction on the skin and joints than the distractor items. This could be an explanation why subjects had to switch to a more serial search strategy in this case. Lederman and Klatzky (1997) found the same asymmetry in their experiments. Although in their setup items were pressed to the subjects fingers they could make finger movements and it could be that also in their case subjects found it easier to detect whether their was higher friction on one finger than lower friction on one of the fingers.

For visual search, Wolfe (1998) showed that there is no clear-cut distinction between parallel searches and serial searches based on response times alone. This is probably also the case for haptic searches. However, differences in the extent to which response times depend on the number of items between haptic search tasks do show that information processing in some tasks is more efficient than in others. This could be due to internal processing differences, but in this study using free exploration conditions, subjects also showed differences in their exploratory strategy. Our results show a shift from a very coarse and efficient search strategy (a hand sweep) to more detailed exploration movements on a smaller scale over three different search conditions. This suggests a gradual change from a search strategy with a 'parallel' character to a more 'serial' strategy. In vision it has been shown that eye movements can provide information on whether search is parallel or serial through the number of fixations and saccades (e.g. Zelinsky & Sheinberg, 1997). However, haptic exploratory movements are not readily comparable with eye movements. Saccades can be planned using information from peripheral vision, but in haptics an item can only be detected upon contact with it. This could be an explanation for the differences in search time between target present and target absent trials and the longer distance that subjects moved over the display. Subjects were very unsure whether they have truly searched the whole display. This also indicates that the criterion of a 1 : 2 ratio between search slopes in target present and absent trials is not appropriate to distinguish between serial and parallel search in this type of search tasks. Experiment 2 showed that when subjects performed a hand sweep they swept on average over the whole width of the display in target present trials, not just half of it. In target absent trials they swept the display more than once to be sure there was no target. The ratios between target absent and target present search slopes found in Experiment 1 for the rough among fine, medium rough among fine and fine among medium rough sandpaper were 0.5, 0.75 and 0.6, respectively. So, only for the rough among fine sandpaper condition, a ratio of 1 : 2 for the target present and target absent slopes was found. Experiment 2 shows that this was not because all items were visited sequentially, since the hand movement data clearly shows a parallel search strategy for this condition. Therefore, a ratio of 1 : 2 between target present and target absent trials does not correlate with a serial search strategy in a search task under free exploration conditions.

Our results show that there are haptic search tasks that can be performed markedly fast and efficient while others are more time consuming. We also showed that changes in search slopes between the different condition were accompanied by search strategy differences between the conditions. In this way we have shown for, the first time, a direct connection between search slopes and type of exploration strategy in haptic search. When search slopes were relatively shallow the search was performed through a strategy with a parallel character, while searches yielding a relatively steep search slope were performed through a more serial strategy. This is an important result, because it is difficult to directly relate haptic search slopes to visual search slopes or to haptic searches that were not performed under free exploration conditions. In visual search tasks, pop-out effect usually means that there is little influence of the distractor items. As was already pointed out in the Introduction, a single hand sweep is the most efficient strategy possible to haptically explore a surface with rough items on it. If the target item can be detected through such a strategy this means that the distractor items have little or no influence. Our results show that when the target item was rough sandpaper and the distractor items were fine sandpaper, subjects used a single hand sweep to search the displays. Therefore, we propose that this condition can be interpreted as a haptic version of the pop-out effect.

Chapter 3

Visually guided haptic search

Accepted for publication as:
Plaisier, M.A., Kappers, A.M.L., Bergmann Tiest, W.M. & Ernst, M.O. Visually guided haptic search. *IEEE Transactions on Haptics*, in press.

Abstract

In this study we investigate the influence of visual feedback on haptic exploration. A haptic search task was designed in which subjects had to haptically explore a virtual display using a force-feedback device and to determine whether a target was present among distractor items. Although the target was recognizable only haptically, visual feedback of finger position or possible target positions could be given. Our results show that subjects could use visual feedback on possible target positions even in the absence of feedback on finger position. When there was no feedback on possible target locations, subjects scanned the whole display systematically. When feedback on finger position was present, subjects could make well-directed movements back to areas of interest. This was not the case without feedback on finger position, indicating that showing finger position helps to form a spatial representation of the display. In addition, we show that response time models of visual serial search do not generally apply for haptic serial search. Consequently, in tele-operation systems, for instance, it is helpful to show the position of the probe even if visual information on the scene is poor.

3.1 Introduction

Tele-operation systems and minimally invasive surgery techniques often involve a combination of haptic (force-feedback) and visual information (e.g. camera images, ultrasound). There are several factors that influence image-guided operations. For instance, to facilitate integration of information from the image and the workspace, the image is often superimposed on the workspace. Integration is even facilitated further if the image is projected in depth (Wu, Klatzky, Shelton, & Stetten, 2005). It has also been shown that performance in image guided surgery is influenced by the shape of the image aperture and that performance is better if the surgeon controls the camera position manually (DeLucia, Mather, Griswold, & Mitra, 2006). It is clear that research into how haptic and visual information are combined is important for optimizing performance through such systems.

In the present study, we aim to provide more insight into how several types of visual feedback influence haptic exploration. One example is visual feedback of finger position. Finger position is important for keeping track of which parts of a scene have already been explored. Of course, finger position can also be perceived through proprioception. Although virtual haptic environments created with force-feedback devices like the PHANToM (SensAble Technologies) often allow haptic exploration through only a single contact point with the virtual environment, this does not necessarily prevent the user from forming a spatial representation of the virtual environment. It has been shown that a spatial representation can be established even through short kinesthetic contact (Klatzky & Lederman, 2003). Although spatial representations can be formed through proprioception, spatial representation through vision is usually better (Cashdan, 1968; Worchel, 1951). It has been shown that humans integrate information from the visual and haptic modality in a statistically optimal fashion (Ernst & Banks, 2002). This means that the modality with the highest accuracy is weighed most heavily in the combined percept. Therefore, we expect that vision will play a dominant role in combined haptic and visual spatial representation. It has also been shown that there is transfer of spatial context from visual to haptic search (Nabeta, Ono, & Kawahara, 2003). In that study subjects first performed a block of visual search trials and later a block of haptic search trials. Some displays in the haptic condition were the same as in the visual conditions, while others were new. Subjects were significantly faster when the display had already been shown in the visual condition. Because of these interactions between visual and haptic perception and the fact that a visual spatial representation is usually better than a haptic one, we expect that providing visual feedback of finger position will make haptic spatial exploration more efficient. In this study, we investigate the influence of visual information on haptic exploration and compare haptic search to visual search.

An important difference between visual and haptic exploration of a scene is that in vision a spatial representation of the scene is readily available, which can be used to plan, for instance, saccades directly to areas of interest. In haptic exploration, this spatial representation is not readily available. Adding this type of spatial information through visual feedback could therefore facilitate haptic exploration. To study the influence of visual feedback on haptic exploration, a haptic search task was designed to which different types of visual feedback could be added. The haptic search paradigm has been extrapolated from the visual search paradigm. In visual search, subjects typically search for a certain target item (e.g. a red dot) among varying numbers of distractor items (e.g. green dots) presented on a screen. Usually, response times are measured as a function of the number of items on the display, but eye movements can also be recorded. In daily life, the visual modality is not the only modality that is used to perform search tasks. When we try to take our keys out from our pocket or a pen out of our bag, we search using touch. Contrary to visual search, only a few studies have addressed haptic search in the past. Recently, however, the haptic search paradigm has been gaining attention (Lederman & Klatzky, 1997; Overvliet, Mayer, et al., 2008; Overvliet, Smeets, & Brenner, 2008; Plaisier, Bergmann Tiest, & Kappers, 2008a, 2009b; Plaisier, Kuling, Bergmann Tiest, & Kappers, 2009).

Although items are normally presented on a screen in visual search studies, there are several different ways to present items for haptic search. One way is by pressing the items onto separate fingers. Items can consist of different types of materials or raised lines (Lederman & Klatzky, 1997; Lederman et al., 1988; Overvliet, Mayer, et al., 2008). Items can also be three-dimensional shapes fixed in a grid and subjects have to explore the different shapes sequentially (Overvliet, Smeets, & Brenner, 2008). Another way of presenting three-dimensional shapes that does not force subjects to explore the items sequentially, is to let subjects grasp a number of shapes simultaneously in the hand (Plaisier, Bergmann Tiest, & Kappers, 2009b; Plaisier, Kuling, et al., 2009). Finally, the way of item presentation most similar to the way this is done in vision, is to present items on a surface (Plaisier et al., 2008a). Items can consist of, for instance, rough patches on a smooth surface. Such a 'tactile display' can be actively explored.

The advantage of using stimuli that are actively explored is that subjects can adjust their exploration strategy in order to optimize their performance. It has been shown that there are typical exploratory procedures (EPs) for extracting object properties (Lederman & Klatzky, 1987) and that object recognition can be impaired by constraining the exploratory movements (Lederman & Klatzky, 2004). Analysis of exploratory movements has shown that haptic object recognition is viewpoint dependent (Newell, Ernst, Tjan, & Bülthoff, 2001; Ernst, Lange, & Newell, 2007). Thus, characterization of the exploratory movements that subjects make in combination with response times provides insight into the search strategy used. In two previous studies we have shown the importance of analyzing exploratory strategy for interpreting response times in haptic search tasks (Plaisier et al., 2008a; Plaisier, Bergmann Tiest, & Kappers, 2009b).

In visual search studies, usually only response times are analyzed to determine which search strategy was used. When the response times do not increase with the number of items in the display, the search strategy is referred to as 'parallel' meaning that all items were processed simultaneously. When items are processed one by one, response times increase with the number of items in the display, and the search strategy is referred to as 'serial' (Treisman & Gelade, 1980; Cave & Wolfe, 1990; Wolfe, 1993; Duncan & Humphreys, 1989; Theeuwes, 1993). We have shown in a previous study that the response time slopes can be very shallow in a haptic search task, while analysis of the exploratory movements that were made clearly indicate that the search strategy was serial (Plaisier et al., 2008a). This suggests that visual search models cannot readily be used to distinguish haptic parallel and serial search based on response times alone. As mentioned before, in vision a spatial representation of the scene is readily available whereas this is not the case in haptics. Adding this type of spatial information could make haptic serial search performance more similar to visual serial search performance.

To investigate how visual information can be used to guide haptic search and which types of visual information are most important for enhancing haptic search efficiency, a haptic display was generated using a force-feedback device. On this display, items were defined by regions with a higher friction coefficient than the background of the display. Frictional forces were chosen to define the virtual display, because friction is a property present in the real world that is perceived through lateral motion when you move your finger over a certain material (Lederman & Klatzky, 1987). Subjects haptically explored the display with one finger only, ensuring that the task could only be performed in a serial manner. In the different conditions, varying amounts of visual information could be provided. The effects of the different types of visual feedback were compared to simulations of two extreme types of search strategies. For the first strategy, it was assumed that subjects moved from item to item along the shortest pathway; in this case exploration was efficient and completely guided by item positions. In the second strategy, it was assumed that subjects scanned the whole display with their finger, so exploration was inefficient and completely independent of item positions.

3.2 Methods

3.2.1 Participants

Ten paid subjects (mean age 25 ± 5 years, 4 male) participated in the experiment. One participant was left-handed, while the others were all right-handed according to Coren's test (Coren, 1993). They had normal or corrected-to-normal vision. All subjects were naive as to the purpose of the experiment and gave their informed consent. None of the subjects reported any known hand deficits.

3.2.2 Apparatus

The set-up consisted of a custom-built visuo-haptic workbench. The haptic stimulus was presented using a PHANToM 1.5A force-feedback device. Subjects placed the index finger of their dominant hand in a thimble-like holder that was connected to the PHANToM. The visual stimulus was presented on a computer screen. The subjects looked via a mirror onto the screen such that the visual and haptic stimuli were spatially aligned as illustrated in Figure 3.1a. The finger position was recorded at 50 Hz by sampling the position of the thimble-like holder as a single point in space.

3.2.3 Stimuli

The haptic working range was restricted to the size of the haptic display (15 × 15 cm) in the horizontal plane. Subjects could not move outside of the haptic display, and the edges of the display felt like a continuous wall. The working range was restricted in height such that subjects could raise their finger 4 cm upwards from the display plane, but they were instructed not

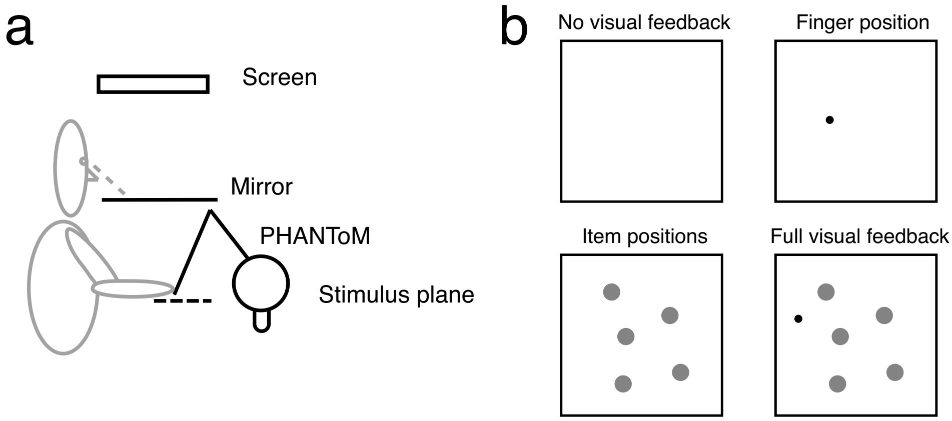

Figure 3.1: a) Illustration of the set-up. Subjects placed their index finger in the holder connected to the force-feedback device that was used to create the haptic display. They viewed the visual stimulus via a mirror projecting the image onto the same plane as the haptic stimulus. b) Examples of the visual display in each of the four conditions.

to lift their finger at all. On the square surface of the haptic display, items consisting of circular areas (1.6 cm diameter) with an increased friction coefficient were placed at random positions (the edges of the items were at least 1.6 cm apart and 1 cm from the boundaries of the display). Both the static and dynamic friction coefficients of the display background were set to 0.2, while distractor items had friction coefficients of 0.5 and the target had both friction coefficients set to 0.8. There could be 3, 5, or 7 items on the display.

There were four different visual conditions, but the haptic display was always defined in the same way. The visual display was represented with a blue square while items were indicated with light-colored disks and finger position with a small sphere. In the first condition, only the square representing the display was shown on the display ('No visual feedback' condition); in the second condition, the square was shown together with the finger position ('Finger position' condition); in the third condition only the square and the item positions were shown ('Item positions' condition); in the last condition the square, the item positions and finger positions were shown ('Full visual feedback' condition). The different conditions are shown in Figure 3.1b. Note that there was never visual information present on which item was the target item.

3.2.4 Experimental design

Subjects were instructed to indicate as fast as possible and accurately whether or not a target item was present. They were informed that the friction coefficient of the target and distractor items would be constant throughout the experiment and that there could at most be one target on the display. They were also told that, like in reality, frictional forces depended on the amount of downward pressure. Responses were made through key presses using keys that were situated next to each other on the keyboard (‘f’ and ‘g’ keys). The response key on the left side corresponded to ‘yes’ and the key on the right side corresponded to ‘no’. To help subjects remember which button corresponded to which answer, the words ’yes’ and ’no’ were shown to the left and right of the visual display, respectively. After pressing a response key, feedback on whether the answer was correct was shown on the screen. Subjects explored the display with their dominant hand, and answered with the other hand. Before the next trial started, subjects moved their finger to the starting position in the upper left corner of the display. During this period, the finger position was shown regardless of the experimental condition.

All conditions were performed in separate blocks of trials. Prior to the experiment, subjects performed a single block of training trials in the full visual feedback mode until they were comfortable with the task and it was clear that they had understood the task. Then, prior to each block of trials they performed at least 20 training trials in the experimental condition of that block. Trials were continued until 9 out of 10 were answered correctly. On average subjects performed 25 $\pm$ 9 training trials and the maximum number of training trials that was needed was 52. Each subject performed all four conditions in a roughly counter-balanced order. Each block consisted of 60 trials (20 trials per number of items) in random order. In half of the trials a target item was present. After 30 trials there was a 5-minute break. The blocks of trials were performed on separate days. Trials that were answered incorrectly were repeated at the end of the block until all trials were answered correctly. If a repeated trial was answered incorrectly then this trial would be repeated again (but this only happened in 25 of the total of 2400 trials). This ensured that there were 10 correctly answered trials for each number of items in each experimental condition. Only the trials that were answered correctly were included in the analysis. Error rates were calculated as the percentage of correctly answered trials of the total number of performed trials.

3.3 Results

The results consist of response time, error rates, and recorded movement tracks. The error rates were well below chance level for all subjects in each of the conditions; statistical analysis (repeated measures ANOVA) of the error rates did not show an effect of condition. There were false negatives (9 % of all trials) as well as false positives (6 % of all trials). Only correct trials were included in the analyses that follow. Figure 3.2 shows a representative selection of tracks over the display of one subject in each of the four

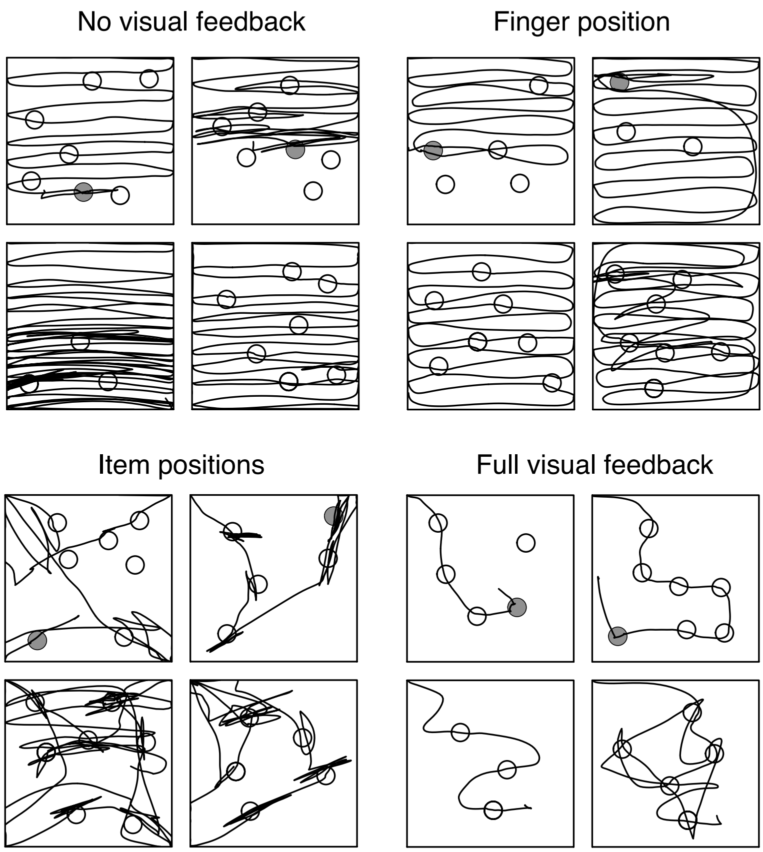

Figure 3.2: A selection of tracked movements over the display from the same subject in each of the four conditions. The upper two panels are always target-present trials (the filled gray disk represents a target), while the bottom two panels show target-absent trials. Trials always started in the upper-left corner.

conditions. For each condition, two target-present and two target-absent trials are shown. It can be seen that there is a clear strategy difference between the conditions in which there was no visual feedback of the item locations compared to the two conditions in which this information was present. In the first case, the subjects systematically scanned the whole display, whereas in the second case exploratory movements concentrated around the item positions. It is clear from the tracks in the 'Item positions' condition that the subject could use visual information about the item positions without visual feedback of the finger position.

3.3.1 Time spent touching the edges

Figure 3.2 suggests that in both conditions without visual feedback on item positions ('No visual feedback' vs. 'Finger position'), subjects touched the edges of the display more often when feedback on finger position was absent. The same holds comparing both conditions with visual feedback on items positions ('Item positions' vs. 'Full visual feedback'). Figure 3.3 shows the percentage of time that the subject spent touching the edges of the display (i.e. finger positions at 2 mm or less from the edges) for each condition. Statistical analysis of the percentages of the duration of a trial that sub-

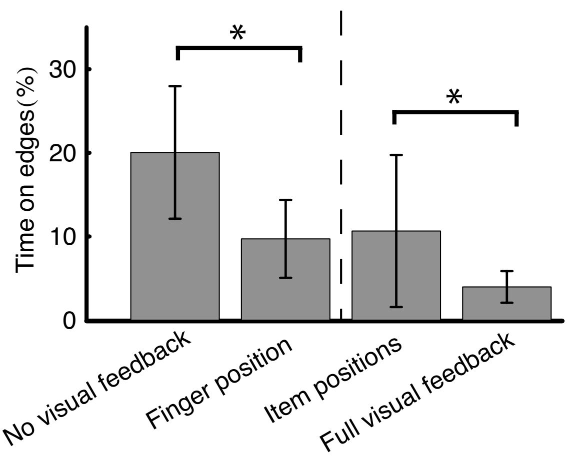

Figure 3.3: Percentage of the duration of a trial subjects spent at distances smaller than 2 mm to the edges of the display, averaged over all subjects for each condition. The error bars indicate the standard deviation of the single subject means. An asterisk indicates that the difference was significant.

jects were touching the edges, showed that there was an effect of condition (repeated measures ANOVA, $F(3, 27) = 15.7$, $p < 0.001$). To determine whether there was an effect of the presence of visual feedback on how much time subjects were touching the edges, post-hoc t-tests were performed to compare the 'No visual feedback' to the 'Finger position' condition and to compare the 'Item positions' condition to the 'Full visual feedback' condition. This analysis showed that the proportion of time subjects touched the edges in the 'No visual feedback' condition was significantly larger than in the 'Finger position' condition ($t = 5.3$, $p < 0.001$) and also significantly larger in the 'Item positions' condition than in the 'Full visual feedback' condition ($t = 2.4$, $p = 0.04$).

3.3.2 Response time slopes

Figure 3.4 shows the response times as a function of the number of items for target-present and target absent-trials in each of the conditions. For the two

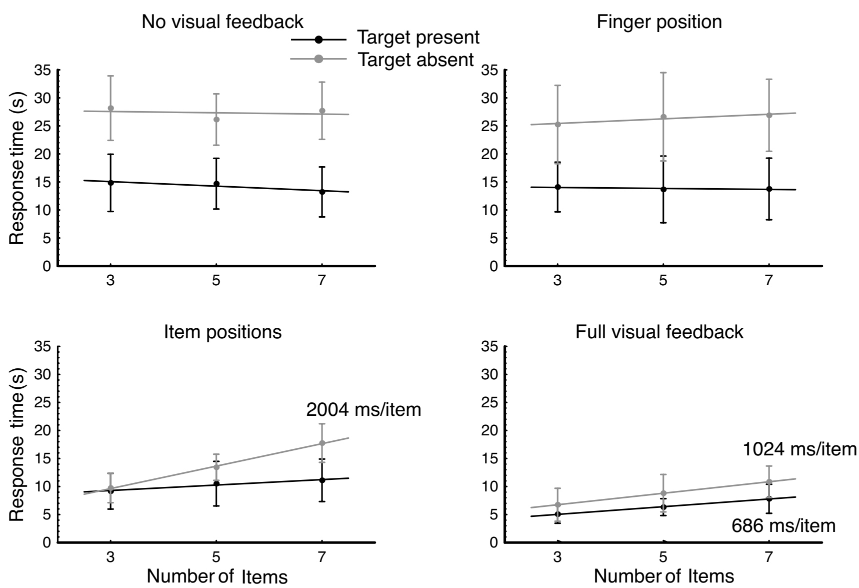

Figure 3.4: Response times averaged over subjects as a function of the number of items for target-present and target-absent trials. Error bars indicate the standard-deviations of the subject means. Solid lines represent linear regression to the mean response times. Slope values are indicated for significant slopes only ($R^2 > 0.9$, $p \leq 0.03$).

conditions without visual information about the item positions the slopes are not significantly different from zero ($p > 0.05$). There is a difference in offset as the target-absent trials yield larger response times than the target-present trials. For the conditions with visual feedback of target positions, the target absent slope was significantly different from zero. The value is indicated in the figure. Both the target present and absent slopes were significantly different from zero for the 'Full visual feedback' condition; the slope values are indicated in the figure and the ratio between the target absent and target present slopes in this last condition was 1.5.

3.3.3 Strategy analysis

In Figure 3.5, the distribution of the distances from the sampled finger position to the center of the nearest item combined is shown for all subjects combined. These distributions can be interpreted as probability density functions of the chance that a finger position was sampled at a certain distance from an item. The bars at distances smaller than the item radius (to the left of the dashed line) represent the time that subjects spent on items. The remainder of the distribution represents the parts of the trials where subjects were moving on the background of the display. It can be seen that this part of the distribution centers on smaller distances for the conditions with visual feedback of item position than for the conditions without this feedback. This means that subjects spent a larger portion of time moving relatively far away from items when visual feedback of item locations was absent than when this feedback was present.

To analyze the differences between the conditions, the distributions were split into distances smaller and larger than the item radius. From the distances smaller than the item radius, the percentage of time that subjects touched an item was calculated (see Figure 3.6a). A large percentage indicates that subjects spent relatively little time on the display background, indicating well-directed movements towards the items. The largest percentage of time was found for the 'Full visual feedback' condition. The distributions of the distances larger than the item radius were analyzed in terms of the mean and the kurtosis, which are shown in Figs. 3.6b and c. A smaller mean indicates that subjects moved on average closer to an item.

The kurtosis is a measure for how heavy the peak in the distribution is; a large value means that a large portion of data was located near the peak and less in the flanks (for comparison: the normal distribution has a kurtosis of 3). Percentage of time on an item, mean distance to an

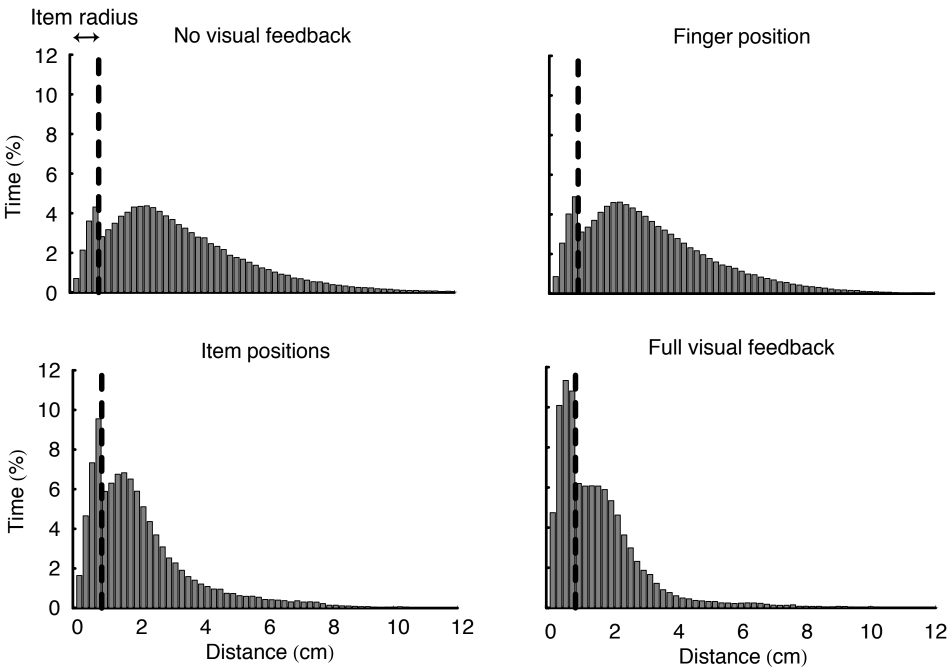

Figure 3.5: Distributions of the distances from the sampled finger position to the center of the nearest item from all subjects in each of the four conditions. The dashed line indicates the item radius (8 mm).

item and kurtosis were calculated from the distributions from each subject. Repeated measures MANOVA was performed on these three measures (Pillai's trace, $F(9, 81) = 13, p \leq 0.001$). Follow-up analysis using univariate tests (ANOVAs) showed that there was a main effect for each measure ($F(1.6, 18) \geq 20, p \leq 0.001$, Greenhouse-Geisser correction was used when appropriate). Post-hoc Bonferroni corrected t-tests showed that most differences between the conditions were significant ($p \leq 0.02$). The non-significant differences between conditions are indicated in Figure 3.6. These results show that when visual feedback of item positions was provided, subjects spent a larger portion of time touching items and less time moving in between items. Furthermore, subjects moved at smaller distances to items and spent less time at distances far away from items when visual feedback of item location was provided.

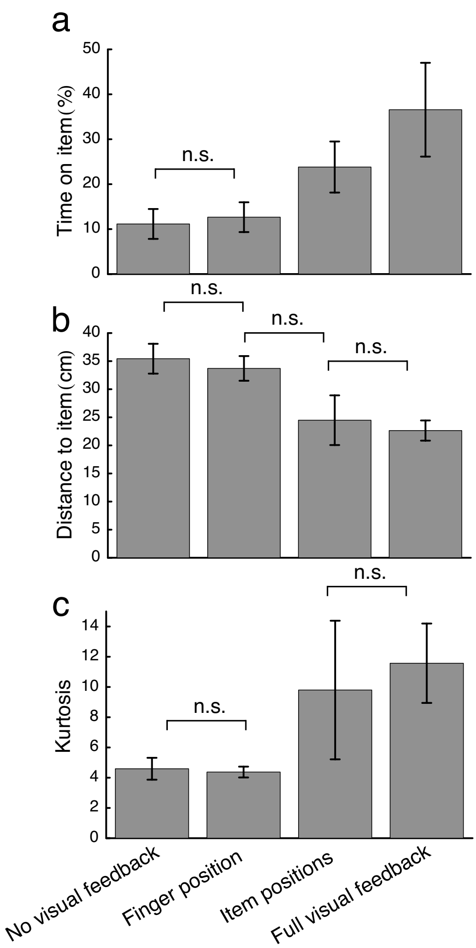

Figure 3.6: Percentage of time on an item (a), distance to the nearest item (b) and the kurtosis (c) of the distribution of sampled distances to the nearest item for each of the conditions averaged over subjects. The error bars indicate the standard-deviation of the single subject means and non-significant differences are indicated.

3.3.4 Global vs. local exploration

When there was no visual information about item positions, subjects scanned the display systematically with their finger. It is possible that they returned to previously visited items after scanning the whole display. To investigate whether subjects did this and whether they were able to use a spatial representation of the items on the display, the tracks from the conditions in which there was no visual feedback of item positions were divided into two parts. To this end, the display was divided into an 8 × 8 grid. Consequently, the grid elements had a height and width of 1.9 cm. This size was in the order of the diameter of an item (1.6 cm), because it can be expected that subjects made scan paths approximately an item diameter apart. Decreasing grid size increases the chance that subjects did not visit a certain element during a trial, while they did search the whole display, which is not desirable. The second part of the track was defined from the moment that all elements in the grid were visited at least once, because from that moment subjects started exploring previously explored parts of the display again. The remaining part of track had to be at least 2 seconds long to be considered as a second part of the track. Scanning direction differed between subjects, but also between trials and even within a trial. This way of defining the track parts works regardless of the subjects' scanning direction. Not all trials had a second part as subjects could answer when they had found a target or immediately after scanning the whole display. Trials without a second part were not included in the analysis. There was a second part in 20% of the trials in the 'No visual feedback' condition and in 31% of the trials from the 'Finger position' condition. Figure 3.7a shows examples of a track with two parts for the 'No visual feedback' condition and for the 'Finger position' condition. It can be seen that particularly in the 'Finger position' condition, well-directed movements towards previously touched items were made during the second part of the trial. In the 'No visual feedback' condition this was not as clearly the case, although in the bottom left panel it can be seen that the subject had a rough idea of where in the display the items were located. The distributions in time of distances from the sampled finger position to the nearest item for the first (light bars) and the second part (dark bars) of trials for all subjects combined are shown in Figure 3.7b. It can be seen that the distributions from the two parts differ mainly in the 'Finger position' condition. The peak from the distribution of the second part is shifted towards smaller distances from items with respect to the peak of the first part. Also, the distribution from the second part of the trials decreases faster for distances far away from items than the distribution from the first part of the trials. This suggests that there was a difference in exploratory strategy between the first and the second part of the trial.

For the distributions from each subject, the percentage of time on an item, mean distance to an item and kurtosis were calculated for the two parts

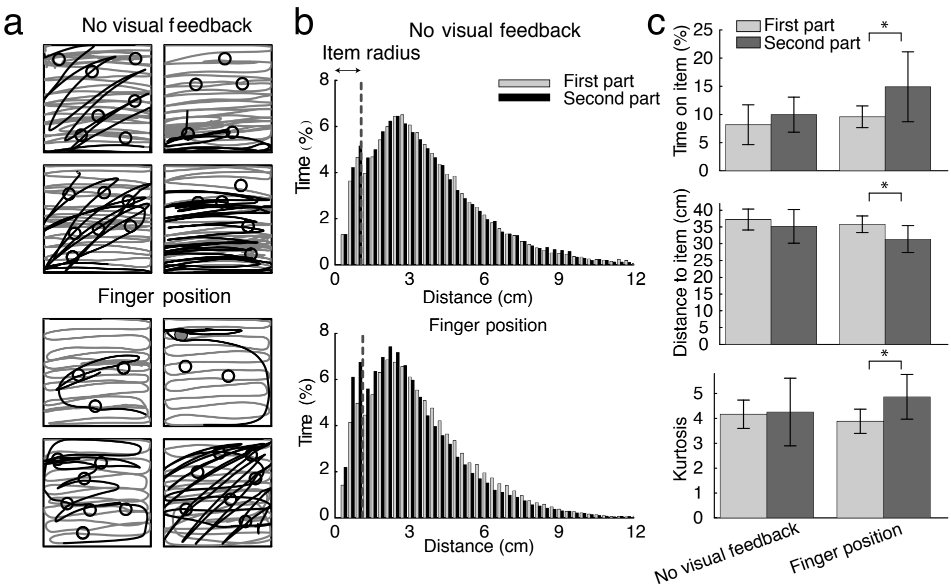

Figure 3.7: a) Four examples of trials with a second part for the 'No visual feedback' and the 'Finger position' conditions. The first part of the track is shown in gray and the second part is shown in black. A target item is indicated with a filled disk. All examples were trials from the same subject. b) Distributions of the distances from the sampled finger position to the nearest item from all subjects for the two stages of the conditions without visual feedback of item positions. c) Time on an item, average distance to an item and kurtosis of these distributions. Error bars represent the standard deviation of the single subject means, and an asterisk indicates a significant difference.

of the trials. Figure 3.7c shows these measures averaged over all subjects. Significant differences between the first and the second part are indicated with an asterisk (paired samples t-tests, $t \geq 2.7, p \leq 0.02$). There were only significant differences between the first and second part in the 'Finger position' condition. In this condition subjects spent a larger proportion of exploration time on items than in the first part. Furthermore, on average, they moved at a smaller distance to items and in combination with the larger kurtosis this indicates that they spent more time near items than further away from items than in the first part of the trial. In Figure 3.7a, it can be seen that sometimes subjects were still systematically scanning after all grid elements were visited. Note that this does make the distributions of the two parts more similar rather than dissimilar. So, the significant differences between the different distributions cannot be due to the criterion we used for splitting up the movement tracks.

3.4 Simulations

Simulations of two extreme search strategies were performed, representing the most efficient and most inefficient strategy. 'Guided search' assumed that the subjects moved with constant speed (corresponding to a movement speed of 10 cm/s or position being sampled every 2 mm at 50 Hz sampling rate) to the nearest untouched item along the shortest pathway. Search was terminated when a target was found or when all items were visited. The resulting distribution of distances to the nearest item is shown in Figure 3.8a.

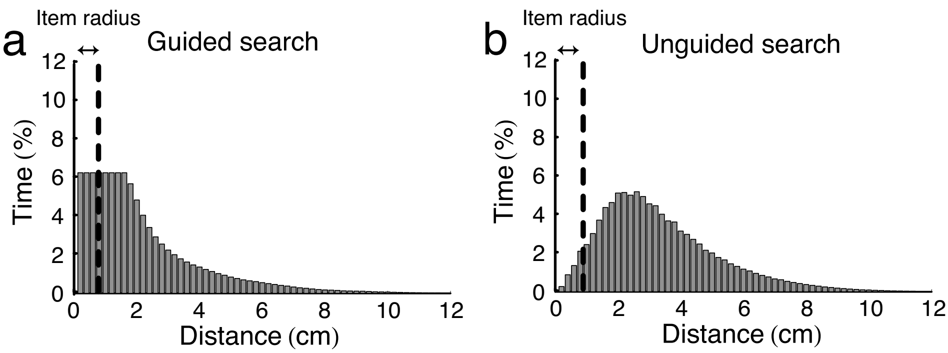

Figure 3.8: Distributions of the distances from the simulated finger position to the nearest item from all displays in the 'Guided search' strategy (a) and the 'Unguided search' strategy (b).

The most inefficient search strategy would be when the whole display was searched, regardless of the item positions. Note that in this case the distribution of the distances is completely driven by the distribution of the items on the displays. The chance that a random point on the display is located at a certain distance from an item is not equal for all distances. For instance, the chance that a random point is very far from an item is quite small. Therefore, in the simulation labeled 'Unguided search', 2500 positions were homogeneously distributed over each display in the set (at 50 Hz sampling rate, this would correspond to a response time of 50 s) and the distance from each position to the nearest item was calculated. The resulting distribution is shown in Figure 3.8b. As the item positions were carefully randomized, the distributions of the distances did not differ significantly for the sets of displays from the different conditions.

Comparison of the distributions of the sampled finger positions (Figure 3.5) to that of the simulations (Figure 3.8) shows that the distribution from the 'No visual feedback' and the 'Finger position' condition resemble the 'Unguided search' simulation, while the 'Full visual feedback' and 'Item positions' conditions are most similar to the 'Guided search' simulation. So, when visual feedback of item locations was present, subjects used a search strategy most similar to the 'Guided search' strategy. Thus, when this feedback was not present, subjects used a strategy similar to the 'Unguided search' strategy. For the 'Finger position' condition it was found that movements during the second part of trials were on average at distances closer to items and had a larger kurtosis than movements during the first part. This shows that the second stage of exploration was shifted towards the 'Guided search' strategy. This indicates that exploration during the second part of trials in this condition was more similar to the conditions with visual feedback on the item positions than the first part. This was, however, not the case for the 'No visual feedback' condition.

In all experimental conditions in Figure 3.5 there is a peak for distances smaller than the item radius, which indicates that subjects spend relatively more time on an item. This peak is absent in the simulations, because a constant movement speed was assumed without distinction between movement on an item or on the background. In the 'Unguided search' strategy the chance that a simulated point was close to or on the center of on item is very small. However, in the 'Guided search' strategy it was assumed that movements were made to the centre of an item and then to the next. This explains why the flat part of the distribution ranges beyond the item radius.

3.5 Discussion and Conclusions

In the present study, we show that adding visual information strongly influences haptic exploratory strategy. In the absence of visual information about item positions subjects systematically scanned the whole display; when information about item positions was added, exploratory movements concentrated around the item positions. Studies into spatial representation have shown that spatial locations and layout can be learned through proprioception. It has been shown that subjects can quite accurately return to a certain target position that has only been briefly touched before (Klatzky & Lederman, 2003). This indicates that the representation of spatial location through proprioception is fairly good. This is in agreement with our finding that subjects can use visual information about item positions in the absence of visual feedback of finger position.

In both conditions without visual feedback of finger position ('No visual feedback' and 'Item location') subjects tended to touch the edges more often than when feedback of finger position was present. The tracks over the display also show that they sometimes followed the edges of the display in this condition. Although it is possible that subjects touched the edges in conditions without visual feedback of finger position because they simply overshot their movement, the fact that subjects often moved along the edges before moving to the next item suggests that subjects used the edges as a reference to re-calibrate their finger position. It has been shown that here is indeed an advantage for creating a spatial representation if an external reference frame (like a bounding square) is provided (Millar & Al-Attar, 2004).

Our data from the conditions without visual feedback of item positions show that subjects sometimes used a two-stage exploratory strategy. First the whole display was scanned and then subjects explored parts of the display again.(Lederman & Klatzky, 1990) have shown that such a two-stage strategy of global exploration followed by local exploration is often present in haptic exploration. An object's shape, for instance, can be explored globally by enclosure, followed by a local exploration procedure like contour following. Interestingly, when there was visual feedback of finger position, exploration in the second stage was clearly different from the first stage. Subjects spent a larger proportion of exploration time on items and moved at distances closer to items. This indicates that subjects had built a spatial representation of the item positions in the display during the first stage and could use this representation to move efficiently back to areas of interest during the second stage. This made the exploratory strategy during the second stage more similar to the strategy used in the conditions with visual feedback on item positions. When there was no feedback of finger position, however, exploratory movements were not correlated more closely to item positions in the second part than in the first part of the track. This shows that forming and using a spatial representation of the display was facilitated by providing visual feedback of finger position.

Because spatial representations can be formed through proprioception alone, the question arises why visual feedback of finger position was required. It has been suggested that visual spatial learning is easier, because in this modality cues like walls of a room that provide a reference frame are readily available (Yamamoto & Shelton, 2007). In another study that was mentioned earlier, it was shown that spatial learning can be aided by providing an external reference frame (Millar & Al-Attar, 2004). In that case subjects explored a map with one hand while touching the external reference frame with the other. In this way, subjects could easily keep track of the position of the exploring finger with respect to the reference frame. In the 'Finger position' condition of the present study, the boundaries of the display and the finger position could be viewed simultaneously; therefore, the position of the finger relative to the display boundaries could also be easily extracted. When the finger position was not shown, extracting this information was much more difficult. This could explain why subjects were able to use a spatial representation of the display in the second stage of exploration in the 'Finger position' condition, but not in the 'No visual feedback' condition. Note, however, that in the present study subjects were not instructed to learn the spatial layout of the display. Therefore, our results do not mean that the layout of the displays could not be learned through proprioception alone. If the subjects were instructed to, they might possibly have been able to do so. Rather, our results show that during a search task, subjects returned to locations where they had previously felt something quite accurately when visual feedback of finger position was provided. It is likely that also in the 'No visual feedback' condition a spatial representation formed during scanning, but probably a much less accurate one than when visual feedback of finger position was available.

Analysis of the response times as a function of the number of items showed that response times were relatively constant for the conditions without visual feedback of item positions. Search strategy was essentially serial in each of the experimental conditions, but in visual search a flat response time slope is usually interpreted as parallel search. This shows that visual search models cannot readily be used to interpret haptic response time slopes. Search strategy analysis showed that there was serial self-terminating search comparable to visual search only when a spatial representation of the display was available. Therefore, in haptic search tasks it is usually important to also analyze the exploratory strategy that was used when interpreting response time slopes (Plaisier et al., 2008a)

Summarizing, visual feedback of item locations could be used to efficiently move from item to item. When this feedback was absent, subjects systematically scanned the whole display. When visual feedback of finger position was provided, they could use the scanning stage to build a spatial representation of the display and move efficiently to items after scanning the whole display. Furthermore, when visual feedback of finger position was absent, subjects used the edges to calibrate their finger position. Finally, response time models from visual search are only applicable to haptic search when a spatial representation of the display is readily available.

Concluding, in tele-operation systems it is clearly most desirable to have full visual feedback, but this may not always be possible as the camera image might be blurred due to fog, for instance. Our results show that providing either visual feedback of finger position only or feedback of item positions can guide haptic exploration. Consequently, in tele-operation systems visual information on the scene can be used to guide exploration even when the probe is not visible. On the other hand, there is also an advantage of showing the position of the probe even if visual information on the scene is poor because the camera image is blurred.

Chapter 4

Salient features in 3-D haptic shape perception

Published as:
Plaisier, M.A., Bergmann Tiest, W.M. & Kappers, A.M.L. Salient features in 3-D haptic shape perception, *Atttention, Perception & Psychophysics*, 71: 421-430, 2009.

Abstract

Shape is an important cue for recognising an object by touch. Several features like edges, curvature, surface area and aspect ratio are associated with three-dimensional shape. To investigate saliency of three-dimensional shape features we developed a haptic search task. The target and distractor items consisted of shapes (cube, sphere, tetrahedron, cylinder, ellipsoid) which differed in several of these features. Exploratory movements were left as unconstrained as possible. Our results show that this type of haptic search task can be performed very efficiently (25 ms/item) and that edges and vertices were the most salient features. Furthermore, very salient local features, like edges, can also be perceived through enclosure, an exploratory procedure usually associated with global shape. Since subjects had to answer as fast as possible, this suggests that speed may be a factor in selecting the appropriate exploratory procedure.

4.1 Introduction

When we reach into our pocket we can easily take out our keys among all other objects we might have in there. However, finding the right key among other keys by touch is much more difficult. Some searches are easy, while others are not. Often it is a specific feature of an object that makes it stand out among the other ones. In the haptic modality such features can be, for instance, material properties, size, weight or shape. This study focuses on the haptic perception of three-dimensional shape and the relative saliency of specific shape features. In this context we consider any shape property that can be used to distinguish two shapes from each other to be a shape 'feature'.

How much an item stands out among other items, i.e. its saliency, has been researched extensively in the visual domain using the visual search paradigm (e.g. Treisman & Gelade, 1980; Wolfe et al., 1989). Usually, subjects are asked to respond as fast as possible whether a certain target item is present among varying numbers of distractor items. Response times are then recorded as a function of the number of items. The additional search time needed per item, or the slope, is a measure of how efficiently the search was performed. When a search is performed at maximum efficiency this slope is near-zero and the target item is said to 'pop out'. For near-zero slopes the search time is thus independent of the number of items and the search is processed in parallel. Search is performed more serially if all items are processed sequentially and response times increase with the number of items. In the target present trials, on average only half of the items are visited before the search is terminated, while in target absent trials all items have to be visited. Therefore, the ratio between the slopes for target absent and target present trials for serial search is often 2. In practice, a wide range of slopes and ratios are found and there is no clear-cut transition between parallel and serial searches (Wolfe, 1998). It would be very valuable to investigate whether such a range of slopes exists in the haptic domain and what the typical values are. This way a framework can be established to compare haptic search slope values, which facilitates the interpretation of these slopes.

When target and distractor item identity are interchanged this sometimes leads to large differences in search slope. Such a difference is labelled a 'search asymmetry'. It has been suggested that these asymmetries arise when two items are distinguished on a single feature that is present in the one item and absent or reduced in the other. When this distinguishing feature is present in the target item and absent in the distractor items the target will ‘pop out’, while the reverse case will yield serial search. Such asymmetries have been used to identify certain features as ‘visual primitives’ (e.g. Treisman & Souther, 1985; Treisman & Gormican, 1988). It is still an open question whether ‘haptic primitives’ exist or even what the definition should be.

The visual search paradigm has been successfully extended to the haptic domain. In the haptic modality there are several ways in which items can be presented to the subjects. For instance, the items can be pressed onto the subjects’ separate fingers (Lederman et al., 1988; Lederman & Klatzky, 1997). These studies showed that especially material properties like roughness might be good candidates for haptic primitives. However, they did not find flat search slopes and in the absence of other information besides response times, these slopes are difficult to interpret. Note that when items are pressed to the fingers, the maximum number of items is, of course, restricted to the number of fingers and also exploratory movements are restricted to small finger movements. The items can also be distributed over a surface allowing subjects to sweep over the surface using their whole hand (Plaisier et al., 2008a). In this study, there were no restrictions on exploratory movements and subjects were free to choose an optimal strategy. For conditions in which the type of movement over the display showed a clearly parallel character, a ratio of 2 was found between slopes for target absent and target present trials. For conditions in which the exploratory strategy had a serial character this ratio was somewhat smaller than 2. Especially when the exploratory strategy was serial, subjects tended to search (parts of) the display repeatedly, because they were uncertain on whether they had searched the whole display. This resulted in a difference in offset, but not a large slope difference. In such a search task, the type of hand movement that was made over the display is a better criterion for distinguishing between parallel and serial search than the ratio between target present and target absent slopes. This shows that exploratory strategies are valuable for interpreting search times and that caution should be taken when using visual search models to interpret haptic search times.

In the haptic modality, the optimal manner of presenting the items depends on the exploratory movements needed to extract a certain type of information. Lederman and Klatzky (1987) investigated which hand movements are typical for extracting various object properties. These included, for instance, lateral motion for roughness perception, pressure for hardness perception and enclosure for global object shape. They have also shown that object recognition can be impaired when these exploratory movements are constrained (Lederman & Klatzky, 2004). Overvliet, Mayer, et al. (2008) found in a search task in which items were three-dimensional shapes fixed in a grid that search times were greatly reduced when the subjects were allowed to enclose the items compared to when they were only allowed to explore them with one finger, showing the effect of constraining exploratory strategies on haptic search times. Therefore, when human performance needs to be optimal, one should be cautious as to the constraints that are put on exploratory movements.

Cutaneous shape perception has been mostly researched in terms of edges and curvature. A sphere is an example of a shape that does not have edges and is defined only by curvature, while a cube does have edges and only flat surfaces (no curvature). Extensive research has been done into the underlying cutaneous signals of tactile shape perception. The mechanoreceptors in the skin have been shown to be sensitive to the edges of stimuli (Phillips & Johnson, 1981). Furthermore, humans can discriminate curvatures that are pressed onto the finger pad (Goodwin, John, & Marceglia, 1991; Jenmalm, Birznieks, Goodwin, & Johansson, 2003) and they can judge the orientation of a cylinder pressed to the finger pad fairly well (Dodson, Goodwin, Browning, & Gehring, 1998). Lederman and Klatzky (1997) found that search for a target item with an edge among distractors without edges pressed to the finger pads was relatively efficient. Note that in all of these studies an item was presented to the finger pads only. In daily life, however, we often hold multiple three-dimensional objects in our hand. These objects can be freely manipulated and rearranged in the hand. Furthermore, curvature and edges are only two of the features that can be used to haptically recognise a three-dimensional shape. Other examples of features that could be used as a cue for shape recognition are surface area, acuteness of the angles and symmetry. In three-dimensional solid shapes several of these features can be present simultaneously and they can be inter-related. Studying perception of isolated shape features is very important, but it is not clear how perception of one feature influences perception of the other. Therefore in the present study, we used well-characterised three-dimensional shapes that differed in several shape features.

To investigate the saliency of shape features in the haptic modality we adopted a search task in which subjects had to grasp multiple items in the hand. The items consisted of three-dimensional shapes suspended from wires. This way of presenting the items allowed the subjects to enclose the shapes and to manipulate and rearrange the items in their hand. Enclosure is, as mentioned before, the typical exploratory procedure for global shape. In this case there is not one object, but there are several items presented simultaneously. The fastest way to explore them would be to grasp them and thereby enclose them in the hand. For difficult searches it could be necessary to explore each item sequentially. To facilitate this, subjects were allowed to release items from the hand. At the same time, releasing items from the hand can be interpreted as an indication of a serial search strategy. This does not mean that search is by definition performed in parallel if no items were released from the hand. However, when items were released from the hand then search had a serial character. Because we left the exploratory movements largely unconstrained, we allowed the subjects to choose the most efficient strategy to optimise their performance.

With this design two experiments were carried out. In the first experiment subjects had to search for a cube among spheres and for a sphere among cubes. If these conditions do not yield the same results then there is a search asymmetry, which would lead to the conclusion that one shape is more salient than the other. In the second experiment three types of shapes were used as target items which were presented among cubes or spheres as distractor items. This will provide insight into the effect of the shape of the distractor items on the efficiency of search for a certain target shape. Note that saliency is in this study then defined with respect to the relative difference between target and distractor items and not as an absolute value. Finally, the results from both experiments are taken into one analysis to investigate the relationship between search efficiency and the difference in several shape features between target and distractor shapes (feature contrast).

4.2 General Method

4.2.1 Participants

Ten paid undergraduate students (6 males, mean age 22± 2 year) participated in both experiments. They were all right-handed according to Coren's test (Coren, 1993) and were treated in accordance with the local guidelines. None of the subjects reported any known hand deficits.

4.2.2 Stimuli and Apparatus

The stimuli consisted of brass shapes, which were suspended from flexible wires. The presented shapes were spheres (radius 0.93 cm), cubes (edge length 1.5 cm), ellipsoids (long radius 1.22 cm, short radius 0.81 cm), cylinders (height 1.63 cm, radius 0.81 cm) and tetrahedra (edge length 3.1 cm). The different shapes are shown in Figure 8.1a. These sizes were chosen such that the volumes of the shapes were equal (3.4 cm^3) so as to prevent weight cues. Although the volume was the same, the shapes differed in several other features. These are summarised in Table 4.1. Here edge acuteness is defined as the smallest dihedral angle (the angle between two planes of the shape). For a sphere and ellipsoid this is then 180°. Vertex acuteness is defined as the solid angle of a vertex of the shape. Curvature indicates the maximum curvature in the shape which is defined as the reciprocal of the radius. The longest axis is defined as the longest cross-section through the shape. The aspect ratio is the ratio of the longest and shortest cross-sections through the shape.

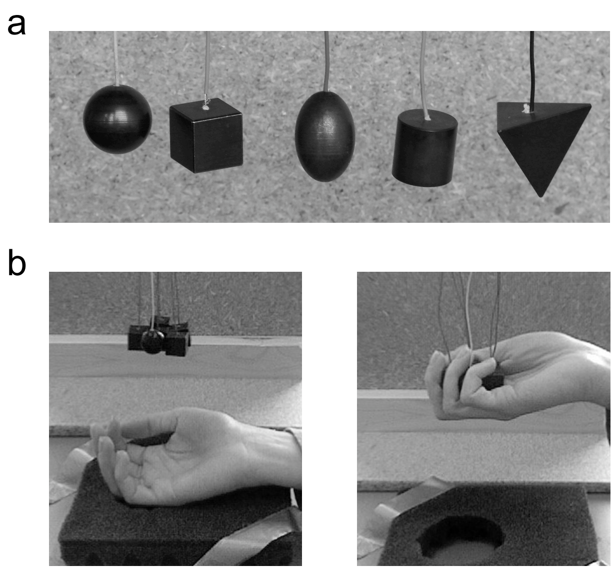

Figure 4.1: Pictures of the stimuli and set-up. a) All of the different shapes used in the experiments. From left to right the sphere, cube, ellipsoid, cylinder and tetrahedron. b) Pictures of a subject grasping the stimuli. In this case the target was a sphere and the distractor items were cubes.

Response times were recorded using a custom built response time measuring device. Time measurement was started when the subject touched any of the items and the measurement was terminated with a vocal response registered with a headset microphone. The resulting response time was then returned by the device with an accuracy of 10 ms (for a detailed description of this device see Plaisier, Bergmann Tiest, and Kappers (2008b)).

Table 4.1: Values of several parameters of the six shapes that were used in the experiments. Note that the volume was constant. Edge acuteness indicates the dihedral angle, vertex acuteness is the solid angle of a vertex, curvature is defined as the reciprocal of the radius, the longest axis indicates the length of the longest cross-section and aspect ratio is the ratio between the longest and shortest cross-sections.

	Shape				
Parameter	Sphere	Cube	Tetrahedron	Cylinder	Ellipsoid
Edge acuteness (degree)	180	90	70	90	180
Vertex acuteness (degree^2)	20627	5157	1791	10313	20627
Edge length (cm)	0	1.5	3.1	5	0
Maximal curvature (cm^{-1})	1.1	0	0	1.2	1.86
Longest axis (cm)	1.9	2.6	3.1	2.3	2.4
Surface area (cm^2)	10.9	13.5	18.7	16.5	11.1
Aspect ratio	1	1.7	1.3	1.4	1.5
Number of edges	0	12	6	2	0
Number of vertices	0	8	4	0	0

4.2.3 Procedure

Prior to the beginning of a trial, blindfolded subjects placed their dominant hand with the palm upwards in a holder. They were instructed to reach upwards and grasp all items simultaneously. This is illustrated in Figure 8.1b. The subjects were instructed to respond as fast as possible whether or not a target item was present by calling out the Dutch equivalents of ‘yes’ or 'no'. It was also emphasised that it was important that the answer was correct. They received feedback from the experimenter on whether the answer was correct. Error trials were repeated at the end of a block of trials and only correct responses were included in the analysis. After initially grasping all items, subjects were allowed to release items from their hand during the trial. It was emphasised that they should only release items from their hand if they thought this was the most efficient strategy. The experimenter scored whether an item was released from the hand during each trial. There were no restrictions on exploratory hand and finger movements.

A total of eight experimental conditions was measured. Search for a cube among spheres (condition 1) and for a sphere among cubes (condition 2) belonged to Experiment 1. The other six conditions (Experiment 2) were searches for an ellipsoid, cylinder or tetrahedron among spheres and searches for these same target shapes among cubes. This means that the distractors were either all spheres or all cubes. All conditions were performed in separate blocks of trials and the subjects were informed of what shape the target and distractor items would have in that particular block of trials. Each block of trials was preceded by a training session. For each condition, subjects performed at least 20 training trials and trials were continued until 10 subsequent trials were correct. It was never necessary to exceed 30 training trials. The subjects were presented with 3, 4, 5, 6 or 7 items. Seven was the maximum number of items that could be held comfortably in one hand. Each condition consisted of 100 trials, 20 trials per number of items. A target item was present in half of the trials. Care was taken that the order in which the eight conditions (of both Experiments 1 and 2) were performed was as close to counterbalanced over subjects as possible. Error rates did not exceed 7 % in any of the conditions.

4.3 Experiment 1

This experiment consisted of two conditions. Subjects had to search for a sphere among cubes or for a cube among spheres. Note that these conditions only differ by interchanging the target and distractor identity. If these conditions yield different results in terms of the slope of the response times as a function of the number of items, this suggests that the one shape is more salient than the other one.

4.3.1 Results

The response times averaged over subjects as a function of the number of items are shown Figure 4.2a. The lines represent linear regression to the response times ($R^2 \geq 0.86$). The target present and target absent slope values for a cube among spheres were 63 ± 8 ms/item (SE) and 200 ± 10 ms/item, respectively. Regression analysis for the sphere among cubes yielded 113 ± 30 ms/item and 520 ± 30 ms/item for target present and absent trials. Note that the ratios between target absent and target present search slopes are rather larger than 2 and they differed considerably between the two conditions. For the cube among spheres this ratio was approximately 3, while for the sphere among cubes it was almost 5.

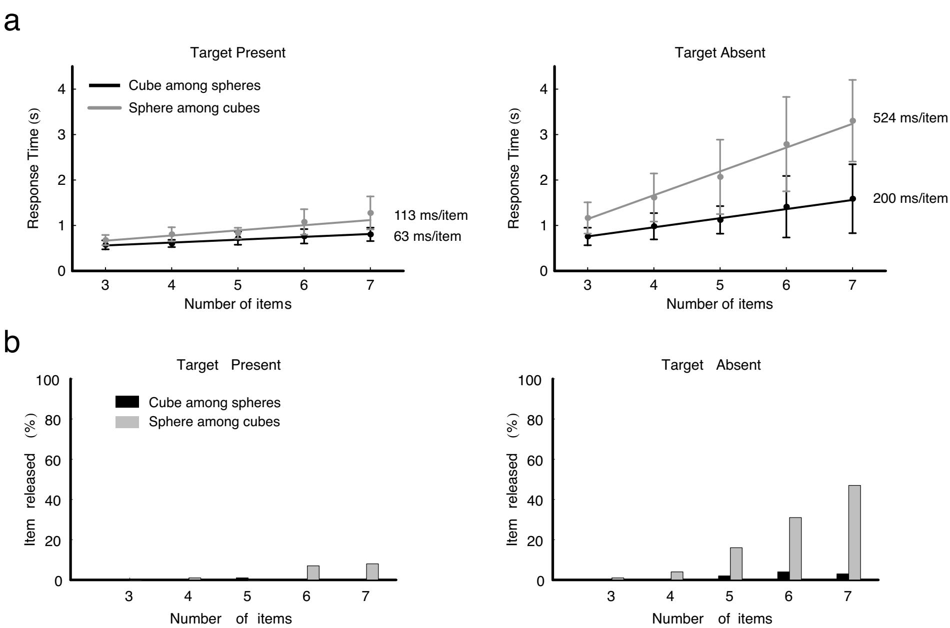

Figure 4.2: Search for a sphere among cubes and a cube among spheres. a) Response times averaged over subjects as a function of the number of items. The error bars indicate the standard error. The left panel shows target present trials for both conditions while the right panel shows target absent trials. The solid lines show linear regression to the data and the slope values are indicated in the figure. b) Bar charts of the percentage of trials in which an item was removed from the hand averaged over subjects for each number of items. Again the left panel shows target present trials and the right panel target absent trials.

A 2 × 2 (condition × target presence) repeated measures Analysis Of Variance (ANOVA) was performed on the slopes of the single subject's response times. This yielded significant main effects for condition ($F(1, 9) = 28.5$, $p < 0.001$,$\eta_p^2 = 0.8$), target presence ($F(1, 9) = 42.9$, $p < 0.001$, $\eta_p^2 = 0.8$) and the interaction term ($F(1, 9) = 14.2$, $p < 0.01$, $\eta_p^2 = 0.6$). The significant effect for condition indicates that there is a search asymmetry. Cube among spheres yields a smaller slope than the reversed condition. The significant effect for target presence shows that slopes were significantly larger for target absent trials, which is commonly found in both haptic and visual search tasks. The significant interaction term indicates that the slope difference between target present and target absent trials depended on the condition. This is clear from the difference in the ratio between the slopes from target present and target absent trials in the two conditions.

Figure 4.2b shows the percentage of trials in which at least one item was released from the hand. It can be seen from this figure that items were mainly released in target absent trials and specifically when the distractors were cubes. A 2 × 2 (condition × target presence) repeated measures ANOVA on the percentage of trials in which an item was released from the hand showed significant main effects for condition ($F(1, 9) = 12.5$, $p < 0.05$, $\eta_p^2 = 0.6$) and target presence ($F(1, 9) = 10.1$, $p < 0.05$, $\eta_p^2 = 0.5$). Also the interaction term between condition and target presence was significant ($F(1, 9) = 13.9$, $p < 0.01$, $\eta_p^2 = 0.6$).

4.3.2 Discussion

The search slopes for target present and target absent trials in the sphere among cubes condition were significantly larger than for a cube among spheres, indicating a search asymmetry. This suggests that a cube among spheres is more salient than a sphere among cubes. It was also found that the ratio between target absent and target present slopes was much larger for the sphere among cubes condition than for the reverse condition. This difference in the ratio between target absent and target present slopes was accompanied by a difference in the number of trials in which items were released from the hand. Analysis of the item release data showed that this happened significantly more often in target absent trials when the distractors were cubes. The analysis of the search slopes together with the item release data indicate that subjects switched their strategy during a trial from only grasping all items in target present trials to releasing items in target absent trials especially when the distractors were cubes. This strategy difference between target present and target absent trials could therefore explain why the ratio between the search slopes was much larger than is usually found in visual search tasks.

The most important difference between the present search task and all previous search tasks reported in the literature is that in the present study item positions could be actively rearranged. The items could be slid along each other by the subject and, therefore, physical interactions between the shapes could have influenced search efficiency. A hand full of cubes is more difficult to manipulate and rearrange than spheres, because cubes do not slide along each other as easily as spheres do. The cubes have, for instance, a larger mutual contact area than spheres, so the frictional forces are larger. A possible explanation for the large difference in target absent and target present search slopes is then that a sphere slides easily out from between the cubes when grasping the shapes, but when there are only cubes this, of course, does not happen. In that case subjects might have been uncertain as to whether there really was no target present and adopted a more serial search strategy leading to larger search times. To investigate the differences in search efficiencies when the distractor items were cubes compared to when distractor items were spheres, Experiment 2 was performed.

4.4 Experiment 2

This experiment consisted of six search conditions. These conditions were searches for an ellipsoid, cylinder or tetrahedron target among spheres and among cubes as distractor items. By comparing search efficiency of one of the target shapes among cubes to search for the same target shape among spheres the effect of distractor shape on search efficiency for each target shape is investigated.

4.4.1 Results

Figure 4.3a shows the response times averaged over subjects as a function of the number of items for each of the target shapes among spheres, while the response times for the target shapes among cubes are shown in Figure 4.3b. The lines represent linear regression to the response times and the obtained slopes are indicated in the figure. This figure shows that search for an ellipsoid among cubes was very efficient, while among spheres it was the most inefficient condition. Also, search for a tetrahedron among spheres was more efficient than search for a tetrahedron among cubes. This

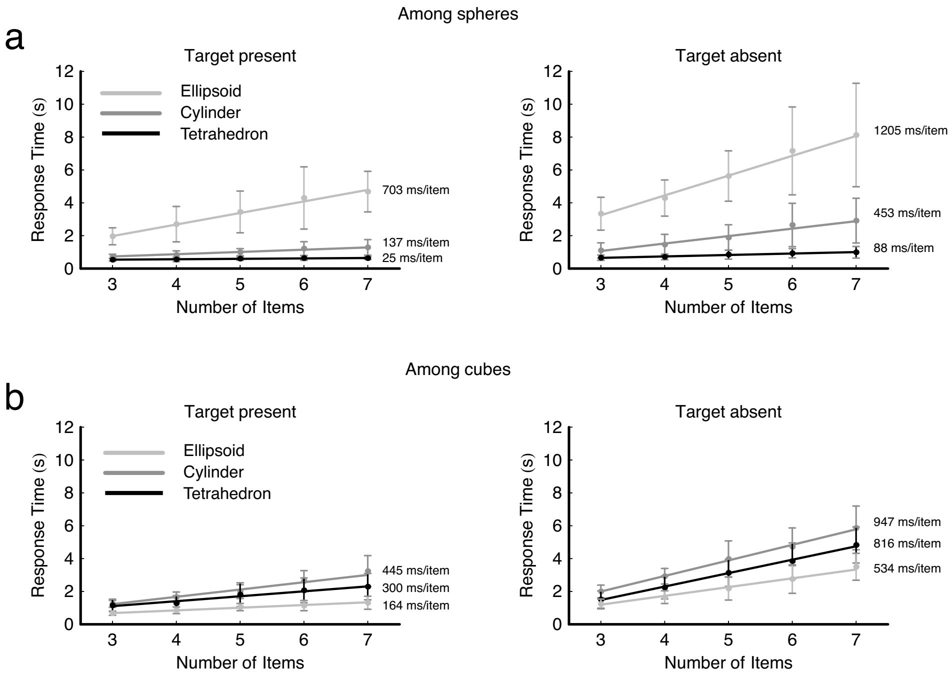

Figure 4.3: Response times averaged over subjects as a function of the number of items for each of the target items. The error bars indicate the standard error. The left panels show target present trials, while the right panels show target absent trials. The solid lines represent linear regression to the data and the slopes are indicated in the figure. a) Distractor items were spheres and b) distractor items were cubes.

indicates that search efficiency depended on whether the distractors were cubes or spheres.

Linear regression was also performed on the single subject's data. Figure 4.5 shows the obtained slopes averaged over subjects. A 2 × 3 × 2 (distractor shape × target shape × target presence) repeated measures ANOVA on these slopes yielded significant main effects for distractor shape ($F(1, 9) = 6.5$, $p < 0.05$, $\eta_p^2 = 0.4$), target shape ($F(1.2, 10.7) = 17.2$, $p < 0.005$, $\eta_p^2 = 0.7$) and target presence ($F(1, 9) = 90.1$, $p < 0.001$, $\eta_p^2 = 0.9$). The interaction terms between distractor shape and target shape ($F(1.3, 11.6) = 51.2$, $p < 0.001$, $\eta_p^2 = 0.8$) and between distractor shape and target presence ($F(1, 9) = 14.5$, $p < 0.01$, $\eta_p^2 = 0.6$) as well as the interaction between the distractor shape and target presence ($F(2, 18) = 5.8$, $p < 0.05$, $\eta_p^2 = 0.4$)

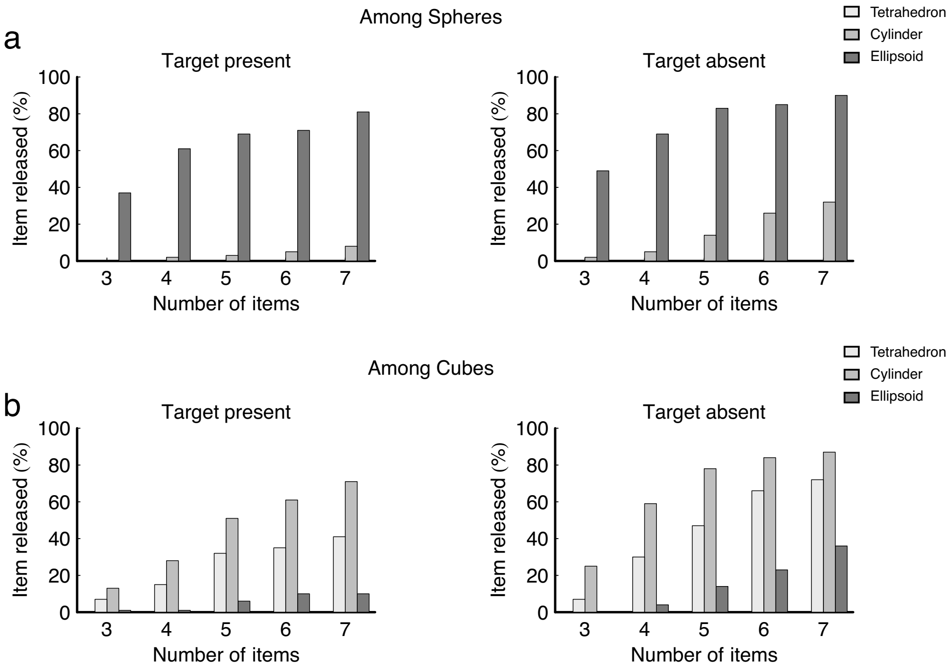

Figure 4.4: Percentage of trials in which an item was removed from the hand for each of the target items as a function of the number of items. The left panels show target present trials and the right panels show target absent trials. a) Distractor items were spheres and b) distractor items were cubes.

were significant. Also the interaction between distractor shape, target shape and target presence was significant ($F(2, 18) = 16.6$, $p < 0.001$ $\eta_p^2 = 0.6$).

Figure 4.4 shows that the percentage of trials in which an item was released from the hand increased with the number of items and that the percentage of item release trials was largest in target absent trials for both distractor shapes. In general, in a larger percentage of trials an item was removed when the distractors were cubes than when the distractor items were spheres. A 2 × 3 × 2 (distractor shape × target shape × target presence) repeated measures ANOVA on the percentage of trials in which an item was released showed significant main effects for target shape ($F(2, 18) = 14.8$, $p < 0.001$, $\eta_p^2 = 0.6$) and target presence ($F(1, 9) = 54.7$, $p < 0.001$, $\eta_p^2 = 0.9$). The factor of distractor shape did not reach significance. However, the interaction between distractor shape and target shape ($F(1.3, 11.7) = 43.5$, $p < 0.001$, $\eta_p^2 = 0.8$) as well as the interaction between distractor shape

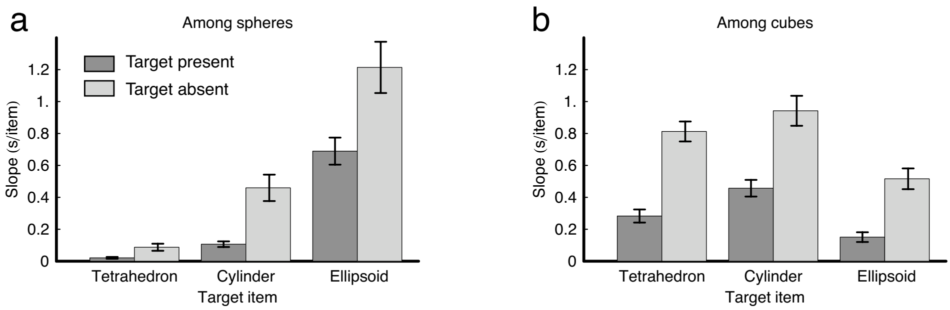

Figure 4.5: Bar charts of the averaged slopes from regression to the individual subject's data for each of the target items. The error bars indicate the standard error. a) Distractor items were spheres and b) distractor items were cubes.

and target presence ($F(1, 9) = 11.3$, $p < 0.01$, $\eta_p^2 = 0.6$) were significant.

For both the search slopes and the percentages of item releases there was an interaction found between target presence and distractor shape. In Figure 4.5 it can be seen that the ratio between target absent and target present search slopes varied for the different conditions. Also, the ratio between the percentage of item releases between target absent and target present trials varied over the conditions. There was a significant positive relationship between the search slope ratios and the item release ratios ($r = 0.38$, p (two tailed) < 0.001).

4.4.2 Discussion

The slopes from the different conditions cover a large range of values and search efficiency ranged from highly efficient (25 ms/item) to quite inefficient (703 ms/item). Note that search efficiency for a certain target depended heavily on the identity of the distractor items. For instance, search for an ellipsoid among cubes was relatively efficient, while search for an ellipsoid among spheres was performed the least efficient of all conditions in the experiment. The reversed result was found when the target was the tetrahedron. Search for the tetrahedron among spheres was performed more efficiently than all other search conditions in the experiment, but search for the tetrahedron among cubes was much less efficient.

These differences in search slopes were accompanied by differences in the percentage of trials in which an item was released from the hand. The conditions for which this percentage was high were also the conditions which yielded relatively large search slopes. This indicates that the search strategy that was adopted depended on the specific combination of target shape and distractor shapes. As was also found in Experiment 1, subjects released items from the hand more often when distractors were cubes than when the distractor items were spheres. Also, analysis of the slopes showed that search slopes were on average larger when the distractors were cubes. There are two possible explanations for this strategy difference. First, the relatively intense stimulation from the edges and vertices of the cubes may make it difficult to find a target among cubes and as a consequence subjects tend to remove an item from their hand more often. Another explanation mentioned earlier suggested that the difference between cubes and spheres as distractors is due to the physical interactions between the shapes when they slide along each other. Also, the interaction between distractor shape and target presence found in the analysis of the slopes as well as the item release data could be related to the sliding of the shapes along each other. A certain target shape might easily slide out from among the distractor items, but in target absent trials this does not happen and subjects adopted a more serial search strategy. Both of these explanations are in agreement with the data and possibly both effects play a role here.

A difference in strategy between target present and target absent trials can explain a large ratio between the associated slopes. The results showed that when there was a large difference in search slope between target present and target absent trials, there was also a large difference between the percentages of item releases for target present and target absent trials. This suggests that a large difference in search slope between target present and target absent trials may indeed have been caused by a change in search strategy. Strategy differences also explain why the general range of search slope values that was found is so large. When items are released from the hand, the extra search time per item is considerable and this will result in a much larger search slope than for conditions in which no items were released. Therefore, haptic search slopes from different conditions of one experiment can span a much larger range than is usually found in vision. This is especially true under conditions of free exploration, where subjects actively explore the stimulus.

4.5 Relationship between shape features and search efficiency

To investigate which specific features of three-dimensional shape were most salient the slopes of Experiment 1 and 2 were included in one analysis. The shape features that were taken into account are: edge acuteness, vertex acuteness, edge length, maximal curvature, length of the longest axis, surface area, aspect ratio, the number of edges and the number of vertices. Note again that the volume was the same for all of the shapes. The values of these parameters for the different shapes are summarised in Table 4.1. Because the search efficiency depended on the specific combination of target and distractor shapes, the absolute difference between the values of the parameters of the target shape and those of the distractor shape were taken for each of the conditions in Experiments 1 and 2, designated by Δparameter. Linear regression was then performed on the target present search slopes as a function of these absolute differences. For each of the parameters this is shown in Figure 4.6.

In the set of shapes that was used the edge acuteness and vertex acuteness were highly correlated, as were the number of edges and the number of vertices (Pearson's $r > 0.9$). To avoid collinearity problems of the pairs of correlated features, the feature with the smaller R^2 value was not taken into the multiple regression analysis. This means that edge acuteness and number of edges were taken into the analysis and vertex acuteness and number of vertices were not. Multiple regression analysis with stepwise variable entry was performed on the aforementioned parameters with the target present slopes as dependent variable. The regression analysis showed that only acuteness of the edges improved the model significantly ($R^2 = 0.91, p = 0.002$). This indicates that this was the dominant parameter for search efficiency. Note that because edge acuteness and vertex acuteness were highly correlated we cannot distinguish between these two features. Therefore, we conclude that edges and vertices are the most salient features in haptic perception of three-dimensional shape. Search is most efficient when the target shape has edges and the distractors do not, or vice versa. The number of edges, however, did not show a significant relationship with the search slopes. Of course, if both target and distractors do or do not have edges, search efficiency will depend on other parameters and these results do not mean that other features cannot be used for haptic shape recognition. These results could be interpreted in terms of similarity between target and distractor items. In vision it has been proposed

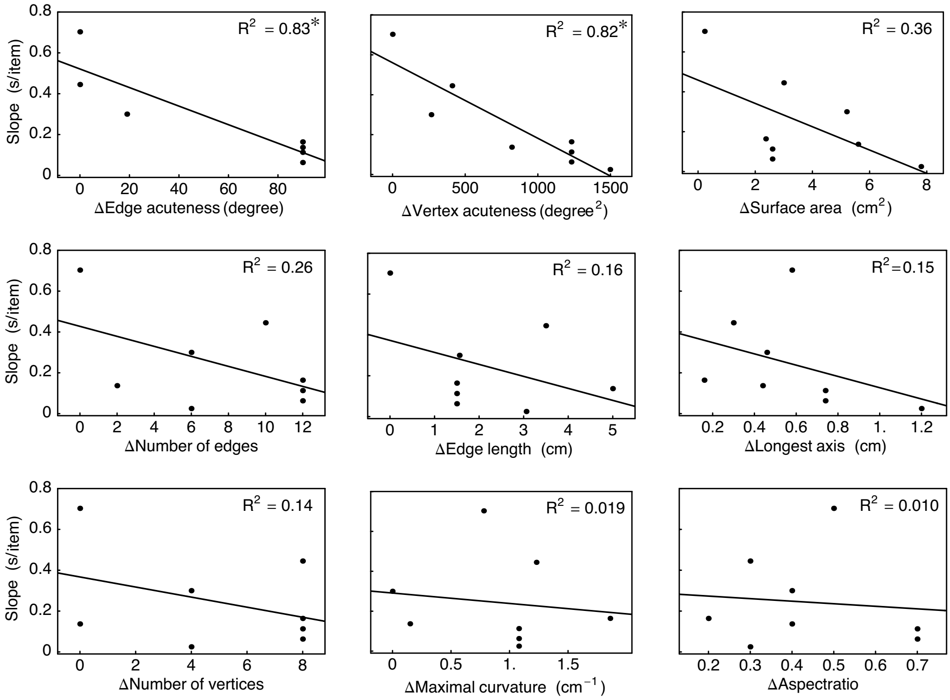

Figure 4.6: Target present search slopes as a function of the absolute difference between target and distractor shape in several parameters of the shapes. These parameters were edge acuteness, curvature, longest axis, surface area, aspect ratio and the number of edges. An asterisk indicates that the relationship was significant ($p < 0.05$). Note that this was only the case for edge acuteness and vertex acuteness.

that difficulty of a search task increases with increasing similarity between target and distractor items (Duncan & Humphreys, 1989). In relation to this it has also been shown that saliency of a target item increases when there is a feature contrast between target and distractors in several dimensions (Nothdurft, 2000). If we extrapolate this to our study, this would indeed predict that a cylinder among cubes is more salient than a tetrahedron among cubes, since the cylinder differs from the cube in both edge acuteness and curvature, while the tetrahedron only differs mainly in edge acuteness.

4.6 General Discussion

It has been shown that haptic object recognition can be very fast and accurate (Klatzky et al., 1985; Klatzky & Lederman, 1995). To mediate fast object recognition there must be certain object features, like shape, that are extracted and processed very efficiently. Unfortunately, only few studies are available on haptic perception of three-dimensional shape. The studies that are availably are mainly concerned with the comparison between haptic and visual shape perception. One of these studies, in which subjects had to haptically explore solid shapes and then recognise the felt shape visually, showed that this task could be performed fairly well (Norman, Norman, Clayton, Lianekhammy, & Zielke, 2004). This suggests that there is some shared underlying representation of shape for both modalities. However, in this study there were also differences found between the modalities. When the shapes were first presented visually and had to be recognised haptically, the authors found that shapes with similar 'global shape' tended to be confused. This could be due to a difference in weighing of global and local shape features between the haptic and visual modality. It has indeed been shown that such a difference exists. Lakatos and Marks (1999) showed that in vision global shape properties are weighed more heavily than local shape properties. In the haptic modality, on the other hand, global and local shape properties were weighed equally.

For haptic perception of global and local shape properties, different types of exploratory procedures (EPs) have been identified (Lederman & Klatzky, 1987). The typical EP for global shape is enclosure, while for local shape the EP is contour following. In the aforementioned study by Lakatos and Marks, subjects wore (splinted) gloves to force them to explore the shape through enclosure alone in an attempt to bias them towards global shape properties. However, this did not produce different similarity ratings between the objects compared to when the objects were explored without the glove. In their paper they suggest that hand movements are not so specialised as is sometimes thought. It could be that also some information about local shape can be extracted through enclosure. Global and local shape properties can also be identified in the present study. The global shape can be interpreted as the shape of the cluster of the individual shapes when grasped together in the hand. Local shape properties are then the shape features of the individual shapes.

In our study also different exploratory strategies were used. When search was performed efficiently, no items were released from the hand and the target could be detected through enclosure (e.g. tetrahedron among spheres). Note that the global shape of the cluster of items was the same for target present and target absent trials. So, in this case *local* shape properties could be detected through an exploratory procedure which is associated with *global* shape. Since subjects were instructed to perform the task as fast as possible, this suggests that speed can be a factor in selecting an EP. Lederman and Klatzky (1990) showed that there is a typical two-stage sequence for haptic object exploration consisting of initially grasping (enclosing) the object followed by a more specialised EP if necessary. In the present study, enclosure was not for all of the search conditions sufficient to find the target. In some, less efficient, conditions the individual shapes were explored through a serial strategy and items would be removed from the hand (e.g. ellipsoid among spheres). This suggests that the most salient local shape features can be perceived through enclosure, while less salient local shape features cannot be perceived this way.

Our results show that edges and vertices are very salient features of three-dimensional shape. Search was performed efficiently if the target shape had edges while the distractor shapes did not, or vice versa. Efficiency of search for an edge has been studied before (Lederman & Klatzky, 1997). In that study, items were presented by pressing surfaces onto the subjects' individual fingers. This could then be a continuous flat surface, or a flat surface with an raised edge on it. In that study no asymmetry was found between search for an edge among no-edge items and the other way around. However, the comparison between the results from that study to those from the present study is not straightforward. The exploratory movements in the present study were very different from the small finger movements that can be made when surfaces are pressed onto the fingers. Moreover, in the present study the items consisted of solid three-dimensional shapes and in such shapes many features are present simultaneously and they can also be inter-related. For instance, a closed three-dimensional shape that does not have edges must necessarily have curvature. The asymmetry between search for a cube among spheres and search for a sphere among cubes that was found in Experiment 1 of the current study can be interpreted in two ways. It could be interpreted as a search for the presence of an edge being more efficient than search for the absence of an edge, or as search for an edge being more efficient than search for curvature. This last interpretation is supported by the results from Experiment 2 which showed that search for a cylinder among cubes was less efficient than search for a tetrahedron among cubes, while a cylinder has curvature and cubes and tetrahedrons do not.

Haptic shape perception has often been studied by studying perception of isolated cues. Examples are research into the saliency of isolated shape features like in the study of Lederman and Klatzky (1997) or measurements perceptual thresholds for, for instance, curvature (Gordon & Morison, 1982; Goodwin et al., 1991; Pont, Kappers, & Koenderink, 1997; Van der Horst & Kappers, 2007). Although these studies do provide important information on the perception of shape features, in three-dimensional shapes many shape features are present simultaneously and they are often related. It is not *a priori* clear how these different features are combined and therefore the present study, in which these features are combined is very valuable. Furthermore, free exploration, as used in the present study, allows the subjects to optimise their perception (Lederman & Klatzky, 1987) and facilitates comparison to haptic exploration in daily life. Moreover, differences in exploratory movements across search conditions can provide information on the saliency of a target item. We have suggested before that the use of global exploratory movements combined with high search efficiency may be a useful definition for a haptic version of the 'pop-out' effect (Plaisier et al., 2008a). According to that definition, edges and vertices would certainly qualify as a shape features that 'pop out'.

Chapter 5

The role of item fixation in haptic search

Published as:
Plaisier, M.A., Kuling, I.A., Bergmann Tiest, W.M., and Kappers, A.M.L. The role of item fixation in haptic search, in *Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 417-421, 2009.

Abstract
Enclosing objects in the hand is a common and efficient way of haptic exploration. Recently, the importance of grasping for more realistic haptic perception of virtual objects has been recognised in haptic interface design. While several studies on haptic perception have addressed haptic exploration of a single object, perception of several objects grasped together in the hand has received almost no attention yet. In this study we focus on the importance of freedom to manipulate the objects in the hand for three-dimensional shape perception. Furthermore, we investigate differences in detection speed for different positions in the grasping hand. Subjects were asked to search for a cube among spheres or for a sphere among cubes. Response times were measured for different locations of target shape in the hand. Also, the way in which the items were fixed was varied from allowing small displacements and rotation of the shapes to rigidly fixed. There were only differences in search times between the different positions in the hand, when the centre item was difficult to access because of the surrounding items. Finally, we show that search was faster when the items were rigidly fixed than when displacement and rotation was possible. This shows that more exploratory freedom does not necessarily make search for a three-dimensional shape faster.

5.1 Introduction

It has been shown that typical exploratory procedures can be identified for haptically extracting certain types of information. Examples are lateral motion for perceiving roughness, lifting the objects for assessing weight and enclosure for global shape perception (Lederman & Klatzky, 1987). Furthermore, it has been shown that constraining the exploratory movements can impair haptic object recognition (Lederman & Klatzky, 2004). In haptic display design, the importance of haptic exploration using the whole hand and enclosing virtual objects for a more realistic perception of three-dimensional objects has been recognised. Devices like the CyberGrasp enable force-feedback for grasping virtual objects (*Immersion Corporation: CyberGrasp*, n.d.). Also, there are displays being developed that give the sensation that an object is held between two hands by applying pressure and shear forces to the palm (Minamizawa, Kamura, Kawakami, & Tachi, 2008). While the importance of haptic perception of objects through grasping objects in the hand has been made clear, relatively few studies have addressed perception of objects enclosed in the hand.

In daily life it is a common situation to have multiple objects grasped together in the hand. We can extract information about the shapes, sizes and materials of the objects in our hand by grasping them and making finger and hand movements. The object positions in the hand can be freely rearranged and the shapes can be rotated over an arbitrary axis. The question arises how important it is that the objects' positions and orientations in the hand can be freely manipulated for extracting information about the objects.

It has also been shown that there is a typical two-stage procedure for exploring an object (Lederman & Klatzky, 1990). First, there is a grasp and lift stage, which is followed by more detailed exploration like contour following. A similar procedure could be identified for haptic exploration of several three-dimensional shapes grasped together in the hand. After grasping the shapes in the hand, they can be explored more extensively by making hand and finger movements or they can even be explored subsequently if necessary(Plaisier, Bergmann Tiest, & Kappers, 2009b). When exploring several objects simultaneously, exploratory freedom can be constrained by fixing the objects' relative positions and orientations in space. In that case the objects cannot be rearranged in the hand and item manipulation is constrained. This constraint also means that if you are trying to find a particular object among other objects in your hand, this so-called 'target item' may be easier to locate when it makes contact with certain parts of the hand than other parts. There are several factors that might cause differences in detection speed. Firstly, it may depend on the number of nerve endings in the skin of the different parts of the hand. Estimations of the number of nerve endings for different areas of the hand showed that density is relatively high for the fingers and substantially smaller for the palm of the hand (Johansson & Vallbo, 1979). It has also been shown that spatial resolution reported is higher for the fingers than for the palm of the hand (Craig & Lyle, 2001). Secondly, when grasping rigid objects, detection speed may also depend on the forces that can be exerted on the objects grasped at different locations in the hand. Finally, response times may also depend on the ability to focus attention to different parts of the hand.

Search tasks have often been used in vision to investigate saliency of certain object features. Typically, the task is to respond as fast as possible whether a certain target item is present among distractor items. Usually, the number of items is varied and response times are measured as a function of the number of items. The slope of this function is then a measure of the efficiency with which the search was performed. This paradigm has been successfully expanded to the haptic domain. In several studies a haptic search task has been used to investigate saliency of several object properties. There are different ways in which the items can be presented. They can be pressed onto the separate fingers (Lederman et al., 1988; Lederman & Klatzky, 1997) or they can be presented in a two-dimensional plane (Plaisier et al., 2008a; Overvliet, Mayer, et al., 2008). Recently, we have added a new type of haptic search task in which three-dimensional shapes are grasped together in the hand (Plaisier, Bergmann Tiest, & Kappers, 2009b). In this case the shapes were suspended from flexible wires which allowed subjects to freely rotate and translate the shapes. Note that this also allowed the subjects to manipulate the item positions with respect to each other.

This last study showed that search for a target shape with edges among distractor shapes without edges or vice versa is highly efficient under relatively unconstrained item manipulation conditions. While that study aimed at investigating saliency of three-dimensional shape features, the present study aims at providing more insight into the influence of constraining translational and rotational freedom of the shapes in the hand on search times. By fixating the items, the subjects' freedom of moving the items in the hand is constrained, while exploratory movements are left unconstrained. This is a new way of constraining exploratory freedom. Constraining exploratory freedom by constraining exploratory movements, has been previously shown to impair haptic information extraction (Lederman & Klatzky, 2004). To investigate whether constraining exploratory freedom by fixating the items leads to an increase in search times, the way in which the items were fixed in space was varied. A consequence is that the relative positions of the items cannot be manipulated like in the previous study. For instance, if the target item is initially in the centre it cannot be moved out from among the other items. In order to see how this affects target detection, we investigated differences in search times for different locations of the target shape in the hand. Note that we did not intend to measure sensitivity of different part of the skin on the hand. Rather we investigated processing of information across different positions in the grasping hand. Therefore, we let subjects perform a search task by actively grasping the stimuli and we let exploratory movements largely unconstrained.

In the present study again a search task where the items consisted of three-dimensional shapes was used, but in this case response times were not measured as a function of the number of items. Rather, the number of items was kept constant, while the target position was varied. Response times were thus measured as a function of target position. Furthermore, there were three different ways in which the items could be fixed. One way of fixation allowed for small displacements and rotation of the items. In the two other fixation methods the items were rigidly fixed to a certain position and orientation with all items in the same plane (bottoms levelled) or with the centre item lowered. In this last condition the global shape of the plane in which the items are fixed is curved and corresponds roughly with the shape of the grasping hand, which might reduce response times. The task was either to search for a cube among spheres or for a sphere among cubes. These shapes were chosen because it has been shown that in this case the target can usually be detected through only grasping the items together in the hand without more detailed exploration being necessary (Plaisier, Bergmann Tiest, & Kappers, 2009b).

5.2 Method

5.2.1 Participants

Eight paid undergraduate students (mean age 20 ± 2 years) participated in the experiment. Four of them were female. Seven subjects were right-handed and one was left-handed according to Coren's test (Coren, 1993). None of the subjects reported any known hand deficits. Participants were treated in accordance with the local guidelines and gave their informed consent.

5.2.2 Stimuli and Apparatus

The stimuli consisted of brass cubes (edge length 1.5 cm) and spheres (radius 0.93 cm). These sizes were chosen such that the volume of both shapes was constant to eliminate possible volume and weight cues. The number of items presented was always five and the shapes were spatially located at the positions indicated in Figure 5.1a. In this figure it can be seen that each position corresponded with a certain region of the hand. Roughly these were: near the thumb (position 1), near the wrist (position 2), near the little finger (position 3), near the middle finger (position 4) and in the centre (position 5). These positions were only roughly defined. Since subjects actively grasped the items, a more precise definition of the item positions is not possible and for the purpose of the experiment this was sufficient.

In some conditions the items were fixed to a metal tube through a rigid connection (Figure 5.1b). This method of fixation did not allow for movement of the items with respect to each other and did not allow for rotation of the shapes. All items were either fixed in the same plane with the bottoms of the shapes levelled (fixed), or the centre item was lowered 0.5 cm with respect to the other items (fixed with centre item lowered). The items could also be fixed in a way that did allow for the items to be displaced slightly and to be rotated. This was accomplished by suspending the shape from wires. The wires were pulled through the tubes leaving 0.5 cm of wire from the bottom of the tube (partially fixed). For each of the three methods of fixation two search tasks were performed. The task was either to find a cube among spheres, or the target could be a sphere among cubes as distractor items. This added up to a total of six conditions.

Response times were recorded using a custom-built response time measuring device. Time measurement was started when the subject touched the stimulus activating the touch sensitive contact of the device. The measurement was terminated with a vocal response registered with a headset microphone. The resulting response time was then recorded with an accuracy of 10 ms (for technical details about this set-up see (Plaisier et al., 2008b)).

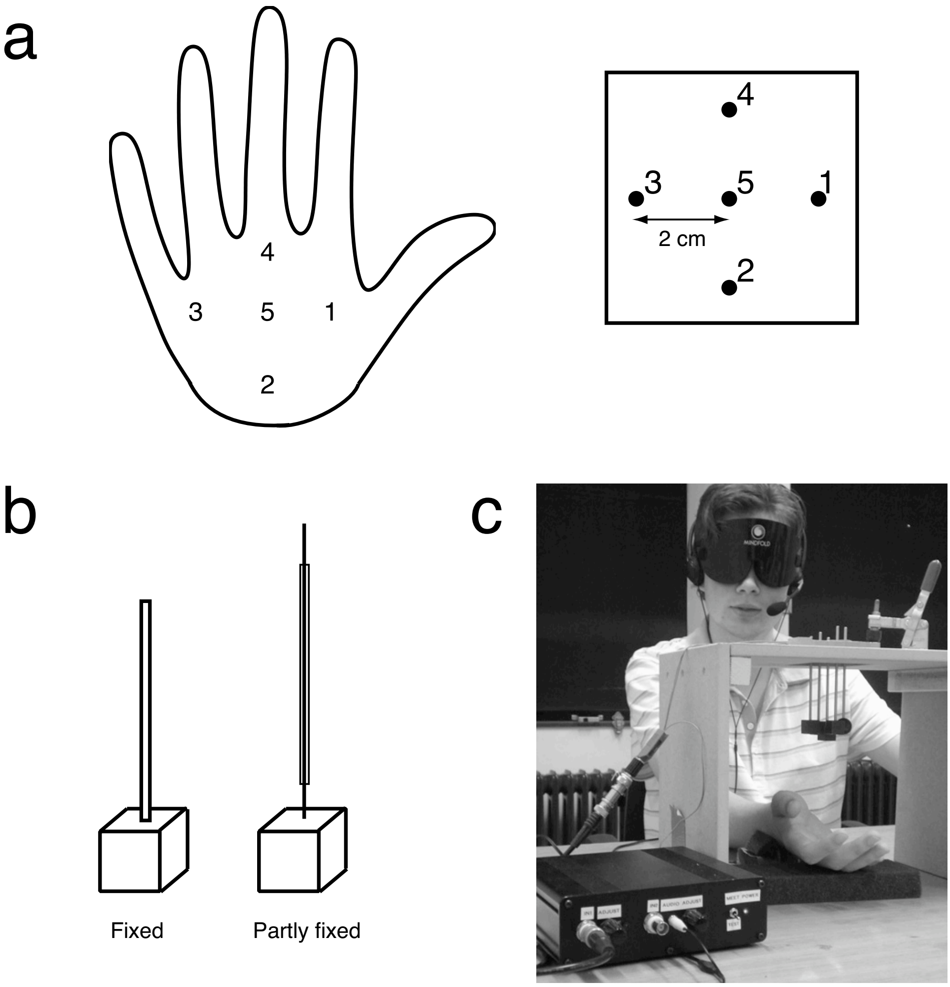

Figure 5.1: a) Shows a right hand with palm facing towards the front and the item positions are labelled according to the positions in the grid at which they could be fixed, shown on the right. b) Schematic drawing of the different methods of item fixation (in this case a cube). In the left-hand image the cube is rigidly fixed to a metal tube (fixed), while in the right-hand image the cube is fixed to a flexible wire which is pulled through a metal tube (partly fixed). The free wire end had a length of 0.5 cm. c) A picture of a subject in the set-up.

5.2.3 Procedure

The subjects were blindfolded and they placed their dominant hand with the palm facing upward in a holder prior to the beginning of a trial (Figure 5.1c). There were also holders for the elbow and for the wrist to keep the starting position of the hand and arm the same prior to each trial. They were instructed to reach upward and grab all items simultaneously. It was also emphasised that they should try to grasp the items in the same way each trial. Other than this there were no restrictions on exploratory hand and finger movements. Each of the subjects could place the fingers in between the tubes to completely enclose all of the shapes simultaneously with the hand. Subjects were instructed to respond whether a target item was present as fast as possible by calling out the Dutch equivalents of 'yes' and 'no'. It was also emphasised that it was important that the answer was correct. The subjects received feedback on whether the answer was correct from the experimenter. Error trials were repeated at the end of the block and only correct responses were included in the analysis.

All conditions were performed in separate blocks of trials and each block of trials was preceded by a training session. To allow subjects to optimise their performance, they were informed prior to the experiment of which shape the target item and which shape the distractor items in that particular block of trials would have. Subjects performed at least 20 training trials and trials were continued until 10 subsequent trials were correct. It was never necessary to exceed 20 training trials. Each condition was performed in a separate block of trials. The target could be located at any of the five item positions. For each target position 10 trials were performed totalling 50 trials. There were also 50 trials in which the target item was absent such that in half of the trials there was a target item present. All of the subjects performed each of the six experimental conditions. Care was taken that the order in which the conditions were performed was as close to counterbalanced over subjects as possible.

5.3 Results

First the target present and target absent trials were analysed regardless of target position. Figure 5.2a shows the response times averaged over subjects for target present and target absent trials for each of the conditions. A $3 \times 2 \times 2$ (fixation method $\times$ target shape $\times$ target presence) repeated measures ANOVA (Analysis Of Variance) was performed on the response

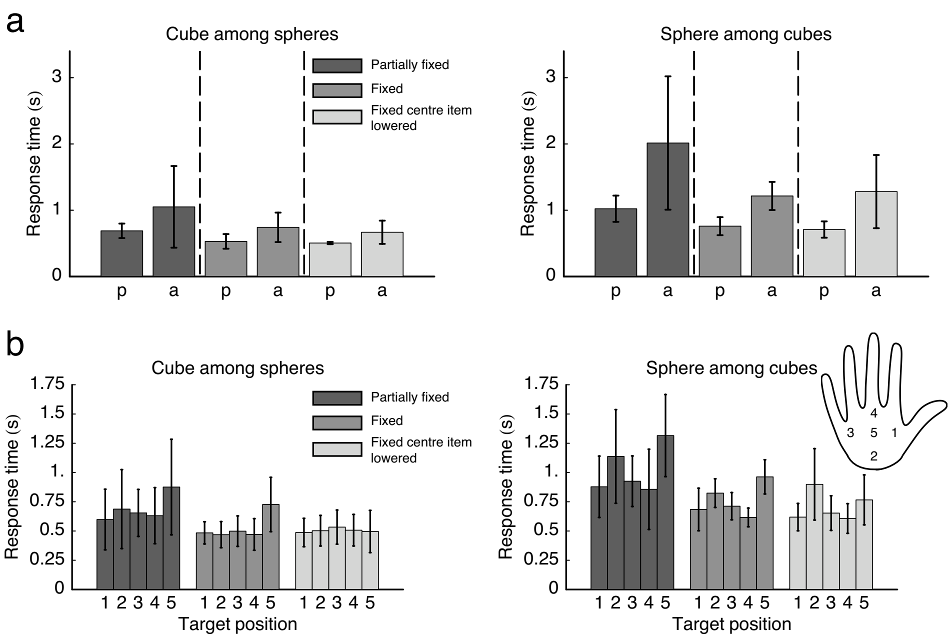

Figure 5.2: The left-hand panels show the results for search for a cube among spheres and the right-hand panels show results for search for a sphere among cubes. Error bars indicate the standard error of the mean. a) Response times averaged over subjects for target present, indicated with p, and target absent trials, indicated with a, in each of the item fixation methods. b) Response times for the different target positions.

times from each of the conditions. Sphericity was assumed according to Mauchly's test for all factors unless indicated otherwise. The ANOVA showed main effects for fixation method ($F_{2,14} = 7.0, p < 0.01$), target shape ($F_{1,7} = 46.4, p < 0.001$) and target presence ($F_{1,7} = 23.0, p < 0.01$). Furthermore, the interaction between target shape and target presence was significant ($F_{1,7} = 13.3, p < 0.01$). These results mean that response times were significantly longer for target absent trials than for target present trials. This is commonly found in search tasks. The interaction shows that the difference in response time between target absent and target present trials was larger when the target was a sphere than when it was a cube. Pairwise comparisons (post-hoc t-tests) with Bonferroni correction between the three ways of item fixation showed that the response times were significantly faster when the items were fixed ($p < 0.05$) and fixed with the centre item lower ($p < 0.05$) compared to the partly fixed conditions.

Response times for the different target positions are shown in Figure 5.2b. The target present trials were analysed further to investigate possible effects of target location. To this end, a $3 \times 2 \times 5$ (fixation method $\times$ target shape $\times$ target position) repeated measures ANOVA was performed on the response times from the target present trials. The ANOVA showed main effects for fixation method ($F_{2,14} = 12.1, p < 0.01$), target shape ($F_{1,7} = 38.9, p < 0.001$) and target position ($F_{4,28} = 11.3, p < 0.001$). There were significant interaction effects between fixation method and target position ($F_{8,56} = 4.2, p < 0.001$) and between target shape and target position ($F_{4,28} = 5.6, p < 0.01$).

A post-hoc analysis was performed on the response times from the three methods of item fixation separately. Three separate 2×5 (target shape $\times$ target position) ANOVAs showed that for each of the methods of item fixation there was a main effect of target shape and target position. Only for the fixed-with-centre-item-lower condition was there a significant interaction between target shape and target condition. Bonferroni corrected post-hoc t-tests showed for the partially fixed condition that when the target was in the centre (position 5), the response time was significantly larger than when the target was near the thumb (position 1) ($p < 0.05$). For the fixed condition post-hoc t-tests showed that response times were significantly larger when the item was in the centre (position 5) than for each of the other positions ($p < 0.01$). For the fixed-with-centre-item-lower condition the interaction between target shape and target position was significant, therefore the analysis of differences between the positions was performed on the response times for the two target shapes separately. In this case no significant differences between the target positions were found.

5.4 Discussion and Conclusion

Our results show that constraining exploratory freedom by fixating the items does not necessarily yield longer search times. Allowing small displacements and rotation of the items (partially fixed) yielded longer response times than when the items were completely fixed to their location and orientation (fixed and fixed with centre item lowered). An explanation for this could be that it is easier to exert forces on the shapes when they are rigidly fixed. In this case search could be performed by grasping and squeezing the shapes. In a previous study where the shapes were suspended from wires and therefore not fixed to a position in the hand, target present response times found for five items were 720 ms for a cube among spheres and 830 ms for a sphere among cubes (Plaisier, Bergmann Tiest, & Kappers, 2009b). These response times are roughly in the same range as for the partially fixed condition (690 ms and 1000 ms), but larger than for the completely fixed conditions in the present study (530 ms and 740 ms for fixed in the same plane and 500 ms and 700 ms for centre item lower). Response times from this previous study, however, should be only compared with caution since in that study the number of items was varied which may influence response times.

The general exploration method of grasping and squeezing the shapes simultaneously was observed for each of the experimental conditions. We chose to use spheres and cubes in this study, because it is known from a previous study that search for a cube among spheres and also the reverse can generally be performed through only grasping all of the shapes simultaneously (Plaisier, Bergmann Tiest, & Kappers, 2009b). In the same study it has been shown that in searches for other target shapes sometimes more detailed exploration is needed (e.g. when searching for an ellipsoid among spheres). In such cases, it is possible that search is faster when items can be rotated or slightly translated. Item positions relative to the hand, however, are no longer constant during a trial for such serial and more time consuming exploratory strategies.

We have also shown that there are differences in search times for the different target locations. When all items were fixed in one plane, detection of a target in the centre was slower than for any of the other target positions. However, when the centre item was lowered as little as 0.5 cm there were no differences found for the response times between the different target positions. This shows that there are no large differences in response times for different positions in the hand, as long as skin contact can be easily made with each of the items. By lowering the centre item, the global shape of a grasping hand is followed and the centre item can be easily accessed.

The finding that response times are larger for search for a sphere among cubes than for a cube among spheres is in agreement with what we have reported in a previous study (Plaisier, Bergmann Tiest, & Kappers, 2009b). In that study cubes and spheres were suspended from wires and could be freely rearranged in the hand. Here it was also found that search for a cube among spheres is faster than the reverse condition. For that way of item presentation also sliding of the shapes along each other might have played a role. It is easier the rearrange a hand full of spheres than a hand full of cubes. This could also explain why search for a cube among spheres is faster than search for a sphere among cubes. As the items were (partially) fixed to a certain position in the present study, this factor could not play a role here. Consequently, the present study clearly supports the idea that detection of edges and vertices is faster than detection of curvature, or the absence of edges and vertices. Lederman and Klatzky (1997) reported that search for the presence of an edge is faster than search for the absence of an edge. These findings support the idea that edges are highly salient features for haptic perception of shape.

An alternative explanation for search for a cube being faster than search for a sphere could involve heat conduction. Although the cube and the sphere had the same volume, they differed in surface area. The surface area of the cube is slightly larger than that of the sphere. The contact area with the skin was possibly larger for the cube resulting in more heat flow into the cube than into a sphere. However, contact area likely varied more with the position of the item in the hand than with its shape. Also skin temperature varies across different areas of the hand. Moreover, it has been shown that heat conductivity is not a very salient feature (Lederman & Klatzky, 1997). Therefore, it is much more likely that search for a cube was faster because it has edges and not because of heat conduction.

In conclusion, there are no large differences in search times for different target locations in the grasping hand as long as the skin contact can be easily made. Search for a cube among spheres is faster than search for a sphere among cubes also when the shapes are fixed to a certain spatial position or orientation. Finally, search for three-dimensional shape is faster when the shapes are rigidly fixed in space. This kind of fundamental knowledge about haptic perception can be useful in future design of especially hand worn haptic devices that allow for enclosing virtual objects with the hand.

Chapter 6

Subitizing in active touch

Published as:
Plaisier, M.A., Bergmann Tiest, W.M., and Kappers, A.M.L. One, two, three, many – Subitizing in active touch, *Acta psychologica*, 131: 163-170, 2009.

Abstract

'Subitizing' refers to rapid and accurate judgement of small numbers of items, while response times and error rates increase rapidly for larger set-sizes. Most enumeration studies have been done in vision. Enumeration studies in touch have mostly involved 'passive touch', i.e. touch without active exploration. In daily life a much more common situation is that of 'active touch', e.g. when we count the number of coins in our pocket. To investigate numerosity judgment in active touch, we let subjects haptically explore varying numbers of spheres. Our results show that enumeration for up to three items is more efficient than for larger numbers of items. We also show that enumeration in this regime was not performed through estimation. Furthermore, it is shown that numerosity information was accessed directly and not through mass or volume cues. Not only do our results show that a haptic version of subitizing exists in active touch, they also suggest similar underlying enumeration mechanisms across different modalities.

6.1 Introduction

From visual studies it is known that people judge the number of items on a display rapidly, accurately and almost effortlessly up to a certain number. This phenomenon is known as subitizing (Kaufman et al., 1949). With increasing set-sizes, enumeration becomes error-prone and response times (RT) increase markedly. Consequently, subitizing is characterised by a sharp upward bend in the slope of the RTs and error rates as a function of set-size. The location of the bend depends on the stimulus, but in vision it is generally at about four items (e.g. Atkinson et al., 1976; Mandler & Shebo, 1982). A total of three processes for numerosity judgement can be distinguished. First, for small numerosities there is the efficient and accurate process labelled 'subitizing' as described above. Secondly, for larger numerosities a more time-consuming and error-prone process is used which is referred to as 'counting'. Finally, there is the efficient process of 'estimation' for approximate numerosity judgement. The term 'enumeration' refers to numerosity judgement in general through any of these processes.

Outside the visual modality, numerosity judgment experiments in audition have also reported evidence for subitizing (Ten Hoopen & Vos, 1979; Camos & Tillmann, 2008). In this case items are usually presented sequentially and not simultaneously as is often done in vision. There has been much debate on whether two separate mechanisms are involved for subitizing and counting. Recently several brain imaging studies have focussed on this question. Pasini and Tessari (2001) suggested left hemispherical specialisation for subitizing and right hemispherical specialisation for counting. A study by Piazza, Mechelli, Price, and Butterworth (2006) reported left hemispherical specialisation for approximate numerosity judgement in both vision and audition. If a similar or even a single mechanism underlies both visual and auditory numerosity judgement, it is likely that this mechanism also extends to the haptic modality.

In touch, enumeration studies have been mostly restricted to 'passive touch', i.e. touch without active exploration. One study in which subjects had to judge the number of fingers stimulated with pins reported subitizing (Riggs et al., 2006), but no subitizing was found in a study where subjects had to report how many vibrators were distributed over the body surface (Gallace, Tan, & Spence, 2006). This raises the question whether subitizing only occurs in touch when stimuli are presented to the separate fingers. In a follow-up study where vibrators were presented to the subjects' fingers Gallace, Tan, and Spence (2008) again reported no indication for subitizing.

However, for both presentation to the separate fingers as well as presentation distributed over the body surface, error-rates were extremely high (up to 90%). These results show that stimulation of the separate fingers does not necessarily lead to subitizing, but it is also possible that vibrators are not a suitable stimulus for investigating tactile numerosity judgement. It is still not unlikely that stimulation of the separate fingers represents a special case in haptic numerosity judgement for which subitizing can occur. Recent studies have shown that there are interactions between spatial and number representations in the parietal cortex (see Hubbard, Piazza, Pinel, and Dehaene (2005) for an overview). It has been shown that areas involved in number processing partially overlap with those involved in finger movements (Pesenti, Thioux, Seron, & De Volder, 2000). Furthermore, finger motor circuits are activated during cognitive tasks such as enumeration and can facilitate cognitive processing (Andres, Seron, & Olivier, 2007; Carlson, Avraamides, Cary, & Strasberg, 2007). Since motor circuits are activated during number processing, an even more important question than that of the importance of stimulating separate fingers is whether subitizing also occurs in active touch.

In daily life, we usually explore objects through active touch. Allowing to actively explore enables the subjects to adopt the most efficient exploration strategy (Lederman & Klatzky, 1987). When exploratory movements are restricted, haptic information processing can be impaired (e.g. Lederman & Klatzky, 2004). Therefore, experiments using active or passive touch do not necessarily yield similar results. We investigated whether subitizing occurs in active touch by letting subjects enumerate varying numbers of spheres grasped in the hand. Response time and error rates were then recorded as a function of the number of items. First a numerosity judgement experiment was carried out to investigate whether two regimes exist in this type of numerosity judgement task. A second experiment was performed to investigate the role of relative discriminability between the presented numbers of items. In the last experiment the role of volume and mass cues was investigated. Finally, we introduce a model to predict response times for numerosity estimation and determine whether this model can describe the data from the second and third experiments.

6.2 General Methods

6.2.1 Subjects

Thirty paid participants (13 male, 22 ± 3 yrs) were randomly distributed over three experimental groups. One participant was left-handed and all others right-handed according to Coren's test (Coren, 1993) and none of them had any known hand deficits. All subjects were naive as to the purpose of the experiment and they were treated in accordance with the local guidelines.

6.2.2 Stimuli and set-up

The stimuli consisted of brass spheres, which were suspended from wires (Figure 8.1a). The wires were flexible enough for subjects to freely lift and rearrange the spheres in the hand. In Experiments 1 and 2 the number of spheres was varied and each sphere had a diameter of 1.86 cm and a mass of 29 g. The size of these spheres was chosen such that varying numbers of spheres could be held comfortably in one hand, while they were large

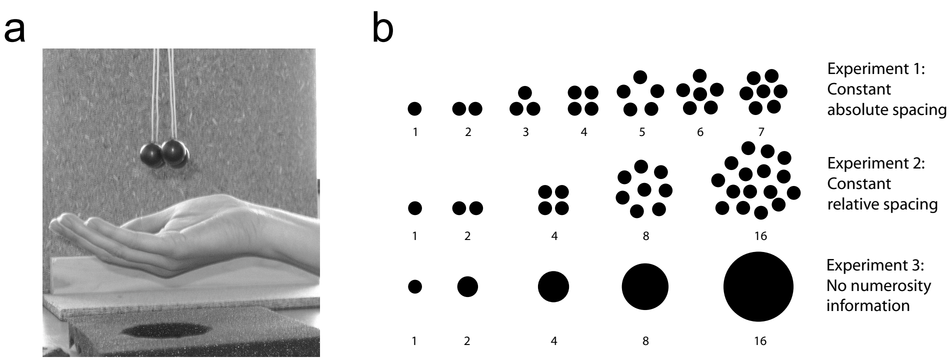

Figure 6.1: Experimental design. a) Blindfolded subjects placed their dominant hand in a holder below the stimulus and grasped upwards. Time measurement started automatically when a subject touched the stimulus. Subjects had to call out the number of spheres, or the number of equivalent small spheres, as fast as possible, terminating the measurement. b) In Experiment 1 the number of spheres increased linearly from 1 to 7, while in Experiment 2 the number of spheres doubled repeatedly from 1 to 16. In Experiment 3 subjects were presented with single spheres having masses and volumes equivalent to the 1 to 16 spheres in the previous experiment.

enough to easily resolve the individual spheres. In Experiment 3 sphere size was varied and the set included spheres with diameters of 1.86 cm, 2.34 cm, 2.95 cm, 3.72 cm and 4.69 cm. Response times were measured using a custom built device. Time measurement was started automatically when a subject touched the stimuli and it was terminated through a vocal response. The RTs were measured with an accuracy of 10 ms. For technical details about this device, see Plaisier et al. (2008b).

6.2.3 Design

Subjects were blindfolded and wore earplugs to eliminate sound cues. They were instructed to grasp the stimuli with their dominant hand and respond the correct number of spheres (Experiments 1 and 2) or sphere size (Experiment 3) as fast as possible. There were no restrictions on exploration strategy nor on hand movements, other than having to initially grasp all items simultaneously. After grasping all items they were allowed to release spheres from their hand during a trial. Whenever an incorrect response was given, the experimenter informed the subjects of what the response should have been and the response time was discarded. Before the experiment was started, a minimum of 20 practice trials was performed until 10 in a row were correct. The number of practice trials never exceeded 30 trials. Each subject performed 25 trials per number of spheres or sphere size. There was a 5 minute break after 50 minutes and none of the experiments took more than 90 minutes.

6.2.4 Analysis

Strategy shifts in cognitive tasks can be detected by regression of a model consisting of multiple linear parts with unknown change points to, for instance, response times (Luwel, Beem, Onghena, & Verschaffel, 2001). To determine the location of a possible transition point in our data, a bilinear function for the response time T as a function of the number of items N consisting of two linear regimes with a discrete transition was used:

$$T(N) = (r_1 N + c_1) H\left(\frac{c_2 - c_1}{r_1 - r_2} - N\right) + (r_2 N + c_2) H\left(N - \frac{c_2 - c_1}{r_1 - r_2}\right) \quad (6.1)$$

Here $H(N)$ is the Heaviside step function and r_1 and r_2 are the slopes, while c_1 and c_2 represent constant offsets. Regression of this function was performed on the response times averaged over subjects from Experiment

1. We also checked whether a bilinear function described the data better than two other models: a bilinear function with slope $r_1 = 0$ and a linear function.

Note that a linear function has 2 free parameters, while the bilinear function has 4 free parameters and the bilinear function with the first slope set to zero has 3 free parameters. Regression of a function with more free parameters is more likely to yield a larger R^2 value, so we cannot compare the performance of these models by looking at the R^2 values only. To determine how well the three functions described the data while taking into account the differences in the numbers of free parameters, the Akaike information criterion (AIC) was used (Akaike, 1974). With this method the best function can be selected by looking at how much information is lost when a certain function is used to describe the data. The model for which this information loss is the smallest is then selected as the best model. This calculation is based on the sums of squares (SS), the number of data points (n) and the number of free parameters (k). Here we used AIC with small sample-size correction (AIC_c), which is defined as:

$$\mathrm{AIC}_c = n \ln\left(\frac{\mathrm{SS}}{n}\right) + 2k + \frac{2k(k+1)}{n-k-1} \tag{6.2}$$

From these AIC_c values Akaike weights (w_a) can be calculated. These w_a represent the relative probability that a certain function in a set of M functions describes the data best. The function with the w_a closest to unity is then determined to perform best. These values are obtained through the following equation:

$$w_a^i = \frac{e^{-\frac{1}{2}\mathrm{AIC}_c^i}}{\sum_{m=1}^{M} e^{-\frac{1}{2}\mathrm{AIC}_c^m}} \tag{6.3}$$

6.3 Experiment 1: Constant absolute spacing

This experiment was a haptic version of a 'classic' visual numerosity judgment experiment in which subjects are presented with varying numbers of items and have to enumerate them. If haptic numerosity judgement is similar to that in vision we expect to find a sharp upward bend in both response times and error rates.

6.3.1 Methods

Subjects were presented with 1, 2, 3, 4, 5, 6 or 7 items (Figure 8.1b). Numbers larger than 7 were not presented, because all subjects had to be able to hold all the spheres comfortable in the hand for the whole numerosity range. They also had to be able to move the spheres in the hand without spheres falling out of the hand unintentionally. Instructions were as described in the General Methods.

6.3.2 Results

The averaged RTs as well as the error rates are shown in Figure 6.2. Note that no incorrect responses were given up to 3 items and that the error rates increased for larger numbers of items. If errors were made the presented number of spheres was always confused only by numbers one sphere less or one sphere more. The low error-rates indicate that subjects were able to perform the task very well and they did not systematically under- or over-estimate the number of spheres. The response times seem to show two regimes. To test this, regression of a linear function, a bilinear function and a bilinear function with the first slope set to zero was performed on the averaged response times. Data points were weighted according to their standard error. Regression yielded $R^2 = 0.78$ for the linear model, $R^2 = 0.99$ for the bilinear model and $R^2 = 0.99$ for the bilinear model with the first slope set to zero. As pointed out before, we cannot conclude from these R^2 values that the bilinear functions perform best, because the functions differ in the number of free parameters. To take the different numbers of free parameters into account, Akaike weights were calculated. This yielded $w_a^1 = 0.97$ for the bilinear function, $w_a^2 = 0.03$ for the bilinear function with the first slope set to zero and $w_a^3 = 0.0003$ for the linear function. This shows that the bilinear function describes the data best with a probability close to one.

Regression of the bilinear function to the response times averaged over subjects yielded a slope transition at 2.9 ± 0.3 (s.e.m.) items. The reported standard error is the value that follows from the regression procedure. The slope of the first linear regime was determined at 167 ms/item, while for the second regime the slope was 839 ms/item. To check how representative the values from the fit to the average data are, the single subjects' data were also analysed by regression of the bilinear function. The location of the transition point, for instance, may vary among subjects. From the fits to the single subjects data it followed that on average the transition point was located

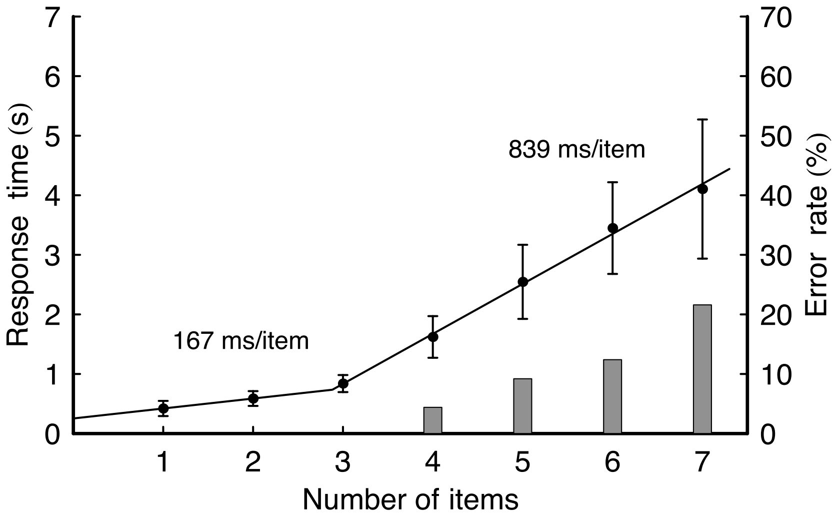

Figure 6.2: Experiment 1: Constant absolute spacing. Response times (dots ± s.e.m.) and error rates (grey bars) as a function of the number of items averaged over subjects. Note that no errors were made below 4 items. The solid line represents regression of the bilinear function and the slope values are indicated in the figure.

at 3.4 ± 0.8 items. Here, the reported confidence interval is the standard error of the values found for the different subjects. The slope values for r_1 and r_2 were 170 ± 30 ms/item and 976 ± 70 ms/item, respectively. The fitted slopes of the first linear part were significantly smaller than for the second part (paired-samples t-test on the slopes of the individual subjects' RTs, $t_9 = 13$, $P < 0.001$). In Figure 6.2 it can be seen that only two data points lie in the first regime. Note that as we did not force the transition point from the first regime to the second in between 2 and 3 items. Instead, we minimised the total sum of squares of the model as a whole end not the two linear parts separately. Therefore, the first slope is not simply a connection of the first two data points but the result of a fitting procedure. Furthermore, regression of this model to the single subjects' data, showed that the location of transition point varied between subjects and was often in between 3 and 4 items. The average value of the slope in the first regime for the single subject fits is almost the same as the slope that resulted from the fit to the response times averaged over subjects. This shows that the value of the first slope found through this procedure is robust.

In some studies the largest number of items (in this case 7 items) is not taken into the analysis because of possible endpoint effects. It is common in numerosity judgment data that for the highest numerosity, response times are lower than would be expected from the trend in the previous data points (e.g. Trick & Pylyshyn, 1993; Watson, Maylor, & Bruce, 2007). Although we did not explicitly inform the subjects of what was the largest numerosity, it is likely they noticed it during the experiment. However, the response time for 7 items did not significantly deviate from the trend in the previous three points[1] and the regression procedure yielded similar parameters if the response time for 7 items was excluded.

6.3.3 Discussion

Our results show that the response times as a function of the number of items show a sharp upward bend at about 3 items. This relation is better described using a bilinear function than a single linear function, or a bilinear model with the first slope set to zero. This shows that the subitizing slope is not zero. This is in agreement with visual numerosity judgement experiments, in which the subitizing slope is also larger than zero (e.g. Akin & Chase, 1978; Oyama, Kikuchi, & Ichihara, 1981; Trick & Pylyshyn, 1993; Trick, 2008).

Analysis of the single subject data showed a transition point on average between 3 and 4 items and the slopes for the first part of the bilinear function were significantly smaller than those of the second part. Together with the error rates, this provides strong evidence that there are two regimes when making numerosity judgements using active touch. Since there was an upward bend in the response times and also error rates increased after the bending point, it can be assumed that subjects were *counting* for numerosities larger than three items. For small numerosities (<4) subjects used a more efficient and accurate enumeration strategy. Note that we chose the sphere size such that up to 7 could be comfortably held in the hand so the spheres not fitting in the hand was no reason for a change in performance after 3 spheres. A possible explanation is that subjects used an estimation strategy to judge small numerosities very efficiently. Note that while the absolute differences between subsequent numerosities were constant, the *relative* differences in numerosity, mass and volume decreased with the number of items, making discrimination progressively harder. A

[1] The response time for 7 items deviated 0.2 s.e.m. from a linear fit through the data for 4, 5 and 6 items

second experiment was performed to investigate the role of relative discriminability on numerosity judgement.

6.4 Experiment 2: Constant relative spacing

In vision, discriminability differences have been suggested as an explanation for the existence of two regimes (Van Oeffelen & Vos, 1982b). Furthermore, Ross (2003) suggested that subitizing might be explained from a Weber fraction for visual discrimination of numerosities. When the relative difference between two subsequent numerosities is smaller than the Weber fraction, subjects might switch from one enumeration strategy to another.

To investigate the influence of relative discriminabilities between the presented numerosities, relative differences between subsequent numerosities were kept constant in this experiment. Each numerosity differed by a factor of two from the previous numerosity in the range. A factor of two was the largest difference between two subsequent numerosities that was present in Experiment 1. If indeed the transition from efficient enumeration to counting that was found in Experiment 1 was caused by the decreasing relative differences between the presented numerosities, we hypothesise that subjects will use the same enumeration strategy over the whole numerosity range in this experiment. Performance should then be roughly constant over the whole numerosity range in the present experiment.

6.4.1 Methods

Stimuli consisted of 1, 2, 4, 8 or 16 spheres (Figure 8.1b), so that the relative volume and mass differences between the subsequent numbers of spheres were constant. In this case it was no problem to present more than 7 items, in contrast to Experiment 1, because pilot experiments showed that subjects used the same exploratory strategy for the whole numerosity range. They only grasped the spheres and did not need to make exploratory movements. Before the experiment started subjects were informed of which numbers of items could be presented.

6.4.2 Results

Figure 6.3a shows that the response times for up to 16 items were all below 1 s and error rates were low for all numerosities. This indicates that enumeration was facilitated by increasing the relative mass and volume differences.

The effect of numerosity was significant (repeated measures ANOVA, $F_{4,36} = 58$, $P < 0.001$) and Bonferroni corrected pair-wise comparisons showed that enumeration was significantly faster for up to 4 items than for 8 and 16 items. So, RTs were still significantly lower in the first part of the stimulus range. For comparison, Figure 6.3b shows the results from this experiment together with the response times from Experiment 1 (in grey). It can be seen that for 1 and 2 items response times are comparable between the experiments, but for larger numbers the response times from experiment 1 increase rapidly.

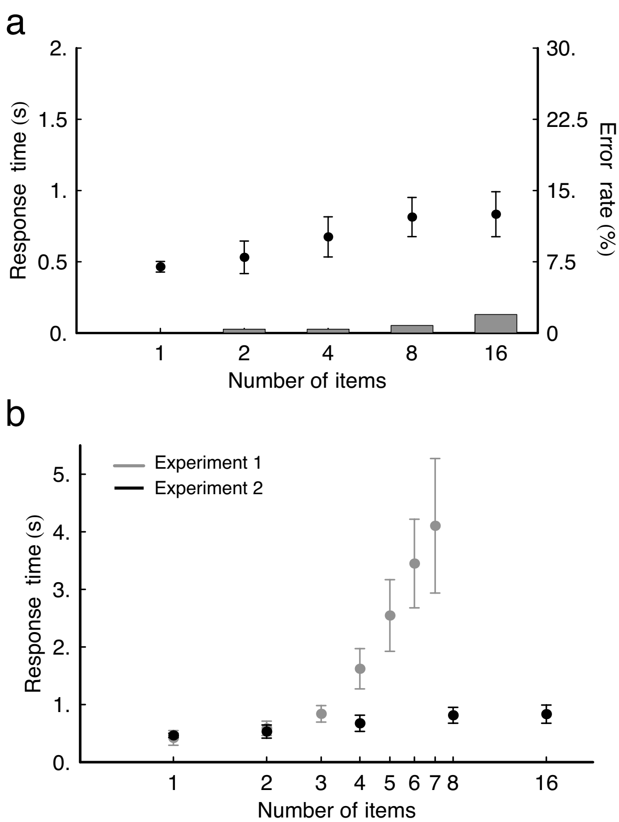

Figure 6.3: Experiment 2: Constant relative spacing. a) Response times (dots ± s.e.m.) and error rates (grey bars) as a function of the number of items averaged over subjects. b) The response times from Experiment 2 plotted together with the response times from Experiment 1(in grey).

6.4.3 Discussion

The results clearly show that when the large numbers of items were spaced further apart, enumeration of large numbers of items (8 and 16) was facilitated markedly. For numerosities larger than 2 the response times from Experiment 1 increased much more rapidly than in the present experiment. This indicates that an estimation process was used to judge these numerosities and subjects did not count the individual items. However, RTs were not constant over the different numerosities. The significantly lower response times for small numbers of items (≤ 4) suggest that a still more efficient enumeration process was applied for small numbers of items. This result suggests that subjects did not use the same enumeration process over the whole numerosity range contradicting the hypothesis that smaller relative differences between the numbers of items at the low end of the stimulus range is the cause of the two regimes found in Experiment 1. Possibly, subjects used volume and mass cues to estimate numerosity for the large numbers of items, but used a different, more efficient, strategy for small numbers of items in the present experiment. This strategy should then not be based on mass and volume cues, but on numerosity information. Experiment 3 was carried out to investigate the role of mass and volume estimation.

6.5 Experiment 3: No numerosity information

To investigate the role of mass and volume cues, subjects were deprived of numerosity information and only mass and volume cues were available. If the efficient performance for small numerosities, or rather small masses and volumes, remains, then it would be an effect of mass and volume estimation. If, on the other hand, the efficient performance disappears, numerosity information would be accessed directly and not through volume and weight estimation.

6.5.1 Methods

In this experiment stimuli consisted of single spheres having volumes and weights equivalent to the varying numbers of spheres in the progressive-spacing experiment. The spheres were labelled 1, 2, 4, 8 and 16 equivalent to the numbers of spheres in the previous experiment (see Figure 8.1b) and subjects were instructed to respond with the correct label. With the removal of numerosity information, this task has in fact become a categorisation task.

6.5.2 Results

The results for this experiment are shown in Figure 6.4. RTs for all numbers of items except 16 were significantly higher than in Experiment 2 ($t_{18} \geq 2.4, P \leq 0.034$) and particularly for 1 and 2 items did the error rates increase substantially. In error trials the presented sphere was always confused with the previous or the next sphere size in the range (e.g. sphere 8 was confused with either sphere 4 or 16 and sphere size 2 was confused with either 1 or 4). The effect of sphere size was significant (repeated measures ANOVA, $F(4, 36) = 7.9, p < 0.001$). This can be seen from Figure 6.4, in which the response times at both ends of the stimulus range are smaller.

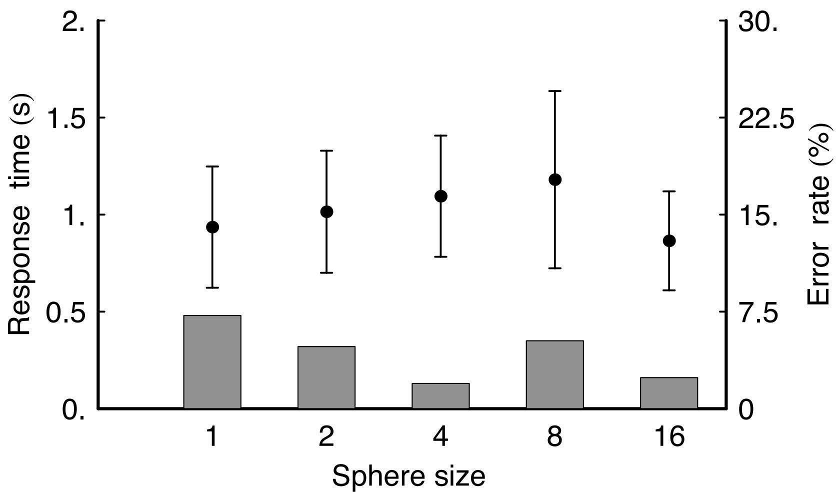

Figure 6.4: Experiment 3: No numerosity information. Response times (dots $\pm$ s.e.m.) and error rates (grey bars) as a function of sphere size averaged over subjects.

6.5.3 Discussion

The results show that in the absence of numerosity information, performance in the first part of the stimulus range deteriorated and the highest error rate was found for the sphere size which was equivalent to 1 item in Experiment 2. The more efficient performance for small numerosities found in Experiment 2 disappeared. Although there was no special regime found in the present experiment, performance was not completely constant over the whole stimulus range. A model based on relative discriminability was designed to describe the expected pattern in the response times for numerosity judgement through an estimation process.

6.6 Estimation model

If subjects used an estimation strategy, we would expect discriminability between the different stimuli in the set to play a role. Discriminability can be assumed to be proportional to the perceived difference between two stimuli. Perceived difference in magnitude usually obeys Fechner's law. According to Fechner's law the perceived difference between two quantities scales with the logarithm of the ratio of these quantities. Therefore, it can be assumed that the discriminability d between quantities x_1 and x_2 is proportional to:

$$d(x_1, x_2) \propto \left| \log \frac{x_1}{x_2} \right| \tag{6.4}$$

In Experiments 2 and 3, the subjects knew which numerosities could be presented. If they used an estimation strategy, it is expected that when judging a stimulus they considered each of the possibilities weighted according to their discriminability. The total response time is then expected to be inversely related to the sum of the discriminabilities. This is then given by:

$$\mathrm{RT}(N) = a + \frac{b}{\sum_{n=i}^{j} |\log \frac{N}{n}|} \tag{6.5}$$

where N is the numerosity that is presented, n is an iterator which runs from the smallest numerosity in the set (i) to the largest one (j) over all numerosities in the set. Free parameters a and b scale the offset and width of the function.

This model was fitted to the response times found in both Experiment 2 and Experiment 3 using non-linear regression. Each data point was weighted according to the corresponding standard error. Figure 6.5 shows the best fit for both sets of response times. The data from Experiment 3 are described fairly well by this model ($R^2 = 0.77$). In this case the model also performs much better than a straight line ($R^2 < 0.01$). On the other hand, the model

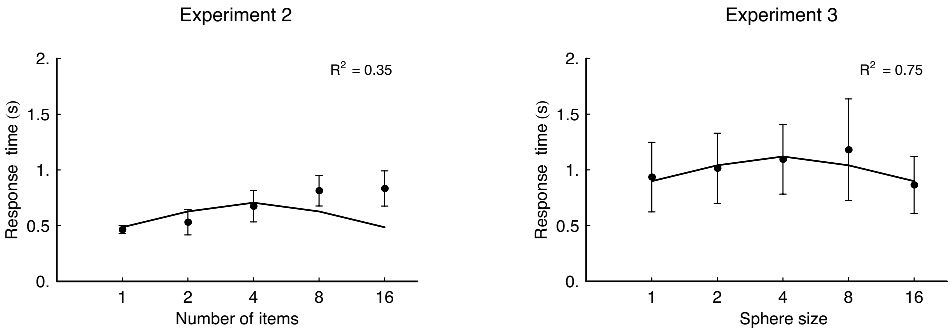

Figure 6.5: Best fit of the estimation model to the response times averaged over subjects for both Experiment 2 with numerosity information present and Experiment 3 without numerosity information. R^2 values are indicated in the figure. It is clear that the estimation model does not fit the data from Experiment 2. This indicates that a different enumeration mechanism is used when numerosity information is present.

is clearly not suitable to describe the data from Experiment 2 ($R^2 = 0.34$). In fact, it performs worse than a straight line ($R^2 = 0.86$). Note that in this case we can directly compare R^2 values between a straight line and our model, because both have the same number of free parameters.

6.6.1 Discussion

The estimation model based on discriminability describes the pattern in the response times when there was no numerosity information fairly well. Note also that this model predicts end effects, i.e. response times at the ends of the stimulus range are smaller. This is often found in numerosity judgement experiments. However, this model cannot describe the data from Experiment 2 in which numerosity information was present. Therefore, the conclusion has to be that other processes than just mass and volume estimation play a role in this experiment [2]. Response times for the largest numerosities in Experiment 2 were similar to those for the largest spheres

[2] Note that weighted regression of the model was used, which means the response times were weighted according to their standard error. Since the standard error for small numerosities was smaller than for the larger numerosities, the distance between the response times and the fit for Experiment 2 is smallest for the first part of the stimulus range. However, this does not mean that the model can describe the data for small numerosities, but not for the larger numerosities.

in Experiment 3. Also, there seems to be an endpoint effect for the largest numerosity in Experiment 2 similar to the end effect for the largest sphere in Experiment 3. This suggests that for the largest numerosities (8 and 16) in Experiment 2 subjects used an estimation process like in Experiment 3. For the smallest numerosity in Experiment 2 there is no end-effect like for the smallest sphere in Experiment 3 indicating that indeed in this regime a different enumeration process was used. The estimation model assumes the same estimation process for all stimuli in the set. Since this model cannot be used to describe the pattern in the response times from Experiment 2, this suggests that subjects used different enumeration processes for small and large numerosities in that experiment. Note that in line with this interpretation, the estimation model can not describe the data from Experiment 1 either. Also in this case different mechanisms are used over different parts of the range of numerosities as there is a transition from subitizing to counting.

6.7 General Discussion

In the visual domain there is still debate on what type of process subitizing actually is. Several explanations have been proposed for this efficient enumeration of small numbers of items. Mandler and Shebo (1982) showed that presenting dots in canonical patterns can increase the subitizing range from 4 to 6 items. Pattern based explanations of subitizing suggest that for small numbers of items pattern recognition plays a role. Similarity between different numbers of items is sometimes linked to this pattern based explanation. Logan and Zbrodoff (2003) showed that when subjects had to rate the similarity between different numbers of items, they rated small numbers of items very dissimilar to each other whereas larger numbers of items tended to be rated increasingly similar. Visual numerosity discrimination experiments have shown that there is a constant Weber fraction of 25% for number discrimination (Ross, 2003). This could explain why there is a transition from subitizing to counting at about 4 items, because this is where the relative difference between two subsequent numerosities drops below the 25% discrimination threshold. Recently it has been shown that this explanation does not hold (Revkin et al., 2008). In the present study we found a more efficient enumeration regime for small numerosities when the relative differences between subsequent numerosities was always a factor of two (Experiment 2). This indicates that also in haptic numerosity judgement, the efficient enumeration of small numbers of items is not caused by large relative differences between the presented numerosities. Such a regime of efficient performance was not present when numerosity information was absent (Experiment 3). In that case a model for predicting response times based on Fechner's law described the response times very well, while this model could not describe the response times from Experiment 2. This shows that numerosity information was accessed directly for small numerosities and not through mass or volume estimation.

Subitizing is sometimes assumed to be a purely parallel process yielding response times that are independent of the number of items. This is, however, not generally found in visual studies. In visual numerosity judgement studies where subjects were not forced to answer within a certain time interval, the slope of the response times as a function of the number of items in the subitizing regime is generally 40 to 100 ms/item, while the counting slope is between 250-350 ms/item (e.g. Akin & Chase, 1978; Oyama et al., 1981; Trick & Pylyshyn, 1993; Trick, 2008). In visual search studies, a continuous range of response time slopes is found from purely parallel to purely serial search (Wolfe, 1998). This is generally explained using search models that assume that visual searches are performed through a combination of parallel and serial processes (e.g. Duncan & Humphreys, 1989; Cave & Wolfe, 1990; Theeuwes, 1993). If we extrapolate this idea to numerosity judgement, then the subitizing slope is expected near the parallel end of the range and the counting slope on the serial end of the range of slopes. In a previous study we have measured a range of search slopes in the haptic modality using the same set-up as in the present study (Plaisier, Bergmann Tiest, & Kappers, 2009b). Like in visual search, also a large range of slope values was found. These ranged from highly efficient search for a tetrahedron among spheres (88 ms/item, target absent), to very inefficient for an ellipsoid among spheres (1200 ms/item, target absent). Because in numerosity judgement always all items have to be processed, comparison to target absent search slopes is appropriate here. The subitizing slope we found in Experiment 1 is somewhat larger than the visual subitizing slopes (167 ms/item), but it is near the parallel end of the range of haptic search slopes. The counting slope found in Experiment 1 (839 ms/item) is on the serial end of the range of search slopes. So, although the haptic subitizing and counting slopes found in the present study may be larger than generally found in vision, they are in agreement with the values one would expect from the range of slopes found in haptic search. Note also that our subitizing slope is substantially smaller than the tactile subitizing slope reported by Riggs et al. (2006) (270 ms/item). The fact that the slope for small numerosities is larger for the haptic modality than in vision does not rule out that similar mechanisms underlie numerosity judgements in this regime, especially since in search tasks response time slopes are generally larger in the haptic modality than in vision.

Information extraction is quite different between these two modalities and before items can be enumerated through any mechanism they have to be individuated. It is not unlikely that this process is less efficient for the haptic modality. Besides being fast the subitizing mechanism is also characterised by accuracy. Regardless of the exact subitizing slope value, our data shows that enumeration of small numerosities is much faster and more accurate than for larger numerosities. Furthermore, we have shown that this enumeration process is not the same as estimation. Therefore, we conclude that for small numbers of items (< 4), regardless of relative differences between presented numerosities, an efficient mechanism analogous to the visual subitizing process is used. For larger numbers of items, either counting or estimation is used, depending on the size of the relative differences between the numbers of items.

An approach to the understanding of visual subitizing that is not based on discriminability, involves FINSTs (Fingers of Instantiation). This explanation is based on the Visual Indexing theory which is the idea that humans have a way to refer to a certain item without having to link it to specific features of the item such as position (see Pylyshyn (2001) for a review). From multiple object tracking experiments it is known that subjects can track up to five items simultaneously (Pylyshyn & Storm, 1988). This leads to the hypothesis that there are five FINSTs that allow for up to five items to be tracked in parallel. It has been proposed that this also explains why small numerosities (< 5) are enumerated more efficiently than larger numerosities (Trick & Pylyshyn, 1993, 1994). Visual experiments involving moving items have shown that also in that case subitizing can occur, even with the addition of distractor items (Alston & Humphreys, 2004). Note that in our haptic experiment the item positions were also not fixed. However, in this case the items could be physically manipulated by the subject and moved as the result of the subject's own action. This is of course never the case in visual experiments and physical item manipulation is specific for the haptic modality. It is possible that a haptic version of visual indexing exists or that visual indexing can also be used to process information that is extracted haptically.

Riggs et al. (2006) have reported evidence for the existence of an accurate and fast regime in haptic numerosity judgement for small numerosities

(< 4) when separate fingers were stimulated with pins. This is in agreement with the present study were we also find that there is evidence for two regimes in numerosity judgement. Gallace et al. (2006) have shown that there is no evidence for the existence of two regimes when vibrators distributed over the body surface had to be enumerated. In another study they have shown that also when vibrators were placed on the separate fingers results were quite similar to when they were distributed over the body surface and again they did not find an indication for two regimes in numerosity judgement (Gallace et al., 2008). However, for both modes of presentation (to the fingers and distributed over the body surface) error rates were very high (up to 90%). Already for a numerosity as small as two the error rate was 40 % when the vibrators were presented to the fingers and from 4 items or larger the error rates are at chance level. This indicates that subjects could not assess the presented numerosities accurately in any part of the numerosity range. This suggest that perhaps vibrators are not suitable stimuli for this type of task. It is also possible that the presentation time (100 ms) was too short. Limiting presentation time is generally not desirable in numerosity judgement studies, since it may actually influence the enumeration process (e.g. Jensen, Reese, & Reese, 1950; Trick, 2008). Short presentation time may force subjects, for instance, to use a faster but less accurate estimation process. Whatever the reason, the large error rates in the Gallace et al. study show that subjects could not enumerate the items correctly. Therefore, analysis of the accompanying response times does not provide insight into the enumeration process and it is not clear what should be concluded from this study.

The existence of two separate mechanisms (subitizing and counting) in visual numerosity judgement has been disputed. Balakrishnan and Ashby (1991, 1992) compared different models and showed that they could not find evidence for the existence of a discontinuity in numerosity judgement of small and larger numbers. They concluded that enumeration was a continuous process and the upward bend is caused by an increasing cognitive load. One problem with this study is that they limited presentation time, which may influence the enumeration process as pointed out earlier. A more important point is that, for instance, an exponential function may approximate the shape of numerosity judgement data quite well. It is however not likely that response times will continue to increase exponentially with the number of items. Such a function can thus only be a good fit in the regime where it approaches a function with two linear parts. Therefore, finding a continuous function that fits the data does not rule out that there are actually two distinct underlying enumeration processes just as well as finding a good fit for a bilinear function does not necessarily mean that there are two distinct processes. Therefore, in Experiment 1 we compared performance of a bilinear function and a linear function to determine whether there was evidence for the existence of two regimes in the data. However, from this data alone it is not possible to conclude that there are two distinct enumeration processes. For drawing conclusions about whether separate enumeration mechanisms underlie performance in these regimes, one could manipulate the numerosity range and look at how performance changes. In vision this has, for instance, been done by Revkin et al. (2008). Therefore, we performed Experiments 2 and 3, which showed that when numerosity information is present, an enumeration process is used for small numerosities which is more efficient and accurate than estimation. This suggests that there exists an enumeration mechanism that is different from both counting and estimation in haptic numerosity judgement, similar to subitizing in vision.

Summarising, our results show that in haptic numerosity judgement there is evidence for the existence of two regimes in terms of response times and error rates. For small numerosities (< 4) enumeration is more efficient and accurate than for larger numbers of items. Furthermore, we have shown that this efficient and accurate performance only occurs when numerosity information is present and does not depend on the relative differences between the numerosities in the range. It is unclear whether the haptic enumeration process is the same as the process underlying visual enumeration. Nonetheless, the pattern in the data is similar to that found in visual numerosity judgement studies. We therefore propose that the efficient enumeration of small numerosities in haptic numerosity judgement can be labelled haptic subitizing in analogy to the visual effect.

Chapter 7

Similar processing in visual and haptic numerosity judgment

Submitted as:
Plaisier, M.A., Bergmann Tiest, W.M., and Kappers, A.M.L. Similar processing in visual and haptic numerosity judgment.

Abstract

'Subitizing' refers to fast and accurate judgement of small numerosities, whereas for larger numerosities either counting or estimation are used. Counting is slow and precise, whereas estimation is fast but imprecise. In this study consisting of 5 experiments we investigated if and how the enumeration process is affected by the relative spacing between the presented numerosities. To this end we let subjects enumerate the number of dots presented on a screen and recorded their response times. Our results show that subjects switch from counting to estimation if the relative differences between subsequent numerosities are large (a factor of 2), but that enumeration in the subitizing range was still faster. We also show this fast performance for small stimuli only occurred when numerosity information is present. This indicates this is typical for number processing and not magnitude estimation in general. Furthermore, comparison with a previous haptic study suggests similar processing in numerosity judgment through haptics and vision.

7.1 Introduction

In visual numerosity judgment, three different enumeration processes can be identified. Small numerosities (≤4) are enumerated fast and error-free through a process that has been labeled 'subitizing' (e.g. Kaufman et al., 1949; Atkinson et al., 1976; Mandler & Shebo, 1982; Trick & Pylyshyn, 1993). The slope of the response times as a function of the number of items in this regime is generally found to be 40 – 100 ms/item (e.g. Akin & Chase, 1978; Oyama et al., 1981; Trick & Pylyshyn, 1993; Trick, 2008). For larger numerosities (> 4) the slower and more error-prone process of 'counting' is used and response times and error rates increase rapidly with the number of items. The slopes of the response times are usually 200 – 400 ms/item in this regime. Humans adults, but also infants and animals, can also judge approximate numbers without counting (e.g. Beran, Taglialatela, Flemming, James, & Washburn, 2006; Whalen, Gallistel, & Gelman, 1999; Dehaene, Dehaene-Lambertz, & Cohen, 1998). This last enumeration process is fast and will be referred to as 'estimation'. Numerosity judgments through estimation become less precise for increasing numerosities and obey Weber's law stating that the accuracy is a constant fraction of the magnitude. Therefore, discriminability of two numerosities is defined by their ratio (Izard & Dehaene, 2008; Gallistel & Gelman, 2000, 1992). It has been suggested that numerosity judgement without counting relies on a mapping between an internal continuous representation of magnitude onto Arabic numerals or number words (Whalen et al., 1999; Moyer & Landauer, 1967). This mapping has some variability as magnitude representations are retrieved from memory. Recently, it has been shown that this mapping can be recalibrated by providing feedback after each numerosity judgment (Izard & Dehaene, 2008).

The question of what kind of a process subitizing actually is has yet to be answered. It has been suggested that it is not a separate process at all. Balakrishnan and Ashby (1992) have suggested that there is no evidence for the existence of a subitizing regime. Others have argued that subitizing is caused by large relative differences between small numerosities (Van Oeffelen & Vos, 1982b). For instance, the relative difference between 2 and 3 is much larger than between 6 and 7. It has been shown that there is a 25 % Weber fraction for the discrimination of large numerosities (8 – 64 items) (Ross, 2003). This would explain a transition to counting above four items, because then the relative difference between subsequent numerosities becomes smaller than the discrimination threshold. Recently, it has been shown that the hypothesis that subitizing is very accurate estimation does not hold (Revkin et al., 2008). In that study, the authors compared enumeration of 1, 2, 3, 4, 5, 6, 7 or 8 items to enumeration of 10, 20, 30, 40, 50, 60, 70 or 80 items. Note that the relative differences between subsequent numerosities were the same for both numerosity ranges. By limiting the response time, subjects were prevented from counting the items. For the first range they found that enumeration of 1 to 4 items was faster and more accurate than for the larger numerosities. In contrast, for the second range there was no clear advantage for numerosities 10 to 40 compared to 50 to 80. This suggests that subitizing is not a Weberian estimation process.

Cordes, Gelman, Gallistel, and Whalen (2001), however, did not find such a discrepancy between subitizing and counting range in a study where subjects were shown a numeral and had to make the corresponding number of key presses with verbal and non-verbal counting. In the verbal counting condition, subjects counted the number of key presses out loud, while in the non-verbal condition they had to say “the" with every key press. The coefficient of variation (ratio between the mean response and the standard deviation) was constant over the whole range in both conditions indicating that there was no special performance for small numbers. This suggests that small numbers are represented in the same way as larger numbers, which contrasts the study by Revkin et al. (2008). This could be due to the fact that in the Revkin et al. study, numbers were represented by a collection of dots, while in the Cordes et al. study numerals were used. Subitizing may only be relevant for processing sets of items. Furthermore, it has been shown that numerosities from the subitizing range are rated as more dissimilar than numerosities outside this range (Logan & Zbrodoff, 2003). These findings clearly show that when dots scattered over a display are shown, for some reason numerosities from the subitizing range are recognized faster and more accurately than larger numerosities.

Although it is not clear what causes the fast and accurate enumeration of small sets of dots, it has been shown that subitizing is not limited to visual numerosity judgement. Subitizing has been shown to occur for up to two items in audition (Ten Hoopen & Vos, 1979; Camos & Tillmann, 2008). Note however, that in this case items are often presented sequentially instead of simultaneously. More recently, subitizing has also been shown to exist in haptic numerosity judgement for both ‘passive touch’ (i.e. touch without active exploration) (Riggs et al., 2006), as well as ‘active touch’ (Plaisier, Bergmann Tiest, & Kappers, 2009a). In this last study, we have addressed the role of the relative differences between subsequent numerosities in the numerosity range. Subjects had to grasp and enumerate 1, 2, 4, 8 or 16 spheres. Note that there was always a factor of 2 between subsequent numerosities. In this case, we found that enumeration was fast for all numerosities, but enumeration of small numerosities (≤ 4) was even faster than for larger numerosities. We compared response times and error rates from this task to a different task in which subjects had to label single spheres varying in size. In this case no clear advantage for small sphere sizes was found. The response times from this second task could be described using a model based on Fechner's law for discriminability. This showed that discriminability followed the psychophysical power law over the whole range of sphere sizes. However, this model could not describe the pattern in the response times from the first task in which numerosity was varied. This suggests that although the relative differences between subsequent numerosities were constant over the whole range, small numbers were recognized faster and more accurately than large numbers. Furthermore, this fast recognition was not mediated through the use of volume or mass cues.

In short, our haptic study showed that numerosity judgement without counting was faster for numbers from the subitizing range than outside this range, even when the relative spacing between numerosities was a factor of two over the whole range and feedback was provided so subjects could re-calibrate their number mapping. These results are in agreement with the study of Revkin et al, suggesting that subitizing is not the same process as estimation of large numbers. Based on this hypothesis, the results from our haptic study should be reproducible in the visual domain. Note that this approach is different from the one Revkin et al. (2008) used. In their study relative differences between subsequent numerosities varied over the stimulus range and subjects were forced to use estimation by limiting response times. Our approach is to make the relative differences between subsequent numerosities constant and larger than the discrimination threshold over the whole range. Therefore, subjects would be able to accurately judge the numerosity without counting over the whole range and will use estimation without being forced to do so. If our haptic data is reproducible in the visual domain, this is further support for the idea that numerosities from the subitizing range are recognized faster than outside this range and that this is not due to the mapping of numbers being increasingly less precise for larger numerosities. Moreover, it would argue for a shared representation of number between the visual and the haptic modalities. This has interesting consequences for the possible mechanisms underlying fast recognition of numbers in the subitizing regime as typical visual explanations, such as pattern recognition would in that case be very unlikely.

In Experiment 1, a 'classic' numerosity judgement task was performed in which we reproduce the well-known upward bend in the response times at about 4 items. To investigate what the effect was of decreasing relative differences for larger numerosities in Experiment 1, Experiment 2 was performed. Here, we presented subjects with numerosities that were chosen such that the relative difference between subsequent numerosities was constant over the whole range (1, 2, 4, 8, 16 or 32 items). Note that in this case relative differences between subsequent numerosities were larger than the discrimination threshold of 25% for judging number without counting. If subitizing were accurate estimation made possible because relative differences are above the discrimination threshold, we would not expect faster performance for small numerosities than for large numerosities. In the next experiment we investigated how response times scale with magnitude in the absence of numerosity information. To this end, numerosity information was removed in Experiment 3 and subjects had to name dots with varying sizes. In this case one could expect response times to be constant over the whole range. However, in our haptic study we found end effects at both ends of the range. We also expect to find such effects here and used a model from our haptic study to account for these effects.

The first three experiments were a transference of our haptic experiments to the visual modality, but in Experiments 4 and 5 we go beyond that study. It has been suggested that have a shared representation of number and physical magnitude (Walsh, 2003). If this is true for numbers outside the subitizing range we expect performance similar to that for dot size recognition. Therefore, in Experiment 4 we investigated whether response times for recognition of numbers outside the subitizing range (8, 16, 32, 64 or 128 items) follow the same pattern as those for dot size recognition. This would indicate that mapping of physical magnitude is shared with mapping of numerosities outside the subitizing range. If discriminability for large numbers follows the power law we do not expect a special regime for the smallest numerosities in the range in this case. Finally, in Experiment 5 numerosities from the subitizing regime were added to the numerosity range from Experiment 4 and we investigated how this affected recognition of the larger numerosities in the range. If discriminability of small numbers is indeed much better than that of large numbers, we expect that adding numbers from the subitizing range will not affect recognizability of the larger numerosities.

7.2 General Method

7.2.1 Participants

Ten paid subjects (age 21 ± 3 years) participated in Experiments 1, 2, and 3. Five of them were female. They performed the three experiments in counterbalanced order. Ten other paid subjects (age 21 ± 2 years) participated in Experiments 4 and 5. Two of them were male. They performed the two experiments in counterbalanced order. All participants had normal or corrected to normal vision. They were treated in accordance with the local guidelines and gave their informed consent.

7.2.2 Set-up and procedure

Stimuli were presented on a 20 inch LCD monitor (Apple Cinema) with a 1050 × 1680 pixels resolution. A mask was placed over the monitor, leaving a circular display area with a diameter of 25 cm. Varying numbers of black dots were presented on a white background. The circular area over which the dots were randomly distributed could be varied and will be referred to as the occupied area. The display was controlled using a LabVIEW program running under Mac OS. Time measurement was started when the dots appeared on the screen and was terminated when a vocal response was registered using a microphone. Through this system, response times were recorded with an accuracy of up to 3 ms.

Subjects were seated in a dark room at a distance of 57 cm from the monitor with their chin in a chin rest. At this distance an image of 1 cm on the monitor corresponded to 1° visual angle. First a fixation cross appeared in the centre of the display. After 1 s the cross disappeared and the stimulus was presented. The stimulus remained visible until a response was registered after which the stimulus disappeared. Subjects were instructed to respond as fast as possible either the number of dots (Experiments 1, 2, 4 and 5) or the dot size (Experiment 3) that was presented. It was also emphasized that it was important that the answer was correct. After each trial the experimenter entered the response into the computer and feedback on whether the answer was correct was shown on the screen for 1 s in all experiments. If the answer was incorrect, also the correct response was shown. Each experiment was preceded by a training session before the experiment was started. Subjects performed at least 20 training trials and training trials were continued until 10 in a row were answered correctly.

7.2.3 Analysis

Because subjects were instructed to respond correctly and therefore minimize their errors, the error rates should be low in all experiments. Also in the subitizing regime the error rate should be roughly zero. Therefore, error rates are shown as an indication that subjects could perform the task correctly and the response times were used for further analysis. Response times of incorrectly answered trials were excluded from the analysis. Also, response times that deviated more than 3 SD from the mean were discarded as outliers. When sphericity was violated in the statistical analysis, Greenhouse-Geisser corrected values are reported. When the analysis involved regression, we report the results from the regression to the response times averaged over subjects. We also report the mean parameter values determined through regression of the model to the single subjects' data. Note that this does not necessarily yield the same outcome. Regression to the data averaged over subjects is more accurate, but it is also important to show that the same trend is present in the data for each subject individually. Therefore, the results from both procedures are reported. In all regression procedures the response times were weighted according to their inverse squared standard deviations.

7.3 Experiment 1

The purpose of this experiment was to validate our experimental paradigm (e.g. Mandler & Shebo, 1982; Trick & Pylyshyn, 1993). In order to do so, we reproduce the classical two regimes in visual numerosity judgement for small and larger numerosities. The slope of the response times as a function of the number of items and the transition point from subitizing to counting may depend on the stimulus and varies among subjects. This experiment was performed to determine these values for the specific stimulus used in this particular experimental design and this pool of subjects.

7.3.1 Method

Stimuli

In this experiment 1, 2, 3, 4, 5, 6, 7, 8 or 9 black dots were presented on a white background. The dots had a diameter of 0.5° and the occupied area had a diameter of 20° . The dots were placed such that their edges were at least 0.8° apart and 0.8° from the edge of the occupied area. Each numerosity was presented 16 times.

Analysis

To accurately determine the values of the slopes in the subitizing and counting regimes without making assumptions about the location of the transition point between the regimes, regression of a bilinear model was used. The bilinear function is given by:

$$T(N) = (r_1 N + c_1) H\left(\frac{c_2 - c_1}{r_1 - r_2} - N\right) + (r_2 N + c_2) H\left(N - \frac{c_2 - c_1}{r_1 - r_2}\right) \quad (7.1)$$

where N is the number of items, $H(N)$ is the Heaviside step function and r_1 and r_2 are the slopes, while c_1 and c_2 represent constant offsets. Note that through this analysis the location of the transition point follows from the intersection of the two linear parts and is given by:

$$N_t = \left(\frac{c_2 - c_1}{r_1 - r_2}\right) \quad (7.2)$$

The last data point at 9 items was not included in the regression analysis, because of possible end-effects. Subjects usually learn what the maximum numerosity is during the experiments, so after counting the first 8 items they already know that the answer is 9. This reduces response times and this might lead to deviations from linearity for the response times of the largest numerosity in the range. Excluding the largest numerosity is commonly done in numerosity judgement studies (e.g. Trick & Pylyshyn, 1993; Watson et al., 2007; Trick, 2008).

7.3.2 Results

The response times and error rates averaged over subjects are shown in Figure 7.1. It can be seen that enumeration was error-free for up to four items. Repeated measures ANOVA on the response times with numerosity as within subjects factor, showed a significant main effect ($F(1.8, 16) = 148, p < 0.001$). Trend analysis showed that there was a significant deviation from linearity ($F(1, 9) > 23, p < 0.001$). Regression of the bilinear function to the response times averaged over subjects and weighted according to their standard deviation, yielded a slope of 46 ms/item for the first part of the stimulus range and a slope of 270 ms/item for the second part of the range

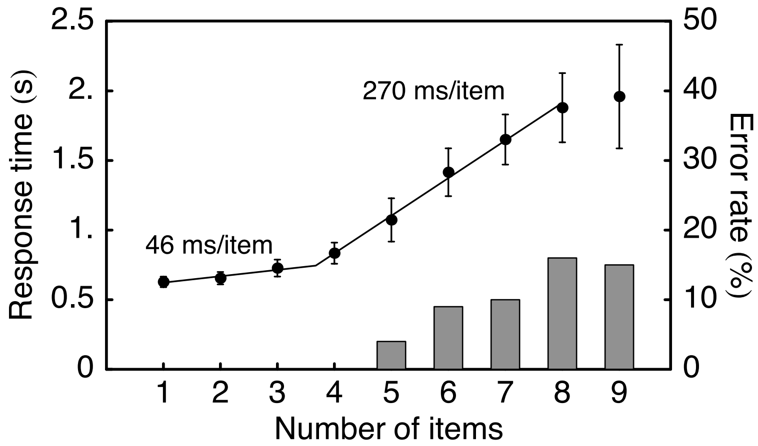

Figure 7.1: Response times (dots) and error rates (bars) averaged over subjects from Experiment 1. The solid line represents the best fit of the bilinear function to the response times averaged over subjects. Slope values are indicated in the figure. The response time for 9 items was not included in the regression analysis. Error bars indicate the standard deviation of the single subject means.

($R^2 = 0.99$). The transition point was located at 3.7 items, so in between 3 and 4 items.

As was mentioned before, the transition point and also the response time slopes may vary among subjects. Therefore, the response times were also analyzed for each subject separately. The bilinear model was fitted to the single subjects' response times. The slopes and transition points from the individual subjects were then averaged. This yielded a slope of 35 $\pm$ 9 ms/item (SE) for the first regime and 272 $\pm$ 17 ms/item (SE) for the second regime. The transition point was located at 3.6 $\pm$ 0.3 (SE) items. For four subjects the transition point was in between 4 and 5 items, three subjects had the transition point in between 3 and 4 items and two of the subjects had the transition point between 2 and 3 items. The overall quality of the fits was good, $R^2 = 0.989 \pm 0.002$ (SE).

7.3.3 Discussion

The values of the subitizing and counting slopes found here are in agreement with the existing literature on numerosity judgement of 40 – 100 ms/item in the subitizing range and 200 – 400 ms/item in the counting range (e.g. Akin

& Chase, 1978; Oyama et al., 1981; Trick & Pylyshyn, 1993; Trick, 2008). Note that this does not necessarily mean that different processes are used for small and large numerosities. There could still be a single underlying process. Rather, these results show that our results are comparable to previous results.

It has been proposed that small numbers are somehow recognized fast and accurately, so there is no need to count them. A possible explanation for a transition from subitizing to counting is then that the relative differences between the subsequent numerosities become successively smaller. When the relative differences are large it may be easy to recognize a certain numerosity. If this were true, it is expected that also larger numerosities can be easily and accurately recognized if the presented numerosities are chosen such that the relative differences are large over the whole range. In that case, there should be no longer an advantage for small numerosities. This was investigated in Experiment 2.

7.4 Experiment 2

The purpose of this experiment was to investigate how response times were influenced by the relative differences between subsequent numerosities in the presented range. The numerosity range was chosen such that there was always a factor of two between subsequent numerosities, because this was the largest relative difference between subsequent numerosities in Experiment 1. We expect that subjects can recognize the different numerosities without counting and response times will be smaller than those found in the counting range in Experiment 1. If an advantage for small numerosities is found, this indicates that subitizing is not related to relative differences between the numerosities. To exclude the possibility that larger response times for larger numerosities were caused by a longer time needed to verbalize these numbers, a control experiment was carried out.

7.4.1 Method

Subjects were shown 1, 2, 4, 8, 16 or 32 dots and they had to respond the number of dots. Subjects were explicitly told which numbers could be presented before the experiment started. Dot diameter was the same as in Experiment 1 (0.5°) and the occupied area had a diameter of 20°. Also a control condition was performed in which subjects were shown digits forming the numbers: 1, 2, 4, 8, 16 or 32, in the centre of the screen and subjects had to respond by calling out the presented number. The height of a digit was 2°.

7.4.2 Results

Response times averaged over subjects and error rates for the different numerosities are shown in Figure 7.2a. It can be seen that the responses were faster for small numerosities (< 4), compared to larger numerosities. Repeated measures ANOVA on the response times showed that the effect for numerosity was significant ($F(1.2, 10.8) = 18.6, p < 0.001$). Trend analysis

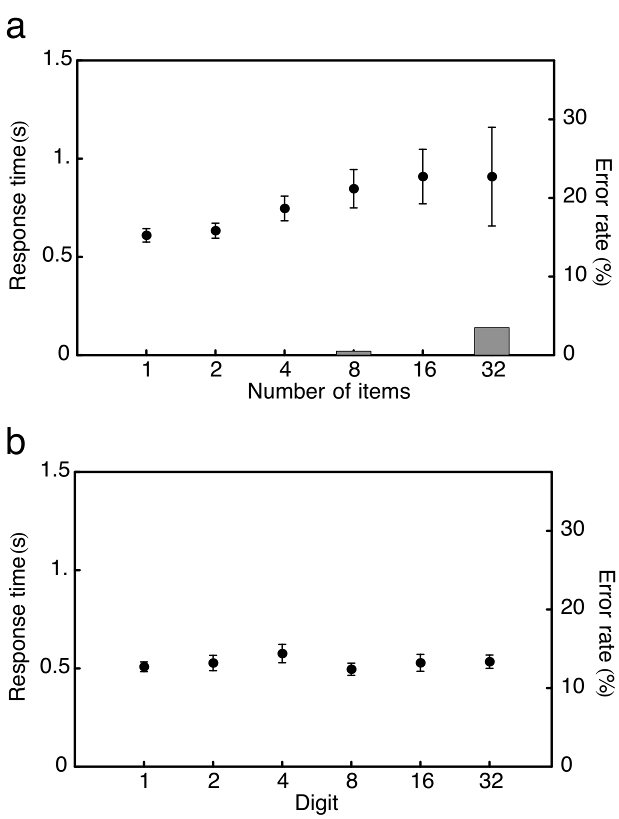

Figure 7.2: a) Response times (dots) and error rates (bars) averaged over subjects from Experiment 2. b) Response times and error rates (these were zero for all numbers) averaged over subjects in the control condition. The error bars represent the standard deviation of the single subject means.

showed that there was a significant linear trend ($F(1,9) = 23.3, p < 0.001$) and a significant cubic trend ($F(1,9) = 24, p < 0.001$) in the response times. This indicates that there was an increase of the response times from small to larger numerosities, but there was also twice a change of direction of the trend. This resulted in the S-like shape in the response times that can be seen in Figure 7.2a. Regression of a linear function yielded a significant slope of 17 ms/item (p = 0.03,R^2 = 0.7).

The results for the control condition are shown in Figure 7.2b. It can be seen that response times are relatively constant over the whole range and no errors were made. Repeated measures ANOVA on the response times showed that there was a significant effect of numerosity ($F(15,45) = 17, p < 0.001$). However, the linear trend was not significant ($F(1,9) = 1.4, p = 0.27$). Pair-wise comparisons showed that there were several significant differences between the different numbers. The largest average difference was 80 ms between numbers 4 and 8 ($p = 0.001$, Bonferroni corrected value).

7.4.3 Discussion

The control experiment showed that there was an effect of numerosity. But more importantly, there was no increase of the response times from small to large numbers. This shows that there was no difference in the time needed to verbalize small and large numbers. Therefore, this cannot explain the advantage in enumeration of small numerosities.

In the main experiment response times were well below 1.5 s over the whole numerosity range, so subjects were clearly not counting the items. From Experiment 1 it can be seen that counting 8 items already takes 2 s. Therefore, we conclude that subjects could recognize the large numerosities (8, 16 and 32) without counting. The results show that when the relative differences between subsequent numerosities are large over the whole numerosity range, subjects can recognize all numerosities without counting. However, there was still an advantage for small numerosities. This shows that small numerosities were recognized faster than large numerosities for reasons other than the relative differences between subsequent numerosities. This is in agreement with what we found in our previous study on haptic numerosity judgement (Plaisier, Bergmann Tiest, & Kappers, 2009a). To investigate what mediates this fast recognition of small numbers, Experiment 3 was carried out in which numerosity information was removed and only other magnitude information was present. It has been suggested that representation of number is shared with magnitude representation. If this fast performance for small numerosities is specific to number representation, we do not expect it to appear for the smallest stimuli in Experiment 3.

7.5 Experiment 3

In this experiment subjects were shown a dot in the centre of the screen. The area of the dot always corresponded to the total area of one of the different numbers of dots from Experiment 2. The dots were numbered accordingly and subjects had to respond the number that was associated with the dot size that was presented. Subjects could recognize the different dots by judging occupied area and luminance. These cues were also present in the stimuli of Experiment 2 and the only difference with respect to the stimuli of Experiment 2 is that the black pixels were all contained within a single disk around the centre instead of distributed over different disks. Consequently, if the fast recognition of small numerosities found in Experiment 2 was mediated by these cues, we expect that we will also find it in this experiment. If the special performance disappears we can conclude that the fast recognition of small numbers is related to black pixels being distributed in a certain way.

7.5.1 Method

Subjects were shown dots that had an area equivalent to the total area of the varying numbers of dots in Experiment 2. They had to respond with the corresponding label. For instance, when subjects saw the dot with area corresponding to the area of 4 dots in Experiment 2 (i.e. dot with diameter 1°), they had to respond by calling out 4. Consequently, the presented dots had a diameter of 0.5°, 0.7°, 1°, 1.4°, 2° or 2.8°. The subjects were shown the different dot sizes together with the labels before the training session was started. This mapping was not visible during the training session or experiment.

7.5.2 Results

Figure 7.3 shows response times and error rates averaged over subjects for the different dot sizes. Error rates were low (<20 %) over the whole stimulus range, indicating that subjects could perform the task correctly. Errors occur over the whole stimulus range in this case and not only for the largest numerosities in the range like in Experiment 2. It can be seen that there is no clear advantage for small numerosities. Although response times increase from 1 to 4 items, they decrease again for 8 and 32 items. Repeated measures ANOVA in the response times showed that the effect of dot size was significant ($F(1.4, 12.6) = 6.8, p = 0.02$). Trend analysis showed that there was a significant quadratic trend ($F(1, 9) = 87.5, p < 0.001$). This means that the trend in the response times had an inverted U-shape, as can be seen in Figure 7.3. There was no significant linear trend. Regression of a linear function to the response times did not yield a significant slope (p = 0.1, $R^2 = 0.5$).

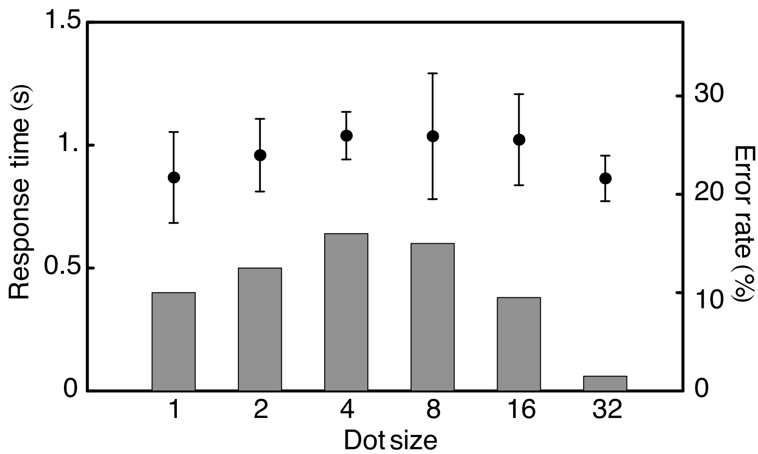

Figure 7.3: Response times (dots) and error rates (bars) averaged over subjects for Experiment 3. The error bars represent the standard deviation of the single subject means.

7.5.3 Discussion

Error rates are generally larger than in Experiment 2, indicating that this task was more difficult. This is not surprising given the fact that numerosity information was removed, so there was less information left in the stimuli. However, when numerosity information was absent, subjects were still able to name the different stimuli correctly and there was a significant trend in the response times. This trend was different from the trend that was found in Experiment 2. When numerosity information was removed there was no longer faster or more accurate performance for small numerosities compared to larger numerosities. Consequently, there was no linear trend, showing that there was no increase of the response times from small to large numbers of items. This suggests that black pixels have to be distributed over several disks to enable fast and accurate performance at the first part of the stimulus range. Response times were, however, not constant over the whole range as indicated by the relatively low R^2 value of the linear function. They decrease at both sides of the stimulus range. This was also the case in our haptic study and we have introduced a model to describe this behavior.

7.6 Model

It has been shown that response times for judging which of two numbers is larger decreases if the difference between the numbers increases (Moyer & Landauer, 1967). This suggest that response times vary with discriminability between numbers. In our paper on haptic numerosity judgement we have introduced a model to describe response times for recognition of a certain stimulus based on discriminability differences between different stimuli (Plaisier, Bergmann Tiest, & Kappers, 2009a). This model describes the pattern of response times only when discriminability follows Fechner's law over the whole range of stimuli. Note that this model describes response times for naming of stimuli that vary in magnitude, not necessarily stimuli differing in numerosity. However, it is often argued that number representation is similar to magnitude representation. Furthermore, it is possible that numerosity is not accessed directly, but through other co-varying cues like luminance. In our haptic study, the model described the pattern in response times very well when subjects had to label spheres differing in size (i.e. when numerosity information was absent). However, as expected, it could not describe the response times when subjects had to enumerate varying numbers of spheres in their hand (i.e. when numerosity information was present), indicating that discriminability did not follow Fechner's law over the whole range of numerosities. If indeed similar processes underlie haptic and visual number recognition, then this estimation model should be able to describe the response times from Experiment 3, but not those from Experiment 2 of the present study.

7.6.1 Derivation

Our model assumes that when a presented stimulus has to be recognized and the correct label has to be given, all stimuli in the range are considered weighted according to discriminability between the presented stimulus and each of the other possible stimuli. In accordance with Fechner's law, discriminability is assumed to be proportional to the logarithm of the ratio between the two compared stimuli. The discriminability d between quantities x_1 and x_2 is thus given by:

$$d(x_1, x_2) \propto \left| \log \frac{x_1}{x_2} \right| \tag{7.3}$$

The total response time is assumed to be inversely proportional to the sum of the discriminabilities. The response time as a function of the presented quantity N can then be described by:

$$T(N) = a + \frac{b}{\sum_{n=i}^{j} |\log \frac{N}{n}|} \tag{7.4}$$

where N is the quantity that is presented, n is an iterator which runs from the smallest quantity in the set (i) to the largest one (j) over all quantities in the set. Free parameters a and b scale the offset and shape of the function. Here, parameter b alone determines the shape of the function, but the average response time over all numerosities in the range (μ) is determined by a combination of a and b:

$$\mu = a + \frac{b \sum_{N=i}^{j} \frac{1}{\sum_{n=i}^{j} |\log \frac{N}{n}|}}{\sum_{n=i}^{j} 1} \tag{7.5}$$

Note that this model predicts that response times decrease towards both ends of the stimulus range. For instance, when the smallest stimulus is presented, there is no smaller one to which it can be compared. Similarly, when the largest stimulus is presented there is no larger stimulus to which it can be compared. Furthermore, if the relative differences between subsequent numerosities are constant, the shape of the function will be symmetrical with the maximum in the middle of the stimulus range. This is illustrated in Figure 7.4. In this figure it can also be seen that the predicted response times will depend on the stimulus range that is presented. Because in this model response times are modeled as a function of the presented range it is crucial that data from the whole range are included in the analysis. This was not the case in Experiment 1, where the last stimulus with the largest numerosity was discarded from the analysis because of possible end-effects. The bi-linear model from Experiment 1 does not predict end-effects and to

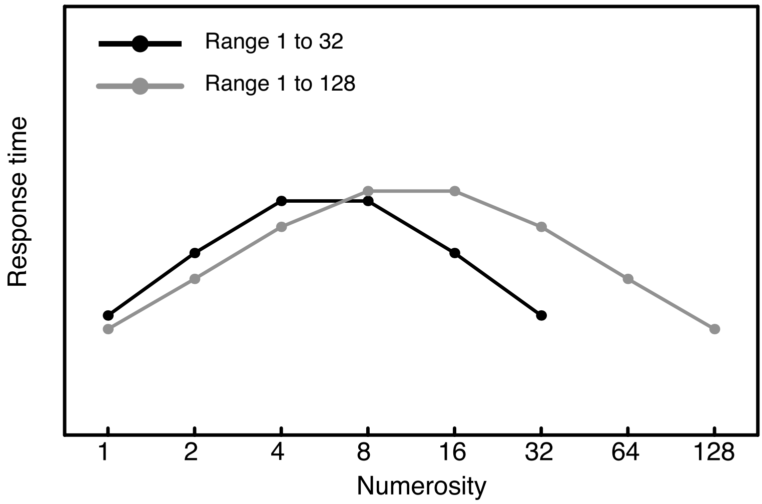

Figure 7.4: Predicted pattern in the response times as a function of the number of items. This is a discrete model and only defined at whole numbers. Therefore the predicted response times are indicated by the dots and these were connected for clarity. Response times for a range from 1 to 32 are shown in black, while those for range 1 to 128 are shown in grey. It can be seen that the predicted response times very much depend on the stimulus range. Note that the scaling in the vertical direction is determined by free parameter b. Therefore, the actual response time may be scaled differently comparing both ranges.

determine the counting slope correctly the last data point should be discarded. The model presented here was fitted to the response times from Experiment 2 and Experiment 3.

7.6.2 Regression analysis

Figure 7.5a shows the response times for the different numbers of items in Experiment 2. The response times for the different dot sizes from Experiment 3 are shown in Figure 7.5b. For both conditions the best fit of the estimation model is represented by the solid line. As can be seen the model cannot describe the data from Experiment 2 (R^2 =0.38) and performs even worse than a linear function. However, it describes the response times from Experiment 3 very well ($R^2 = 0.96$) and much better than a single linear function ($R^2 = 0.5$). The values of the fitting parameters were $b = 2.1$ s and $\mu = 0.7$ s.

Again, the regression analysis was also performed on the data from the single subjects. Averaging the R^2 values from each subject in Experiment 2 yielded $R^2 = 0.009 \pm 0.0009$ (SE). So the model cannot describe the relation between numerosity and response time. This is in agreement with the result from the regression to the response times averaged over subjects. For Experiment 3, this analysis yielded $R^2 = 0.6 \pm 0.09$ (SE), indicating that the model can describe the data in this case. The resulting fitting parameters averaged over subjects were $b = 2.8 \pm 0.3$ s and $\mu = 0.94 \pm 0.04$ s (SE).

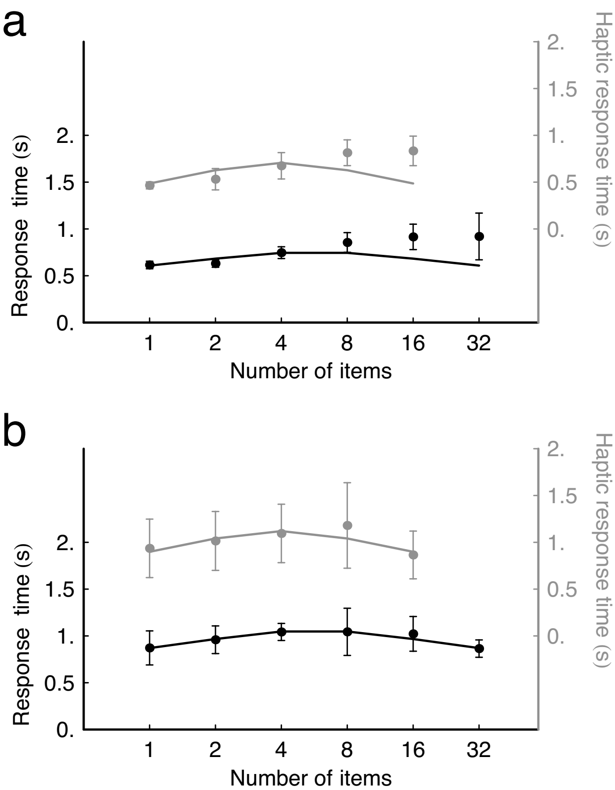

Figure 7.5: Response times from Experiment 2 (a) and Experiment 3 (b) with the best fit of the estimation model (solid line). The response times from the haptic study are plotted in grey. In that case the maximum number was 16. Note the upward shift of the axis.

7.6.3 Discussion

Our analysis shows that our model describes the response times for Experiment 3, where no numerosity information was present. As expected, it does not describe the data from Experiment 2. In Experiment 2, recognition of small numerosities (< 4) was faster than for the larger numerosities. This suggests that discriminability for small numerosities is much larger than for large numerosities even though the relative differences were the same. This is in agreement with what we have reported previously in haptic numerosity judgement. In Figure 7.5 the response times from our haptic study are plotted in grey. Note that for clarity the axis for the haptic response times is shifted upwards. It can be seen that the haptic response times correspond relatively well with the response times from the present visual study, although in the haptic case the stimulus range ended at 16 items. In both modalities, faster performance for numerosities from the subitizing range was found than outside this range. In both cases this faster performance disappeared when stimuli were coded in physical magnitude. This suggests that in both cases response times for the first part of the stimulus range were smaller than for the last part of the stimulus range, but only if numerosity information was present. This indicates that discriminability was better for numerosities from the subitizing range than for larger numerosities. This raises the question whether response times follow a similar pattern as those for magnitude estimation when only numbers larger than the subitizing range are shown.

7.7 Experiment 4

In this Experiment we investigated whether discriminability of numbers larger than the subitizing range follows Fechner's law. Therefore, we removed the numerosities in the subitizing regime from the range of numerosities that was used in Experiment 2 and extended the range to larger numerosities. In this experiment we prevented subjects from using other cues like occupied area, density and luminance by using the same method as Izard and Dehaene (2008) recently reported[1].

[1] This manipulation of the stimuli was not applied in Experiments 1 and 2, to keep the results of these experiments comparable to those of previous studies in which this manipulation was usually not done. Comparison of Experiments 2 and 3 already shows that the pattern in the response times of Experiment 2 are not likely caused by luminance or occupied area estimation only.

7.7.1 Method

The set-up and task were as described in the General Method section. Subjects were presented with 8, 16, 32, 64 or 128 dots randomly distributed over the occupied area. They were explicitly told which numbers could be presented. There were three different types of trials. In one third of the trials dot size (0.15° diameter) and occupied area were kept constant (20° diameter). In another third of the trials the occupied area was varied such that dot density was constant for all numerosities (0.15° dot diameter and occupied area ranged from 5.4° to 21.5° diameter). In the last third of the trials the dot size was varied such that the total luminance was constant for all numerosities (dot diameter varied from 1° to 0.25° and occupied area was 21.7° diameter). All three trial types were interleaved randomly so that only numerosity was a reliant cue in all trials.

7.7.2 Results

Repeated measures ANOVA with numerosity and trial type as factors showed an effect of numerosity ($F(1.3, 12) = 7.8, p = 0.012$) and of trial type ($F(2, 18) = 4.8, p = 0.022$). There was no interaction between both factors ($F(3.2, 29) = 0.98, p = 0.46$) and the quadratic trend was significant ($p = 0.018$). Post-hoc tests (paired t-tests with Bonferroni correction) did not show significant differences between the trial types ($p \geq 0.07$). This indicates that there were no significant differences in the shape of the response times for the different trial types. To be certain of this, regression of the estimation model was performed for the three trial types separately. This analysis yielded $b = 5.9$ and $\mu = 1.2$ s for the trials with varying dot sizes, $b = 5.9$, $\mu = 1.2$ s for the trials with varying occupied area and $b = 6.0$ $\mu = 1.1$ s for the trials in which occupied area and dot size were constant ($R^2 \geq 0.7$). The lack of significant differences in the shapes of the response times allowed us to collapse the three different trial types. Regression to the data with all trial types collapsed yielded $b = 5.9$ and $\mu = 1.1$ s ($R^2 = 0.8$). Figure 7.6 shows the response times and error rates averaged over subjects for all numerosities. It can be seen that the response times follow a pattern similar to that found in Experiment 3. The solid line represents regression of the estimation model to the response times averaged over subjects. For comparison, regression of a linear function did, like in Experiment 3, not yield a significant slope ($p = 0.1$) and performed much worse($\mathrm{R}^2 = 0.4$) than our model.

Regression of our model to the single subject response times yielded

$R^2 = 0.7 \pm 0.08$ (SE), averaged over all subjects. The values of the shape parameter and the average response time were $b = 5 \pm 2$ s (SE) and $\mu = 0.9 \pm 0.03$ s (SE), respectively.

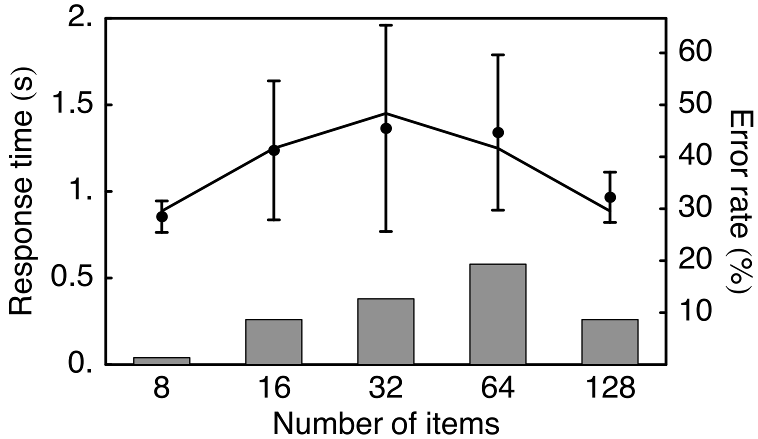

Figure 7.6: Response times (dots) and error rates (bars) averaged over subjects from Experiment 4. The solid line represents the best fit of the estimation model to the response times averaged over subjects. Error bars indicate the standard deviation of the single subject means.

7.7.3 Discussion

These results show that our model can indeed describe response times when numerosity information is present when all numerosities are larger than the subitizing range. This indicates that discriminability between subsequent numerosities is constant over this range of numerosities. Note that this conclusion is also supported by the analysis of the three trial types separately and the conclusion does not change depending on whether we collapse the three trial types or not. In Experiment 5 we investigated whether response times for recognition of large numbers are influenced by the presence of numerosities from the subitizing regime in the presented range of numerosities.

7.8 Experiment 5

In this experiment we investigated whether numerosities from the subitizing regime are taken into consideration during the estimation of larger numerosities. If they are, then adding them to the numerosity range should yield the inverted U-shaped pattern from Experiment 4, but now symmetrical around 8 and 16 (the middle of the range). However, if they are not taken into consideration, then the pattern in the response times should be the same as found in Experiment 4. In this last case we can conclude that small numbers are not taken into consideration or discarded very fast when a large number is presented.

7.8.1 Method

Subjects were presented with 1, 2, 4, 8, 16, 32, 64 or 128 dots randomly distributed over the occupied area. Again subjects were explicitly told which numbers could be presented. Luminance and dot density cues were removed as described in the Method section of Experiment 4. In the trials where dot density was constant for all numerosities, the occupied area now ranged from $1.9°$ to $21.5°$ diameter and in the constant luminance trials the dot size ranged from $2.8°$ to $0.25°$ diameter.

7.8.2 Results

Repeated measures ANOVA with numerosity and trial type as factors showed an effect of numerosity ($F(1.9, 17) = 23.4, p < 0.0001$), but not of trial type ($F(2, 18) = 2.6, p = 0.099$). Therefore, the data from the three different types of trials were collapsed. Response times and error rates averaged over subjects are shown in Figure 7.7. It can be seen that from numerosity 8 and larger the response times follow a similar pattern as found in Experiment 4. The estimation model was fitted to the response times averaged over subjects for different numerosity intervals. The interval over which the quality of the fit is best, indicates the range of numerosities that is included in the estimation process. As was shown earlier, the shape of the model depends on the range of numerosities (Figure 7.4). There were six intervals ranging from 1 to 128, 2 to 128 and so on to the interval from 32 to 128. The R^2 values that were found were 0.3, 0.5, 0.8, 0.9, 0.7 and 0.2, respectively. The optimum in the quality of the fit was thus found over the interval from 8 to 128, i.e. all numerosities well outside the subitizing regime. Regression of the model over this interval is represented with the solid black line in Figure 7.7. The value of the shape parameter and the average response times were found to be $b = 3.5$ s and $\mu = 1.2$ s, respectively. For comparison, regression of a linear function was performed for the whole range of stimuli and over the interval from 8 to 128 separately. Over the whole range the resulting R^2

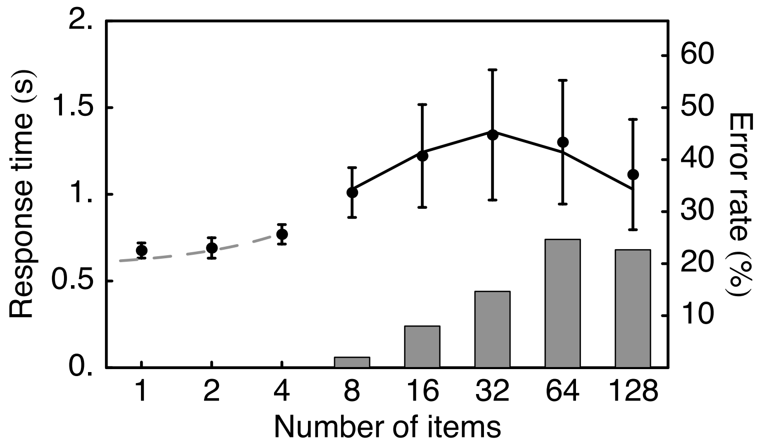

Figure 7.7: Response times (dots) and error rates (bars) averaged over subjects from Experiment 5. The solid black line represents the best fit of the estimation model to the response times. The grey dashed line is the first linear part from the fit of the bilinear function to the data from Experiment 1 plotted on an logarithmic scale. Error bars indicate the standard deviation of the single subject means.

value was 0.15 and for the interval from 8 to 128, R^2 was 0.12. This shows that our model describes the data much better than a linear function.

Also, regression to the single subjects' data was performed. This yielded on average $R^2 = 0.6 \pm 0.1$ (SE), so the model fitted the data well. The shape parameter and average response time were found to be $b = 4 \pm 1$ s (SE) and $\mu = 1.2 \pm 0.8$ s (SE), respectively.

As the same subjects participated in both Experiments 4 and 5 and they performed the experiment in counterbalanced order, the fitting parameters were compared between the experiments. Paired-samples t-tests yielded no significant differences ($p \geq 0.07$) between the experiments for both parameters.

The dashed grey line in Figure 7.7 is the result from the fit for numerosities in the subitizing regime from experiment 1, re-plotted on a logarithmic scale. Because of the logarithmic scaling, the linear function is now curved. It can be seen that the line fits also the response times from this experiment, even though different subjects participated in both experiments. This shows that the response times for numerosities in the subitizing regime were not affected by the difference in the presented numerosities between this experiment and Experiment 1.

7.8.3 Discussion

The results show that adding numerosities from the subitizing regime did not significantly change the response times for numerosities outside the subitizing range. The pattern in the response times was symmetrical around 32 items, which was the middle numerosity between 8 and 128 (i.e. the numerosities outside the subitizing regime). This indicates that numerosities from the subitizing regime were not taken into consideration when numerosities outside the subitizing range were presented. Furthermore, the response times in the subitizing range were comparable to those found in Experiment 1. This indicates that the subitizing process was relatively unaffected by the differences between the numerosity ranges used in Experiment 1 and Experiment 5. These results show that numbers from the subitizing range are not taken into consideration or were discarded very fast when a numerosity outside the subitizing range is shown and vice versa.

7.9 General Discussion

The results from Experiments 1, 2 and 3 are show in agreement with the results from our haptic study (Plaisier, Bergmann Tiest, & Kappers, 2009a). Note that the stimuli differ in many ways between the haptic study and the present visual study. In the haptic case, spheres were grasped and could be actively rearranged in the hand. In vision there is no such active control over the positions of the dots. In the case of vision, on the other hand, pattern recognition may play a role. Pattern recognition has been suggested as an explanation for subitizing (Mandler & Shebo, 1982). Pattern recognition does not seem applicable to the haptic case as the positions of the spheres were not fixed. Moreover, pattern recognition is not likely to have played a role in the study on tactile subitizing where varying numbers of fingers were stimulated (Riggs et al., 2006). The fact that despite these differences, numbers up to three or four are recognized faster and more accurate than larger numbers in vision as well as haptics suggests that the underlying reason may be the same in both modalities. This has interesting implications for the possible processes underlying numerosity judgement, as these should be processes that extend across both modalities. Consequently, pattern recognition not a very likely explanation.

From Experiments 4 and 5 it is clear that numerosities from the subitizing range are not taken into consideration when numerosities larger than the subitizing range are shown. This in line with the idea that subitizing means that subjects almost instantaneously know which numerosity is presented. This does not only mean that subjects perform practically error-free in the subitizing regime, they also know very quickly whether or not the presented numerosity can be subitized. The results from Experiments 2 and 5 both show that even if the relative spacing between subsequent numerosities is large over the whole numerosity range, there is an advantage for enumeration of small numerosities. So constant relative magnitude differences between the numerosities do not enable subitizing for larger numerosities. It was mentioned before that pattern recognition is also not a likely explanation. Still, it seems that numerosities from the subitizing regime are recognized as 'subitizible' very efficiently. It has been shown that numerosities from the subitizing range are rated as more dissimilar than numerosities from outside that range (Logan & Zbrodoff, 2003). This would explain why adding numerosities from the subitizing regime did not affect the response times for recognition of larger numerosities (Experiment 5) much. Now the question arises of what enables this fast recognition of small numerosities?

An explanation for the subitizing mechanism that does not involve discriminability or pattern recognition is based on visual indexing theory (see Pylyshyn (2001) for a review). According to this theory humans can refer to an item without linking it to a specific feature like position. From visual tracking studies, it was found that subjects can track up to 5 items simultaneously and it is hypothesized that the number of items that can be referred to simultaneously in this way is limited to 5 (Pylyshyn & Storm, 1988). This idea can also be used to explain why numerosities smaller than 5 can be enumerated faster and more accurately than larger numerosities (Trick & Pylyshyn, 1994). Although there is no evidence that a process like haptic indexing exists, it is possible that the idea of 'indexing' is not limited to the visual modality.

In conclusion, we have shown that there is an advantage for judging of small numerosities (< 4) over large numerosities even if the relative differences between subsequent stimuli is a factor of 2 over the stimulus range. This advantage was not mediated by recognition of the numerosities through judgment of density, occupied area or luminance. Furthermore, the faster performance for the smallest stimuli in the range disappeared when numerosity information was removed. This supports the idea that subitizing does not reflect very accurate estimation mediated through large differences between subsequent numerosities. Furthermore, we would like to propose that similar processes underly haptic and visual numerosity judgment.

Chapter 8
Bimanual number processing

Accepted for publication as:
Plaisier, M.A., Bergmann Tiest, W.M., and Kappers, A.M.L. Grabbing subitizing with both hands: Bimanual number processing, *Experimental Brain Research*, 2010, DOI 10.1007/s00221-009-2146-1.

Abstract
Visual judgment of small numerosities (< 4) is generally assumed to be done through subitizing which is a faster process than counting. Subitizing has also been shown to occur in haptic judgment of the number of spheres in the hand. Furthermore, interactions have been shown to exist between visually perceived numbers and hand motor action. In this study we compare enumeration of a set of spheres presented to one hand (unimanual) and enumeration of the same total number of spheres presented divided over the two hands (bimanual). Our results show that, like in vision, a combination of subitizing and counting is used to process numbers in active touch. This shows that numbers are processed in a modality-independent way. This suggests that there are not only interactions between perception of numbers and hand motor action, but rather that number representation is modality-independent.

8.1 Introduction

For exact numerosity judgment, two enumeration processes have been identified: a fast and highly accurate process labeled ‘subitizing’ for enumeration of small numbers of items (< 4) and a slower and more error-prone process referred to as ‘counting’ for larger numbers of items (Kaufman et al., 1949). This distinction between judgment of small and large numbers has not only been shown to exist in vision (e.g. Atkinson et al., 1976; Mandler & Shebo, 1982; Trick & Pylyshyn, 1993; Trick, 2008), but also in audition (Ten Hoopen & Vos, 1979) and more recently in touch (Riggs et al., 2006; Plaisier, Bergmann Tiest, & Kappers, 2009a). The finding that subitizing occurs in touch is particularly interesting because it has been shown that parieto-frontal brain circuits dedicated to number processing partially overlap with those dedicated to hand and finger movements (e.g. Pinel, Piazza, Le Bihan, & Dehaene, 2004; Piazza et al., 2002). The existence of interactions between visually perceived numbesr and hand motor actions in terms of corticospinal excitability of the hand muscles and grip opening/closing has been shown in behavioral studies (Andres, Davare, Pesenti, Olivier, & Seron, 2004; Andres et al., 2007; Moretto & Di Pellegrino, 2008). For mediating these interactions it has been proposed that there is an analogue representation of magnitude in the parietal cortex (Walsh, 2003). The existence of these interactions shows that visually perceived numbers can evoke action. What about numbers perceived through action? The question arises of whether there are not only interactions between perceived numbers and action, but whether numbers perceived through active touch are processed in a similar way as visually perceived numbers. If so, this is an indication that magnitude representation in the parietal cortex is modality-independent. Although subitizing occurs in several modalities, it is not yet clear what kind of a process it actually is and how it is dissociated from counting. One thing that is clear from visual studies is that when observers are shown a field of dots, they do not simply add all the dots one by one to arrive at the total. Rather, they seem to enumerate small groups of dots and sum the groups to arrive at the total (Van Oeffelen & Vos, 1982a, 1984). Consequently, enumeration of large fields of dots can be affected by the spatial arrangement of the dots.

In a previous study on haptic numerosity judgment, we have shown that subitizing occurred for up to three items when subjects were asked to enumerate a number of spheres grasped together in the hand (Plaisier, Bergmann Tiest, & Kappers, 2009a). Now the question arises whether subitizing and counting are implemented in a similar way for visual and haptic perception of numbers. If this is the case, the group-and-add strategy observed in visual studies should also be possible the haptic case. To answer this question we presented subjects with varying numbers of spheres that were explored using active touch. In order to cluster sets of items together, we presented a set of spheres to each hand of the subjects. Either one set of spheres was presented to the left or right hand of the subjects (unimanual trials), or two sets were presented to each hand simultaneously (bimanual trials).

The results from the unimanual trials were used to model number processing in the bimanual case. We hypothesize three mutually exclusive outcomes. The first possibility is that subitizing is inhibited because information from both hands is combined in an inefficient way. The second possibility is that subitizing does occur, but subitizing or counting is used depending on the total number of spheres and not cluster size. Finally, there is the possibility that it depends on cluster size which enumeration mechanism is used. If this third hypothesis is true then this shows that configurational effects found in vision also occur in touch and that numbers can be processed through a combination of subitizing and counting. Such similarities between haptic and visual enumeration would be strong evidence for a modality-independent model of number processing and consequently that magnitude representation is modality-independent.

8.2 Method

8.2.1 Participants

Ten paid subjects participated in the experiment. All participants were right-handed according to Coren's test (Coren, 1993) and none of them had any known hand deficits. All subjects were naive as to the purpose of the experiment and signed a declaration of informed consent.

8.2.2 Set-up

The items consisted of brass spheres (1.86 cm diameter, 29 g) suspended from flexible wires (Figure 8.1a). These same spheres were used in our previous study into haptic numerosity judgment (Plaisier, Bergmann Tiest, & Kappers, 2009a). A custom-built device was used to measure the response times. Time measurement was started automatically when a subject touched the stimuli and it was terminated through a vocal response. The response times were measured with an accuracy of 10 ms. For technical details about this device, see Plaisier et al. (2008b).

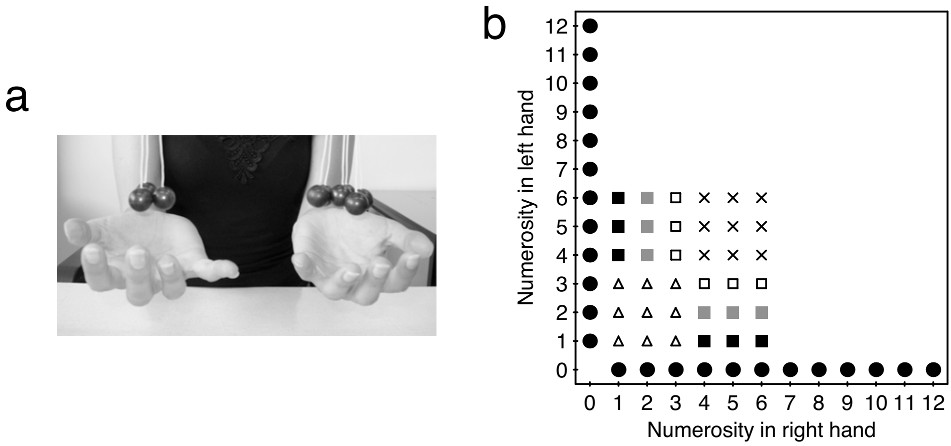

Figure 8.1: a) Picture of a subject grasping upwards to start a bimanual trial. b) Overview of all combinations of numerosities presented to the left and the right hand. The dots indicate unimanual clusters and the other symbols indicate the bimanual cluster combinations. Here subitizing up to three items is assumed. In that case the triangles indicate that both clusters were in the subitizing range and crosses indicate that both clusters were in the counting range. Squares indicate that there was one cluster from the counting range and one from the subitizing range, where a filled square indicates that the smallest cluster had 1 item, a thick outlined square indicates this cluster had 2 items and a thin outlined square indicates that the smallest cluster had 3 items.

8.2.3 Experimental design

Subjects were blindfolded and wore earplugs to eliminate sound cues. They placed their left hand in a holder on the left side and the right hand in a holder on the right side. Sets of spheres could be suspended above these holders. The experimenter informed the subject before the trial started whether the spheres were on both sides and otherwise on which side the spheres were. If there was only one set of spheres the subjects were instructed to grasp upwards with the corresponding hand and respond the correct number of spheres as fast as possible. When there were two sets of spheres, subjects were instructed to grasp upwards with both hands simultaneously and respond the total number of spheres (i.e. the sum of the spheres in the left and right hands). After each trial, subjects were told what the correct number of spheres was. There were no restrictions on exploratory strategy nor on hand movements, other than having to initially grasp all items simultaneously. After initially grasping all items it was allowed to release spheres from their hand during a trial, but subjects were instructed to only do this if they thought that this was the fastest strategy.

Subjects were presented with any number of spheres from 1 to 12. These numerosities were presented to one hand or divided over both hands. There were 33 cluster combinations all of which were presented 5 times with one cluster in the left hand and the other in the right hand and also 5 times vice versa (Figure 8.1b). Note that these combinations allow comparison of enumeration of each total numerosity in the unimanual case to the bimanual case, except for when the total numerosity was 1. Each subject performed 330 trials divided over three blocks of trials of approximately 1 hour. Trials were performed in pseudo-random order such that each numerosity was presented roughly the same number of times in each block. The blocks of trials were performed on different days or with a break of at least 2 hours in between. Before the first block of trials was started, subjects performed 20 practice trials and practice trials were continued until 10 in a row had been answered correctly. It was never necessary to exceed 30 practice trials. Error trials were repeated at the end of the block to ensure an equal number of correct trials for all cluster combinations.

8.2.4 Analysis

Because subjects were instructed to minimize the number of errors, error rates should be generally low and the response times are used for further analysis. Only response times from correctly answered trials were included in the analysis. For the unimanual trials we assume a bilinear function for the response times as a function of the number of spheres (Plaisier, Bergmann Tiest, & Kappers, 2009a). The slope of the first linear part represents the subitizing slope and the slope of the second linear part represents the counting slope. This function is defined as:

$$T_{\text{uni}}(N) = (s_s N + c_1) H\left(\frac{c_2 - c_1}{s_s - s_c} - N\right) + \\ + (s_c N + c_2) H\left(N - \frac{c_2 - c_1}{s_s - s_c}\right) \tag{8.1}$$

Here N is the presented number of spheres, $H(N)$ is the Heaviside step function, s_s and s_c are the subitizing and counting slopes, respectively, and c_1 and c_2 represent constant offsets. Note that through this analysis the location of the transition point follows from the intersection of the two linear parts and is given by:

$$N_t = \left(\frac{c_2 - c_1}{s_s - s_c}\right) \tag{8.2}$$

Regression of this function allowed the slopes to be determined without making assumptions about the transition point which is determined by the fitting parameters from the linear parts. The slope values from the unimanual trials were used to model the bimanual response times.

8.3 Bimanual models

For the bimanual trials, three hypotheses were discussed in the Introduction. The case in which subitizing does not occur and counting is used over the whole numerosity range is represented by the 'No subitizing' model (Figure 8.2 a). The response time as a function of the number of spheres in the left hand (n_1) and the right hand (n_2) is given by:

$$T_{\text{bi}}(n_1, n_2) = c_1 + s_c(n_1 + n_2) \tag{8.3}$$

The second possibility is that subitizing can be used and the enumeration process depends on the total numerosity only (Figure 8.2b). This means that if the total numerosity is in the subitizing range the total is subitized, while if the total is in the counting range it is counted regardless of the sizes of the clusters in the two hands. This 'Total numerosity dependent' model is expressed in terms of the response time function for the unimanual case as:

$$T_{\text{bi}}(n_1, n_2) = T_{\text{uni}}(n_1 + n_2) \tag{8.4}$$

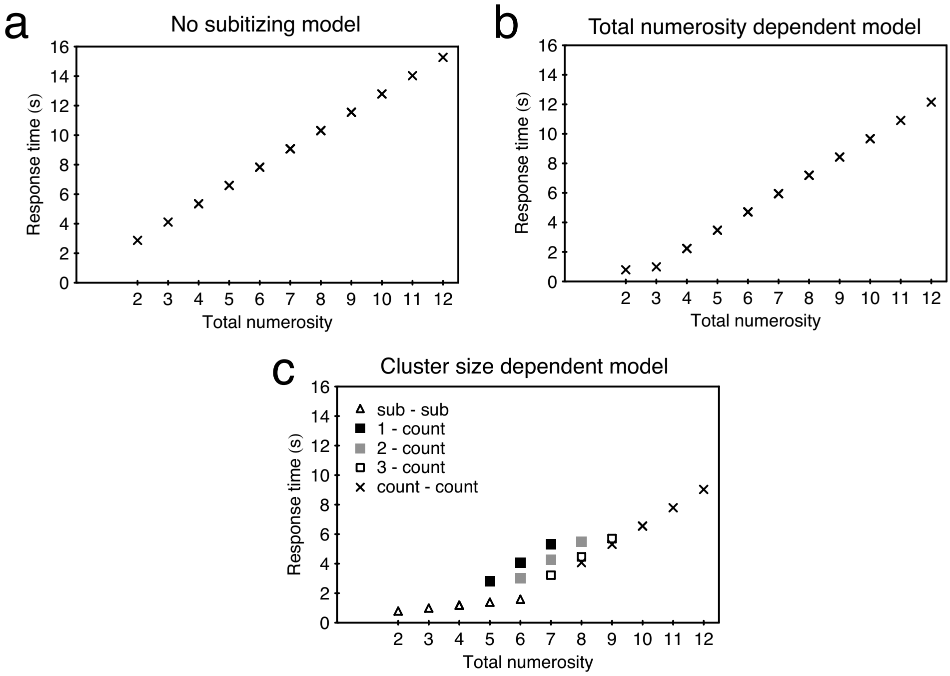

Figure 8.2: Predicted response times for bimanual trials from three different models using the slopes and offset values from the unimanual trials. Regression parameters determined from the unimanual trials were entered into each of the models to arrive at predictions of the absolute response times. There were different combinations of clusters that summed up to the same total number of spheres. The predicted response times can therefore fall on top of each other. 'No subitizing' indicates the prediction of the response times if subitizing does not occur in bimanual number processing. 'Total numerosity dependent' shows the predicted response times in the case that subitizing can be used and the enumeration mechanism that is used depends only on the total number of spheres. 'Cluster size dependent' indicates the prediction of the response times in the case that subitizing can be used and the selected enumeration process depends on the cluster size in each hand. In this case response times depend on the specific cluster combination and therefore different plot symbols were used for different cluster combinations. See text for further explanation of the models.

The last hypothesis is that subjects use either subitizing or counting depending on the cluster size. Both clusters are summed to arrive at the number of spheres. The expected response times from this 'Cluster size dependent' model are shown in Figure 8.2c and are given by:

$$T_{\mathrm{bi}}(n_1, n_2) = T_{\mathrm{uni}}(n_1) + T_{\mathrm{uni}}(n_2) - c_1 \tag{8.5}$$

The time that is needed to sum the two clusters to arrive at the total is neglected. The constant offset c_1 is subtracted once, because otherwise it would be included twice in the response time.

8.4 Results

8.4.1 Unimanual trials

Error rates were overall low (< 11 %) and no errors were made for up to 3 items. To test whether there was an advantage for the left or right hand for subitizing or counting the average response times for the subitizing and counting range were calculated for the left and right hand separately for each subject. Repeated measures ANOVA with hand and numerosity as factors was performed on these values. As expected there was a main effect of numerosity ($F(1.6, 14) = 266, p < 0.0005$, Greenhouse-Geisser corrected values). There was no effect of hand ($F(1, 9) = 0.03, p = 0.86$), nor was there an interaction between hand and numerosity ($F(3.0, 27.3) = 1.8, p = 0.17$, Greenhouse-Geisser corrected values). To determine the subitizing and counting slopes the data from the two hands were collapsed and averaged over subjects (Figure 8.3a). The bilinear function (eq. 8.1) was fitted to the averaged response times weighted according to the inverse squared standard deviations ($R^2 = 0.90$). The transition point was found to be at 3.3 ± 0.2 items, which is in between 3 and 4 items as expected. The resulting subitizing and counting slope values were 0.20 ± 0.03 s/item and 1.2 ± 0.2 s/item, respectively. The uncertainties reported here indicate the SE of the fitting parameters and result directly from the fitting procedure. The slope values found here are comparable to the subitizing and counting slopes of 0.16 s/item and 0.84 s/items found in our previous study on haptic numerosity judgment (Plaisier, Bergmann Tiest, & Kappers, 2009a).

8.4.2 Bimanual trials

Again error rates were overall low (< 7 %) meaning that the bimanual trials were performed accurately. The response times averaged over subjects as a function of the total number of spheres is shown in Figure 8.3b. In this case several cluster combinations were possible to arrive at the same total number of spheres. Response times from combinations with the cluster size in the left and right hand reversed (e.g. 3 − 1 and 1 − 3) were collapsed. Comparison of the pattern in the response times to those predicted by

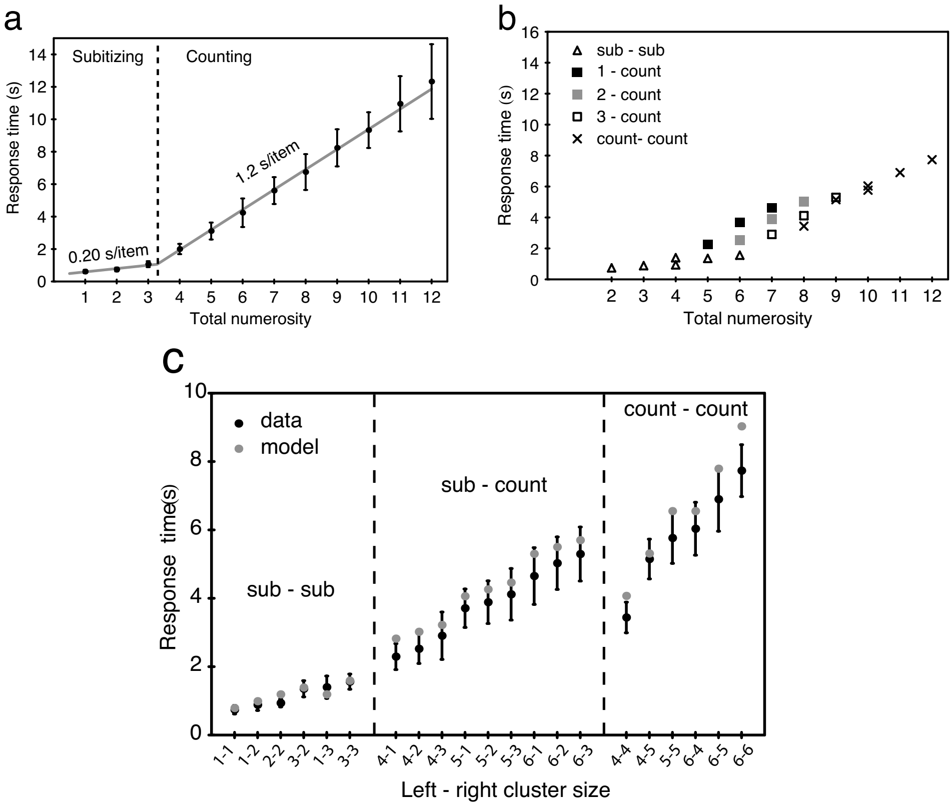

Figure 8.3: a) Response times averaged over subjects as a function of the number of spheres. The spheres were all presented to either the left or the right hand. The error-bars indicate the SD of the single subject means. The solid line represents weighted regression of the bilinear function to the response times. The resulting slope values are indicated in the figure. b) Response times averaged over subjects for the bimanual trials as a function of the total presented numerosity. Plot symbols correspond to those in Figure 8.2. c) Response times averaged over subjects as a function of the presented cluster combination (black dots). The error-bars indicate the SD of the distribution of the single subject means. The light gray dots indicate the predicted response time from the 'Cluster size dependent' model.

the models in Figure 8.3 suggests that the 'Cluster size dependent' model performs best. R^2 values were calculated to compare performance of the three models. Note that there were no free parameters in the models and therefore R^2 values can be negative and even smaller than -1. Negative values indicate that the data is better described by the mean than by the model. The largest possible R^2 value is 1, which means that the data follows the model exactly. For the 'No subitizing' model this yielded $R^2 = -14$, for the 'Total numerosity dependent' model this yielded $R^2 = -0.55$ and $R^2 = 0.93$ for the 'Cluster size dependent' model. This analysis clearly shows that indeed the 'Cluster size dependent' model performs best. The response times as a function of the cluster combinations and the model predictions are shown in Figure 8.3c. Considering the fact that there were no free parameters, the R^2 value is remarkably high which indicates that this model predicts the absolute response times very accurately. Also, the good performance of the model at describing the absolute response times shows that the time needed to sum the numerosities from both hands is indeed negligible. In fact, in can be seen in Figure 8.3c that when there were clusters from the counting range, the predicted response times were somewhat larger than the measured response times indicating that bimanual number processing occurs partially in parallel.

8.5 Discussion

From the bimanual trials it was concluded that the 'Cluster size dependent' model performs best. This not only shows that subitizing occurs for bimanual processing of numbers but also that each cluster is either subitized or counted depending on the cluster size. Both clusters are summed to arrive at the total. Consequently, clustering the items enables subitizing for up to six items and reduces response times for larger numerosities considerably compared to the unimanual case. Note that clustering the items also reduced response times for trials in which both cluster sizes were in the counting range, compared to the unimanual trials. An explanation for this is that the response times in the counting range for the unimanual case are a combination of subitizing 3 items and counting and adding the remaining items. In the bimanual case 3 items are subitized and the remaining items are counted for both hands. This way a total of 6 items were processed through subitizing and less items remained to be counted than in the unimanual case. This clearly demonstrates that enumeration of numerosities from the counting range is performed through a combination of subitizing and counting also when all items are in one hand.

The fact that subitizing is used in combination with counting has also been suggested in vision where a group-and-add procedure is found to be used (Van Oeffelen & Vos, 1982a, 1984). This fact complicates dissociation between the activated brain areas for both processes. In a brain imaging study by Piazza et al. (2002) no evidence was found for the existence of a neural network dedicated specifically to subitizing and that was not activated during counting. This is in agreement with the idea that subitizing is actually a sub-process of counting. Our results show that this is also the case in haptic numerosity judgment.

In conclusion, we have shown that response times are reduced considerably in bimanual number processing and that the subitizing range can be extended up to six items. Furthermore, we have shown that subitizing is involved in the processing of numerosities from the counting range. This shows that the group-and-add strategy found in vision is also used in touch. This provides strong evidence that numerosity is processed in a highly similar way for vision and active touch. Consequently, there is not only an influence of perception of numbers on hand motor action, but rather, numbers perceived through hand motor action are processed in the same way as visually perceived numbers. This suggests that magnitude representation is modality-independent.

Chapter 9

Haptic object individuation

Accepted for publication as:
Plaisier. M.A., Bergmann Tiest, W.M., and Kappers, A.M.L., Haptic object individuation, *IEEE Transactions on Haptics*, in press.

Abstract

Item individuation, i.e. how we decide which parts belong to one object and which to another, is an important aspect of haptic perception and can be crucial in, for instance, design of interfaces where different buttons have to be distinguished. We aim to provide insight into how objects grasped together in the hand are individuated. Subjects were asked to grasp varying numbers of shapes together in the hand and respond fast and accurately the number of shapes. First we investigated the effect of item size on numerosity judgment. In two other experiment the effects of heterogeneity in size and shape of the items were investigated. It was found that numerosity judgment in terms of response times, error rates and object handling was similar in all three experiments. Therefore, we conclude that size and shape features that are used for object recognition do not play a role in item individuation.

9.1 Introduction

We often hold several objects together in our hand and we can extract all sorts of information about these objects like shape, size, weight and the material they are made of. Another type of information that can be extracted is numerosity information, i.e. we can perceive how many objects we have in our hand. For judging how many objects we may hold in our hand it is necessary to individuate these objects. When we enclose an object with our hand it is in contact with different parts of our hand. When there are several objects enclosed in the hand we have to decide which parts belong to one object and which to another object. This is actually a quite complex task which we perform daily without even noticing. The question of how objects are individuated using touch is of course a fundamental one for understanding how haptic perception works. At the same time the outcome may also be of importance for applications. Object individuation can be crucial in situations where different buttons have to be distinguished, for instance on a touch screens. Although haptic object individuation is important, almost nothing is known about it. Objects features like size, weight or shape can be used to recognize a certain object e.g. (Plaisier, Bergmann Tiest, & Kappers, 2009b; Lederman & Klatzky, 2004, 1997, 1987; Klatzky et al., 1985). The same features may also play a role in individuation. In this study we aim to provide insight into the role of such features in haptic object individuation using a numerosity judgement task.

Although research into numerosity perception has been focussed on vision, numerosity perception in touch has been gaining attention recently (Riggs et al., 2006; Plaisier, Bergmann Tiest, & Kappers, 2009a). In numerosity judgment, two processes for exact numerosity judgment have been distinguished. For small sets of items of up to three or four items, a fast and accurate process known as ‘subitizing’ is used (Kaufman et al., 1949). For larger sets of items the more time consuming and error-prone process of ‘counting’ is used. Response times and error-rates as a function of the number of items increase rapidly in the counting range. Consequently, the transition from subitizing to counting is characterized by a sudden increase in error rates and the slope of the response times, usually at 3 or 4 items e.g. (Atkinson et al., 1976; Mandler & Shebo, 1982; Trick & Pylyshyn, 1993; Trick, 2008). Counting is believed to have a serial character and enumeration is achieved by adding each item to a running total. Subitizing is believed to have a more parallel character, because the response time per item is smaller. It is, however, still unclear what kind of a process subitizing exactly is. What is clear from vision, is that not all sets of items can be subitized. Concentric circles or squares cannot be subitized at all and for such stimuli counting is used over the whole numerosity range (Trick & Pylyshyn, 1994). The reason that subitizing does not occur, while the items can be accurately counted is believed to be that the individuation process is impaired.

In a recent study, we have shown that subitizing occurs in active touch (i.e. touch with active exploration) for up to three items when subjects were asked to judge the number of spheres grasped in their hand (Plaisier, Bergmann Tiest, & Kappers, 2009a). In that study, subjects judged a number of equally sized spheres. In daily-life, however, we often hold objects differing in size or shape in our hand and individuation of the objects might be a more complex task. Therefore, individuation of the different items might be impaired. On the other hand, individuation might also be facilitated in that case because these types of feature differences may make the items easier to distinguish. If individuation of the different objects is impaired for such heterogeneous sets of items held in the hand, it is expected that the subitizing slope value will be larger or that subitizing does not occur at all. If individuation is facilitated it is expected that the subitizing slope is smaller for heterogeneous sets of items.

In Experiment 1 we investigated whether the transition point from subitizing to counting is affected by the size of the items. When the items are very large they will not fit comfortably in the hand and this will ultimately affect the subitizing range or prevent subitizing from occurring at all. On the other hand, if the items are smaller than in our previous study on haptic subitizing (Plaisier, Bergmann Tiest, & Kappers, 2009a), the subitizing range may extend beyond three items. In Experiment 2 it was investigated whether items that are heterogeneous in size can be subitized. Finally, in Experiment 3 we investigated whether sets of items heterogeneous in shape can be subitized. If the subitizing slope or range is affected in these last two experiments, it can be concluded that shape and/or size features play a role in item individuation.

9.2 General Methods

9.2.1 Subjects

Ten paid participants (mean age 22 ± 2 yrs) took part in each of the experiments. Four of them were male and all subjects were right-handed according to Coren's test (Coren, 1993). None of them reported any known hand deficits. All subjects were naive as to the purpose of the experiment. They signed a declaration of informed consent.

9.2.2 Stimuli and set-up

The stimuli consisted of brass shapes suspended from flexible wires (Figure 9.1a). Three sizes of spheres were used. The smallest spheres had a diameter of 1.48 cm and mass of 14.5 g. The medium size spheres had a diameter and mass of 1.86 cm and 29 g, respectively. The largest spheres had a diameter of 2.34 cm and a mass of 58 g. Note that there is a factor of two in volume between the subsequent sphere sizes. In Experiment 1 subjects were presented with varying numbers of the small spheres. In Experiment 2

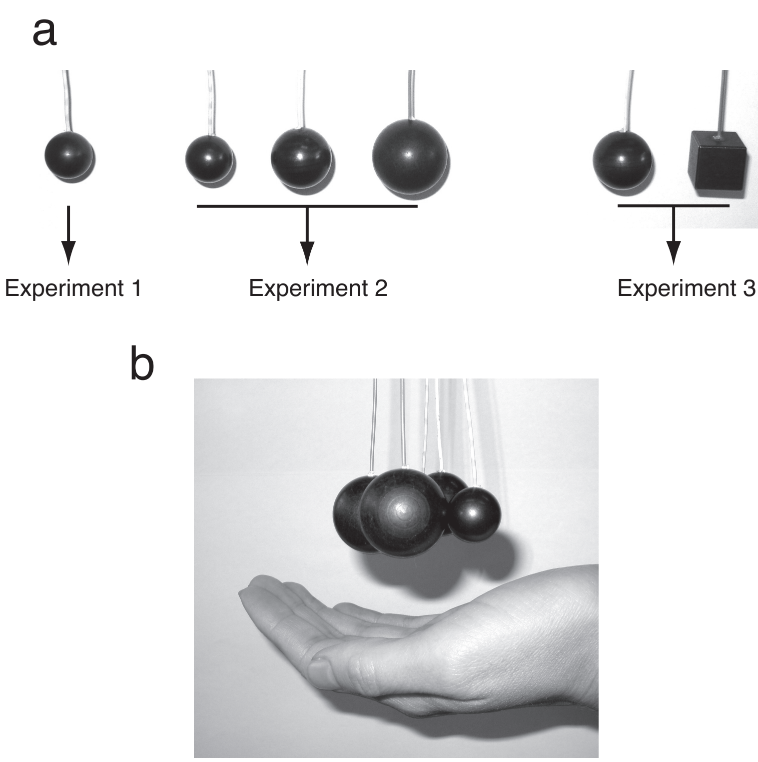

Figure 9.1: a) Picture the shapes used in each of the experiments. b) Picture of a subject reaching upwards to grasp a set of spheres heterogeneous in size from Experiment 2.

subjects were presented with a random combination of the small, medium and large spheres. In this heterogeneous case, the spheres never all had the same size. If the presented number was one, then randomly one of the sphere sizes was presented. Finally, in Experiment 3 random combinations of the medium spheres and cubes were used. The cubes had an edge length of 1.5 cm, such that the volume and consequently also the mass was the same as that of the medium size spheres. There were never sets of items presented that all had the same shape and in the case of one item randomly a sphere or cube was presented. The total number of items presented could be 1, 2, 3, 4, 5, 6, 7 or 8 in each of the experiments.

Response times were measured using a custom built device. Time measurement was started automatically when a subject touched the stimuli and it was terminated through a vocal response. The response times were measured with an accuracy of 10 ms. For technical details about this device see (Plaisier et al., 2008b).

9.2.3 Procedure

Subjects were blindfolded and wore earplugs to avoid sound cues. They were instructed to grasp the stimuli with their dominant hand and respond the correct number of items as fast as possible (Figure 9.1b). There were no restrictions on exploration strategy nor on hand movements, other than having to initially grasp all items simultaneously. After grasping all items they were allowed to release items from their hand during the remainder of the trial and the experimenter recorded whether or not items were released from the hand. Whenever an incorrect response was given, the experimenter informed the subject of what the correct number of items was.

All subjects performed each of the experiments in roughly counterbalanced order. Each experiment was performed in a separate block of trials with a duration of approximately 50 minutes. The blocks were performed on different days or with a break of at least an hour in between. Before the experiment was started, a minimum of 20 practice trials were performed until 10 in a row were performed correctly. It was never necessary to exceed 30 practice trials. Each subject performed 10 trials per number of items and incorrectly answered trials were repeated at the end of the block to ensure 10 correctly performed trials for each numerosity. Incorrectly answered trials were discarded from the analysis of the response times and the item release data. Error rates were calculated as the percentage of incorrectly answered trials of the total performed number of trials.

9.2.4 Analysis

Response times

Strategy shifts in cognitive tasks can be detected by regression of a model consisting of multiple linear parts with unknown change points to, for instance, response times (Luwel et al., 2001). In our previous study on haptic numerosity judgment we have used regression of a bi-linear model to determine the subitizing and counting slopes as well as the transition point between both regimes (Plaisier, Bergmann Tiest, & Kappers, 2009a). The same method was used in the present study. The bilinear function for the response time T as a function of the number of items N consisting of two linear regimes with a discrete transition is given by:

$$T(N) = (r_1 N + c_1) H\left(\frac{c_2 - c_1}{r_1 - r_2} - N\right) + \\ + (r_2 N + c_2) H\left(N - \frac{c_2 - c_1}{r_1 - r_2}\right) \tag{9.1}$$

Here $H(N)$ is the Heaviside step function and r_1 and r_2 are the subitizing and counting slopes, while c_1 and c_2 represent constant offsets. Because there may be individual differences for the slope values and the transition point, regression of this function was performed on the response times averaged over subjects as well as on the response times from each individual subject. The largest numerosity (8) was not included in the analysis, because of possible end-effects. Subjects usually notice what the largest numerosity is and response times for this number are therefore generally smaller than predicted from the bi-linear model. It is common to exclude the largest numerosity in order to get a more accurate estimate of the counting slope in numerosity judgment studies e.g. (Trick & Pylyshyn, 1993; Watson et al., 2007).

Item release rates

During the experiment the experimenter scored for each trial whether items were released from the hand. Typically this rate starts at 0% for small numerosities and increases (usually up to 100%) for larger numerosities. Releasing items from the hand can be interpreted as an indication that items were processed serially. We have shown that when subjects had to search for a cube among spheres, which is an example of parallel search, they would sporadically release items from their hand, whereas when they had to search for an ellipsoid among spheres (an example of serial search) they would almost always release items from their hand (Plaisier, Bergmann Tiest, & Kappers, 2009b). Therefore, in the counting range we expect subjects to release items from their hand more often than in the subitizing range. Note that it does not necessarily mean that processing was parallel when no items are released from the hand. However, when items were released it clearly indicates that processing was not parallel. To analyze these data a cumulative gaussian was fitted to the item release rates as a function of the number of items. The cumulative gaussian is given by:

$$\text{cdf}(x) = 50 + 50 \text{ erf}\left(\frac{x-\mu}{\sqrt{2}\sigma}\right) \tag{9.2}$$

where erf is the error function, μ is the mean of the distribution and σ is the standard deviation. The value of μ indicates the 50% interval and is a measure for the transition from numerosities for which no items are released to numerosities for which items are almost always released from the hand. The value of σ indicates the steepness of the curve and indicates at what rate subjects switched from not releasing items from the hand to always releasing items. The smaller the value of σ the more accurate μ is determined. By fitting the cumulative gaussian to the single subject's data the values of μ and σ were determined for each subject. An example of the

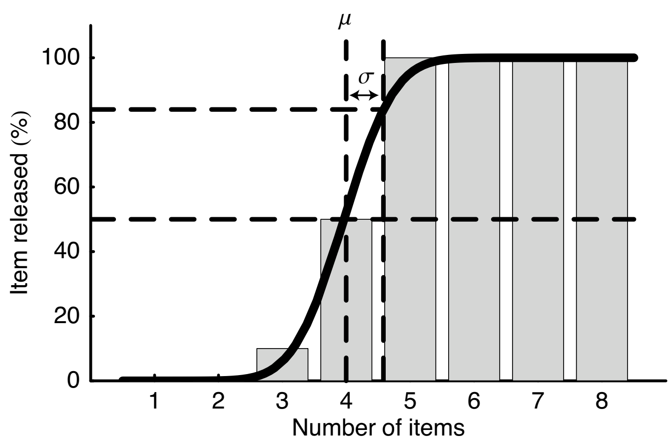

Figure 9.2: The bars indicate the item release rate as a function of the number of items from one subject in Experiment 1. The solid line indicates the fitted cumulative gaussian and μ and σ are indicated.

item release data from one subject in Experiment 1 is shown in Figure 9.2. There is no theoretical reason to expect that the distribution of the items release rates is symmetrical. However, the goodness of fit values found in the different experiments indicate that the cumulative gaussian approximates the data fairly well.

9.3 Experiment 1: Effect of item size

In this experiment we investigated the effect of sphere size on the subitizing range. If the subitizing range of up to 3 items reported in our previous study (Plaisier, Bergmann Tiest, & Kappers, 2009a) is limited by the size of the items we expect that this range will be extended up to larger numbers when the presented items are smaller. If the items are smaller they can be comfortably held in the hand up to larger numerosities. If the subitizing range is not extended for smaller items, we can conclude that the upper bound of the subitizing range is not limited by bio-mechanical constraints determining how many items can be held in the hand.

9.3.1 Methods

Subjects were presented with 1, 2, 3, 4, 5, 6, 7 or 8 small spheres (Figure 9.1a). Instructions were as described in the General Methods.

9.3.2 Results

The response times and error rates averaged over subjects as a function of the number of items are shown in Figure 9.3a. Repeated measures ANOVA on the single subject response times showed a significant effect of numerosity ($F(1.6, 14) = 114, p < 0.0001$) and trend analysis showed significant deviations from linearity ($p < 0.0001$). Regression of the bilinear model (eq. 9.1) on the averaged response times (the solid line in Figure 9.3a) yielded the transition point from subitizing to counting in between 2 and 3 items. The resulting subitizing and counting slopes were 0.18 s/item and 1.1 s/item, respectively. Regression of this model to the single subject response times showed that the transition was on average at 2.6 ± 0.3 (SD) items ($R^2 > 0.97$). The subitizing slope was on average 0.2 ± 0.1 (SD) s/item and the counting slope 1.1 ± 0.3 (SD) s/item.

In Figure 9.3a it can be seen that error rates are overall low (< 8 %) and no errors were made up to 3 items. Figure 9.3b shows a confusion matrix.

On the one axis the presented numerosity is indicated and on the other axis the perceived numerosity. The gray scale indicates the percentage of trials in which the presented numerosity was perceived as a certain other numerosity. The correctly performed trials are therefore on the diagonal. Note that most trials were performed correctly and that both under- and over-estimations occurred.

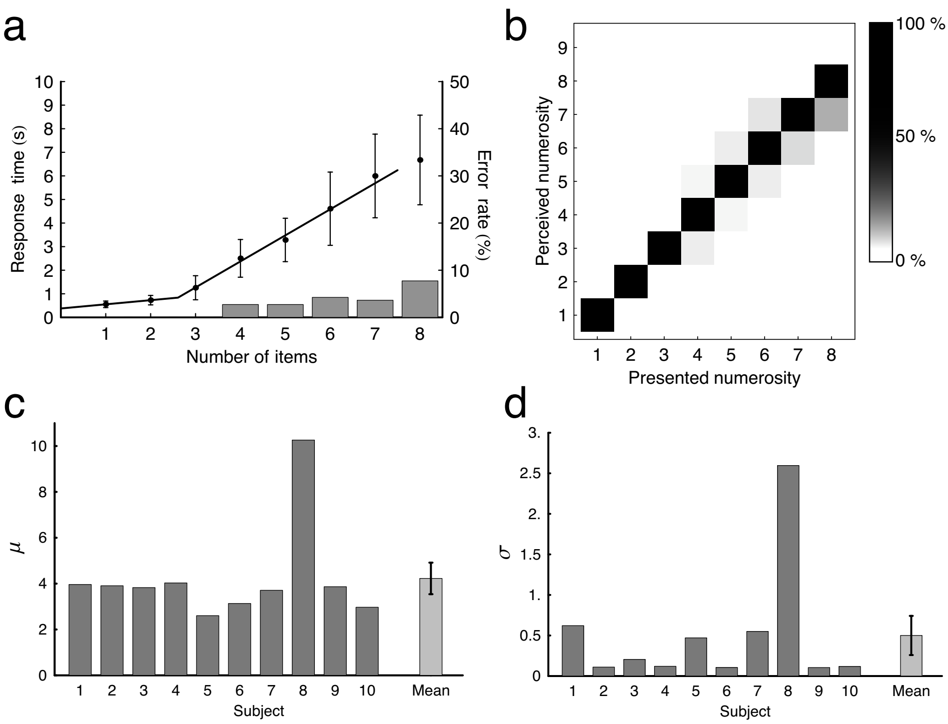

Figure 9.3: Experiment 1: Effect of item size. a) Response times (dots ± SD) and error rates (gray bars) as a function of the number of items averaged over subjects. The solid line represents regression of the bilinear function. b) Confusion matrix. The gray scale indicates how often a presented numerosity was perceived as a certain numerosity. c) Numerosity for which items were released from the hand in 50% of the time, i.e. the location of the mean of the cumulative gaussian. The last bar is the average over subjects and the error bar indicates the standard error. d) Rate at which the percentage of trials in which items were released increases with the numerosity, i.e. the standard deviation of the cumulative gaussian. The last bar is the average over subjects and the error bar indicates the standard error.

As described in the General Methods section cumulative gaussians (eq. 9.2) were fitted to the item release data of each subject. The quality of fits was good overall ($R^2 > 0.78$). In Figure 9.3c and d the values for μ and σ of the item release data are shown for all subjects. It can be seen that for most subjects the value of μ was around 4 items, but there was one subject for which this value was substantially larger. Note that since the maximum numerosity was 8 and this subject did not reach an item release rate of more than 20% for any numerosity, the indicated value is based on extrapolation. Also the value of σ was largest for this subject. This subject did, however, show the same pattern in the response times and error rates as the other subjects. So, apparently this subject counted the items without releasing them from the hand.

9.3.3 Discussion

The results clearly indicate that subitizing occurred in this experiment. Based on the response times together with the error rates and the item release data it can be concluded that the transition from subitizing to counting was located at 2 or 3 items depending on the subject. This shows that the subitizing range was certainly not extended up to larger numbers by using smaller items. Also the subitizing and counting slopes are comparable to the subitizing and counting slope values of 0.16 s/item and 0.8 s/items reported in our previous study (Plaisier, Bergmann Tiest, & Kappers, 2009a). We can conclude that the upper bound of the subitizing range is not determined by bio-mechanical constraints. Thus there must be another reason for the limit of the subitizing range. Most likely there is a cognitive reason. This is also supported by the fact that the subitizing range extends to roughly 3 items in both vision (Mandler & Shebo, 1982; Trick & Pylyshyn, 1993) and touch (Riggs et al., 2006; Plaisier, Bergmann Tiest, & Kappers, 2009a).

9.4 Experiment 2: Size heterogeneity

Again subjects were instructed to enumerate varying numbers of spheres, but in this case the sets of spheres were heterogeneous in size. Note that because all spheres were made out of the same material, the mass varied with the volume. If subitizing occurs and response time slopes, error rates and item release data are similar to that in Experiment 1 we can conclude the mass and volume cues do not play a role in haptic object individuation.

9.4.1 Methods

Subjects were presented with 1, 2, 3, 4, 5, 6, 7 or 8 spheres heterogeneous in size (Figure 9.1a). Sphere sizes and instructions were as described in the General Methods.

9.4.2 Results

Figure 9.4a shows response times and error rates averaged over subjects as a function of the number of items. Repeated measures ANOVA on the single subject's response times showed a significant effect of numerosity ($F(1.2, 10.8) = 99, p < 0.0001$). Trend analysis showed significant deviations from linearity ($p < 0.0001$). Regression of the bilinear model on the averaged response times (the solid line in Figure 9.4a) yielded the transition point from subitizing to counting in between 2 and 3 items items. The resulting subitizing and counting slopes were 0.18 s/item and 1.2 s/item, respectively. Regression of this model to the single subjects response times showed that the transition was on average at 2.5 ± 0.3 (SD) items ($R^2 > 0.97$). The subitizing slope was on average 0.2 ± 0.1 (SD) s/item and the counting slope 1.2 ± 0.2 (SD) s/item.

As can be seen in Figure 9.4a no errors were made up to 3 items and error rates are overall low (< 11 %). A confusion matrix is shown in Figure 9.4b. As in the previous experiment most trials were performed correctly and under-estimation as well as over-estimation occurred.

The values for μ and σ of the item release data are shown for all subjects in Figure 9.4c and d ($R^2 > 0.93$). It can be seen that again for most subjects the value of μ was around 4 items. The values of μ and σ are similar to those found in Experiment 1. Again subject number 8 had a much larger value of μ than the other subjects and in this case the largest item release rate over the whole range was 80% for this subject.

To analyze effects of the ratio of small, medium and large spheres in the different combinations summing up to the same total number of items, the response times were split up according to sphere size which was most numerous. To ensure a reasonable number of trials for all possible sphere size combinations the data from the different subjects were collapsed. In Figure 9.5 response times for trials with mostly small spheres, mostly medium spheres and mostly large spheres are shown. It can be seen in this figure that for one item the response times for each sphere size are almost identical. Combinations for which there was not one sphere size that was most numerous were not included. Note that not all of the three sphere sizes needed to be in the set. Consequently, more trials are included for the odd numbers of items than the even numbers of items. Roughly equal numbers of trials were included for each of the spheres sizes. Averaged over all numerosities in 40% ± 7% (SD) of the included trials most items were small spheres, in 30% ± 3% of the trials most items were medium spheres and in 31% ± 8% most items were large spheres. The response times for one item were not

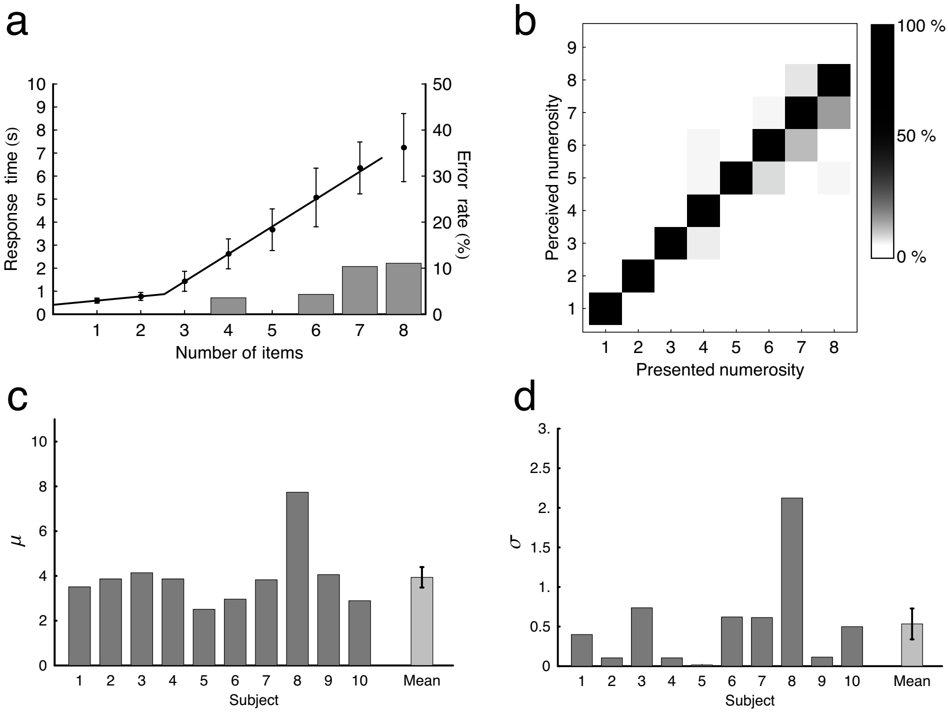

Figure 9.4: Experiment 2: Effect of size heterogeneity. a) Response times (dots ± SD) and error rates (gray bars) as a function of the number of items averaged over subjects. The solid line represents regression of the bilinear function. b) Confusion matrix. The gray scale indicates how often a presented numerosity was perceived as a certain numerosity. c) Numerosity for which items were released from the hand in 50% of the time, i.e. the location of the mean of the cumulative gaussian. The last bar is the average over subjects and the error bar indicates the standard error. d) Rate at which the percentage of trials in which items were released increases with the numerosity, i.e. the standard deviation of the cumulative gaussian. The last bar is the average over subjects and the error bar indicates the standard error.

included in the analysis, because in this case there was only one sphere size presented and not a combination of sphere sizes. In Figure 9.5 there are no clear differences visible between the different item combinations. Repeated measures ANOVA with total numerosity as the repeated factor showed no effect of the most numerous sphere size ($F(2, 10) = 1.1, p = 0.4$).

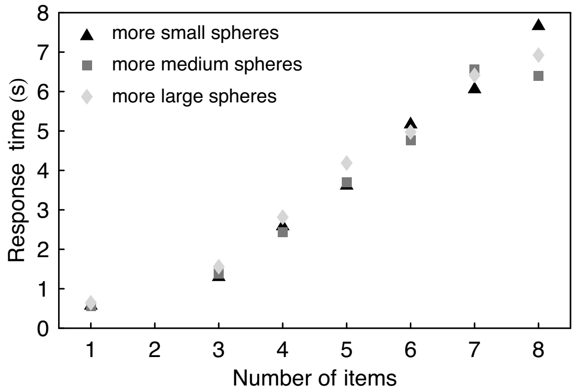

Figure 9.5: Experiment 2: Effect of size heterogeneity. Response times as a function of the number of items averaged over all subjects. Triangles indicate that most items were small spheres, squares indicate that most were medium spheres and diamonds indicate that most items were large spheres.

9.4.3 Discussion

The results show that subitizing is possible up to 2 or 3 items for sets of items that vary in volume and mass. Note that the volume and mass differences were considerable here; a factor of 4 between the small spheres and the large spheres. For example, two items could be comprised of 1 small sphere and 1 medium sphere, but also of 1 medium sphere and 1 large sphere. There is a difference of a factor of 2 in volume and mass between these two combinations of items that each sum up to two items. For larger numerosities the volume differences between different combinations of the three sphere sizes can even be much larger. Clearly, variations in volume and mass do not affect the highly efficient individuation mechanism that is necessary to enable subitizing.

9.5 Experiment 3: Shape heterogeneity

In this last experiment, we investigated whether shape heterogeneity affects haptic object individuation. In this case subjects were presented with a combination of spheres and cubes. In a previous study, we have shown that a sphere among cubes or vice versa is very salient (Plaisier, Bergmann Tiest, & Kappers, 2009b). This indicates that cubes and spheres are highly distinguishable and therefore we specifically chose to present a combination of these two shapes.

9.5.1 Methods

Subjects were presented with a mixture of spheres and cubes (Figure 9.1a). The total numbers of shapes presented were 1, 2, 3, 4, 5, 6, 7 or 8. Instructions were as described in the General Methods.

9.5.2 Results

The response times averaged over subjects as a function of the number of items are shown in Figure 9.6a. Repeated measures ANOVA on the single subject response times showed a significant effect of numerosity ($F(1.3, 12) = 114, p < 0.0001$) and trend analysis showed significant deviations from linearity ($p < 0.0001$). Regression of the bi-linear model to the single subject response times showed that the transition was on average at 2.6 ± 0.3 (SD) items. The subitizing slope was on average 0.2 ± 0.2 (SD) s/item and the counting slope 1.1 ± 0.3 (SD) s/item ($R^2 > 0.97$). Regression of the bi-linear model on the averaged response times (the solid line in Figure 9.6a) yielded the transition point from subitizing to counting in between 2 and 3 items. The resulting subitizing and counting slopes were 0.18 ms/item and 1.1 ms/item, respectively.

It can be seen in Figure 9.6a that no errors were made up to 3 items and error rates were overall low ($< 11\%$). The confusion matrix shows again that most trials were performed correctly and that both under- and over-estimations occurred (Figure 9.6b).

In Figure 9.6c and d the values for μ and σ of the item release data are shown for all subjects ($R^2 > 0.78$). Again the value of μ was around 4 items for most subjects and overall the values of μ and σ were similar to those found in Experiments 1 and 2. The subject with the larger values for μ and σ than the other subjects was the same as in the the previous two

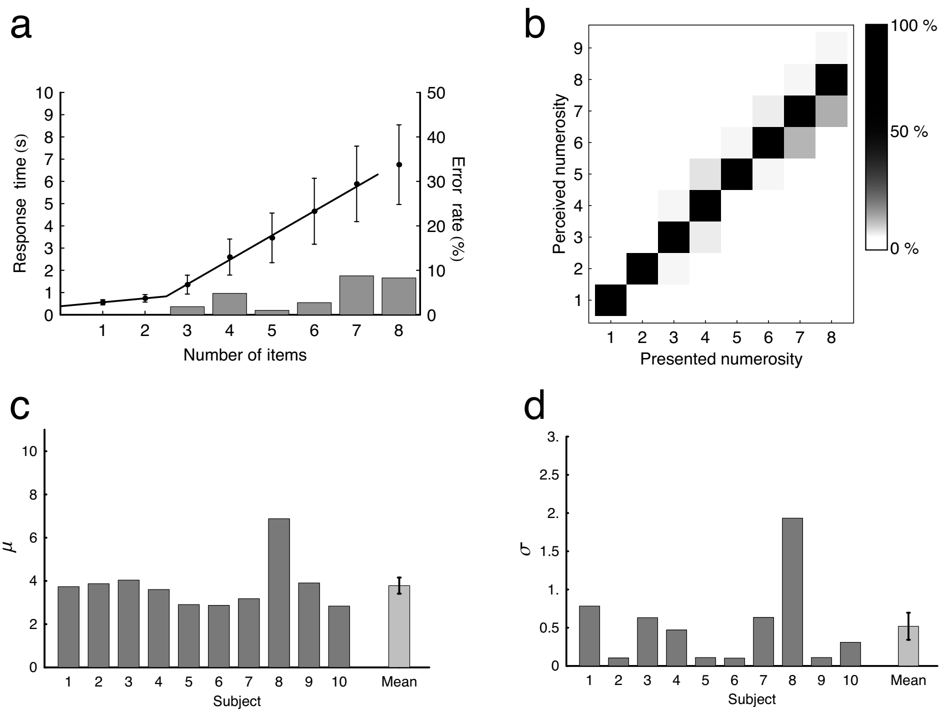

Figure 9.6: Experiment 3: Effect of shape heterogeneity. a) Response times (dots ± SD) and error rates (gray bars) as a function of the number of items averaged over subjects. The solid line represents regression of the bilinear function. b) Confusion matrix. The gray scale indicates how often a presented numerosity was perceived as a certain numerosity. c) Numerosity for which items were released from the hand in 50% of the time, i.e. the the location of the mean of the cumulative gaussian. The last bar is the average over subjects and the error bar indicates the standard error. d) Rate at which the percentage of trials in which items were released increases with the numerosity, i.e. the standard deviation of the cumulative gaussian. The last bar is the average over subjects and the error bar indicates the standard error.

experiments. In this case the largest item release rate for this subject was 50%.

We analyzed whether there was an effect of combinations with more cubes than spheres compared to those with more spheres than cubes. Like in Experiment 2 the data from all subjects were collapsed and the data were split up according to for which shape there were more items. The

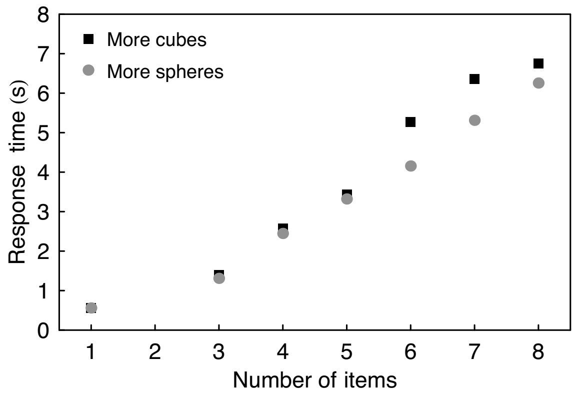

Figure 9.7: Experiment 3: Effect of shape heterogeneity. Response times as a function of the number of items averaged over all subjects. Squares indicate that there were more cubes than spheres and disks indicate there were more spheres than cubes

average response times for each total numerosity for trials with more cubes and those with more spheres are shown in Figure 9.7. It can be seen that response times at one item are similar for both shapes. Furthermore, it can be seen that there is a difference mainly for the larger numerosities between both shape combinations. The combinations with an equal number of spheres and cubes were not included in the analysis (hence there is no response time for 2 items in Figure 9.7). Also the response times for 1 item were not included in the analysis, because there was always either a sphere or a cube and not a combination of both shapes. Averaged over all numerosities in $50\% \pm 4\%$ (SD) of the trials there were more spheres and in $50\%\pm4\%$ there were more cubes. A paired t-test (response times were paired according to total numerosity) showed that response times for combinations with more cubes were longer and this difference was marginally significant $(t = 2.5, p = 0.051)$.

9.5.3 Discussion

Our data show that subitizing occurs for sets of items that are heterogeneous in shape. Note that we chose spheres and cubes, because these are very distinguishable shapes. So, there was no question whether subjects did perceive the shape differences. This indicates that the individuation mechanism that allows subitizing does not rely on shape cues. Response times were, however, for the largest numerosities somewhat larger when there were more cubes than spheres. This is in agreement with what we have found in a previous study where response times in a search task were larger when the distractor items were cubes compared to when they were spheres (Plaisier, Bergmann Tiest, & Kappers, 2009b). One explanations for this is that cubes do not slide as easily along each other as spheres do and are therefore more difficult to rearrange in the hand. This is consistent with the finding that there were mainly differences in the response times for large numbers of shapes, because for large numerosities subjects release items from the hand and rearrange the items in the hand.

9.6 Overall analysis

The results from all three experiments are very similar in terms of response time slopes, error rates and item release data. Therefore, we performed a statistical analysis on the data from all three experiments to investigate whether this showed any differences between the experiments. Repeated measures ANOVA with experiment as a factor was performed on the subitizing slope, the counting slope and the transition point from each subject. There was no effect of experiment for any of these values ($F(2,18) \leq 1.02, p \geq 0.38$). Also the error rates averaged over all numerosities were compared using repeated measures ANOVA. Again there was no effect of experiment ($F(2,18) = 0.17, p = 0.85$). Finally, the μ and σ values from the item release data were compared across experiment. Also for both of these values no effect of experiment was found ($F(1.10, 10.2) = 1.1, p = 0.35$ for μ and $F(1.3, 11.4) = 0.3, p = 0.6$ for σ).

9.7 General Discussion

Our results from Experiment 1 show that the subitizing range does not increase when the items are smaller. This indicates that the location of the transition point from subitizing to counting is not determined by item size. Furthermore, our results show that subitizing also occurs when the items are heterogeneous in size (Experiment 2) or shape (Experiment 3). The extent of the subitizing range and item release rate as well as the error rates show that performance was very similar across all three experiments. In fact, statistical analysis of the results from all of the experiments did not reveal any differences. This indicates that enumeration was not affected by item heterogeneity. Therefore, object individuation was not impaired or facilitated by size or shape heterogeneity. This means that the highly efficient individuation mechanism that has to be available for subitizing to occur does not rely in any way on size or shape features.

This raises the question of what information is used to individuate objects. One possible candidate is whether or not the parts touching the skin can move relative to each other. If they can, they are likely to belong to different objects, whereas if the cannot, chances are that they belong to the same object. In a previous study we have shown that being able to slide items along each other and rearrange them in the hand does not facilitate object recognition (Plaisier, Kuling, et al., 2009). In that study, cubes and spheres were fixed in a grid such that subjects could enclose all items with the hand, but not rearrange them. Their task was to find, for instance, a cube among spheres. The introduction of small object movements did not facilitate extraction of object features like edges. However, in the present study we found that object individuation does not rely on such object features. Therefore, the findings of the previous study do not preclude the possibility that moving objects in the hand may play an important role in object individuation. If the objects were fixed in space, however, they would be spatially distributed and this may also be used for object individuation. Spatial distribution and object motion will be difficult to disentangle, but both are likely to play a role.

In conclusion, haptic object recognition has been shown to be accurate and fast (Klatzky et al., 1985; Klatzky & Lederman, 1995). It relies on features like shape, material and size (Plaisier, Bergmann Tiest, & Kappers, 2009b; Lederman & Klatzky, 1997; Plaisier et al., 2008a). In daily life, when we have a set of keys grasped together in our hand, then a key that has a different shape from the others is easy to recognize. So although we may use shape and size features to recognize objects, these features do not play a role in numerosity perception. Shape and size features may be used to find the correct key to open your front door, they are not used for determining how many keys you have in your hand.

Chapter 10

Summary and Conclusions

In this thesis, a series of investigations into how the haptic system processes information about multiple objects was presented. In order to do so, we have successfully extrapolated the visual search and numerosity judgement paradigms to the domain of active touch. Furthermore, two new types of haptic stimuli were introduced. The first consisted of a surface on which items could be placed, whereas the second consisted of multiple objects that can be simultaneously grasped in the hand. In both cases we have shown that analysis of exploratory strategy is important for interpreting response times. We have introduced a way to analyse these exploratory movements and used numerical simulations to validate this approach. Finally, we have modelled the response times in order to draw conclusions about the underlying processes. The conclusions drawn from each chapter are summarised below.

10.1 What have we learned about haptic search?

We have shown that, similar to visual search, for haptic search tasks there is a range of search efficiencies. In Chapter 2 blindfolded subjects were instructed to search for a target item which differed in roughness from the surrounding distractor items. Items were distributed over a plane. In this case we found response time slopes as low as 20 ms/item when the target was rougher than the distractor items. In the reversed situation (i.e. the target was less rough than the distractor items), the slopes increased to 260 ms/item, indicating a search asymmetry. We also showed that these differences in search slope were accompanied by changes in search strategy. In the condition where the search slope was small, a single hand sweep over the display was sufficient to determine whether a target item was present. In the less efficient condition, a more detailed search strategy was used. Interpretation of absolute search slopes is difficult especially in the absence of typical values of such slopes in the literature. Our study shows that analysis of the exploratory hand movements provides insight into the search strategy. By relating haptic search slopes to parallel and serial search strategies we showed, for the first time, that pop-out effects occur under free manual exploration.

In the case of a serial search strategy, subjects tended to search parts of the display repeatedly. This indicates that subjects were uncertain about which parts of the display had been searched. Under such conditions could adding visual spatial information about the display facilitate haptic search considerably. This was investigated in a controlled way in Chapter 3 using a set-up that incorporated a force-feedback device to display the haptic stimulus spatially aligned with a visual stimulus displayed using a screen and mirror. In this case, the haptic display was explored using the index finger only. Subjects were instructed to determine whether a target was present among distractor items. While the target was recognisable only haptically, visual information on finger position or possible target positions could be given. The results show that subjects could use visual information on possible target positions even in the absence on feedback on finger position. When there was no feedback on possible target locations, subjects scanned the whole display systematically. When feedback on finger position was present, subjects could make well-directed movements back to areas of interest. This was not the case without visual information on finger position, indicating that showing finger position helps to form a spatial representation of the display. These findings are important for the design of tele-operation systems. In such systems, a robot is operated from a distance, often using a combination of haptic information presented using a force-feedback device and visual information presented on a screen. These systems can be used for exploring environments which are hazardous to humans, or as surgical robots for minimal invasive surgery. Our results show that visual information on the scene can be used to guide exploration even when the probe is not visible. On the other hand, there is also an advantage of showing the position of the probe even if visual information on the scene is poor because the camera image is blurred.

Furthermore, in Chapter 3 we show that response time models of visual serial search do not generally apply for haptic serial search. In the case of no visual feedback the response time slopes were not significantly different from zero, while search could only be performed serially. This was due to the fact that subjects were unaware of the locations of the items until they touched them. Therefore, they systematically scanned the whole display and response times were roughly independent of the number of items on the display. This shows that one should be cautious when interpreting haptic search times in terms of visual search models. Already in Chapter 2 we showed that it is important to interpret haptic search slopes in combination with the observed exploratory strategy.

So far, items were distributed over a plane, but in daily-life we often hold several three-dimensional objects grasped in our hand. In this case shape can be an important cue for recognising objects by touch. Several shape features like edges, curvature, surface area and aspect ratio are associated with three-dimensional shape. In Chapter 4, we investigated saliency of three-dimensional shape features when several shapes are grasped together in the hand. The target and distractor items consisted of shapes (cube, sphere, tetrahedron, cylinder, ellipsoid) which differed in several shape features. The results showed that this type of haptic search task can be performed very efficiently (25 ms/item) when the target shape had edges while the distractor shapes did not or vice versa. This indicates that edges are a highly salient shape feature. Furthermore, such very salient local features can be perceived through enclosure, an exploratory procedure usually associated with global shape. Since subjects had to perform the task as fast as possible, this suggests that speed may be a factor in selecting the appropriate exploratory procedure.

The stimulus design of Chapter 4 allowed for the items to be rotated and rearranged in the hand. Such active control over item position and orientation is typical for haptic search, but is never possible in visual search. In Chapter 5, we investigated how this freedom to manipulate the objects in the hand affects three-dimensional shape perception. Furthermore, we investigated differences in detection speed for different positions in the hand while grasping objects. To this end, subjects were instructed to search for a cube among spheres or for a sphere among cubes and response times were measured for different locations of target shape in the hand. The way in which the items were fixed to a certain position was varied from allowing small displacements and rotations of the shapes to rigidly fixed. There were only differences in search times between the different positions in the hand, when the centre item was difficult to access because of the surrounding items. It was also shown that search was faster when the items were rigidly fixed than when small displacements and rotation was possible. An explanation for this difference is that it was easier to exert forces on the rigidly fixed items than on the partly fixed items.

10.2 What have we learned about haptic numerosity judgement?

In vision it is generally found that numerosity judgement is fast and accurate for small numbers of items, whereas response times and error rates increase rapidly for larger set-sizes. The process used to judge small numerosities is known as subitizing. In Chapter 6, we investigated numerosity judgement through active touch. Subjects were asked to grasp a set of spheres in their hand and judge the numerosity. It was found that enumeration for up to three items was error-free and more efficient (167 ms/item) than for larger numbers of items (839 ms/item). In the same chapter, we also show that enumeration of small numerosities was not performed through accurate estimation enabled through the large relative differences between subsequent numerosities in this regime. Furthermore, it was shown that numerosity information was accessed directly and not through mass or volume cues. Not only do these results show that a haptic version of subitizing exists in active touch, they also suggest similar underlying enumeration mechanisms across different modalities. To investigate this further, a visual version of this study was carried out.

In Chapter 7, we showed that also in vision, judgement of small numbers is faster than for large numbers even if the relative differences between subsequent numerosities is large over the whole numerosity range. This in agreement with the idea that subitizing is not the same as very accurate estimation. Furthermore, we show that numerosities from the subitizing range are dissimilar from numerosities outside this range. This means that for some reason, small numbers are recognised faster than larger numbers, regardless of relative differences between the numerosities. In vision it has been suggested that pattern recognition might enable this fast recognition of small numerosities. This is, however, an unlikely explanation in the case of haptic numerosity judgement, as in that case the items were not fixed to a position. Since the results were similar in both modalities it seems probable that a similar reason underlies the occurrence of subitizing in both modalities.

In vision it has been suggested that subitizing is used in combination with counting. A field of dots is judged by grouping together small numbers of dots and these groups are enumerated separately and added together to arrive at the total. If numerosity judgement is performed through a similar process in haptics as in vision, it is likely that such a group-and-add strategy is also used in haptic numerosity judgement. This was investigated in Chapter 8. In that study, the spheres were split into two groups by presenting a set of spheres to each hand simultaneously. We compared enumeration of a set of spheres presented to one hand and enumeration of the same total number of spheres presented divided over the two hands. Our results showed that, like in vision, a combination of subitizing and counting is used to process numbers in active touch. This again shows that numbers are processed in a similar way in touch as in vision. This suggests that number representation is modality-independent.

So far, we showed that subitizing occurs when subjects have to judge a number of identical spheres. However, in daily life objects vary in size and shape. Before the number of items can be judged, the items have to be individuated, i.e. decide which part belongs to one object and which to another. For subitizing to occur, individuation must be fast. In Chapter 9, we investigated whether item individuation is affected by item size and shape. Subjects were asked to grasp varying numbers of shapes together in the hand and again respond fast and accurately the number of shapes. First, we investigated the effect of item size on numerosity judgement by using smaller spheres than were used in Chapter 5. In this case, the subitizing range was not extended to larger numerosities. This shows that the upper limit of the subitizing range is not only determined by how many items can comfortably fit in the hand. In two other experiments, the effects of heterogeneity in size and shape of the items were investigated. In this case subjects had to enumerate a mixture of spheres and cubes, or spheres of different sizes. It was found that numerosity judgement in terms of response times, error rates and object handling was comparable in all three experiments. This indicates that size and shape features that are used for object recognition do not play a role in item individuation.

10.3 What have we learned about haptic perception of objects?

In the Introduction it was pointed out that this thesis was aimed at providing insight into object saliency as well as object individuation. With respect to object saliency we have shown that there are certain object features like roughness or the presence of edges that can be detected without having to explore all objects in the set sequentially. When a target differs from the distractors in terms of such a feature, search can be performed in parallel. The fact that subitizing occurs when three-dimensional objects are grasped in the hand shows that object individuation is fast and efficient. Assuming that counting is a purely serial process, we can conclude that there is a parallel component to subitizing as it is more efficient than counting. Because object individuation has to be performed before objects can be enumerated, the occurrence of subitizing suggests that object individuation of at least small numerosities can be performed in parallel. So, both feature extraction and object individuation can be fast and efficient. This means that when we hold multiple objects the haptic system can detect the presence of a salient feature, but also that the haptic system can determine that this feature belongs to one of several objects without having to explore all objects sequentially.

In the last chapter of this thesis it was shown that object individuation is not affected by heterogeneity in shape or size of the objects. This suggests that objects are not individuated by coupling different features to different objects. Consequently, feature extraction and object individuation can be performed independently. In conclusion, we have learned about haptic perception of objects that object recognition and object individuation can both be fast, but that these are independent processes.

10.4 Afterthoughts and some speculations

The results from this thesis lead to some interesting new questions with respect to haptic object representation. Object recognition and individuation may be performed independently, but to form a full object representation, features will have to be coupled to the correct object. New experiments will be necessary to answer the question of how this coupling is implemented. Inevitably this line of research will lead to the highly ambitious question of “What defines a haptic object?". How do we decide that there are several small objects instead of one large object? Based on the results from this thesis, my guess is that congruency of the movements of the different parts of an object plays an important role in this process. Such a mechanism would allow object individuation regardless of feature differences between the objects. Further research is necessary to see whether this speculation is true and to gain insight into how different parts of a single object are integrated into one haptic object representation.

It may seem that at the end of this thesis, we are left with more questions than we started out with. Note, however, that these questions were already there, but it is because of the results from this thesis that these questions have become apparent. The innovative experimental paradigms and stimulus design in this thesis open up a field of research that is much too large to cover in one thesis. Although we are nowhere near answering the question of what defines a haptic object yet, at the end of this thesis we are definitely a step closer.

References

Akaike, H. (1974). New look at the statistical model identification. *IEEE Transactions on Automatic Control*, *AC-19*, 716–723.

Akin, O., & Chase, W. (1978). Quantification of three-dimensional structures. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 397–410.

Alston, L., & Humphreys, G. W. (2004). Subitization and attentional engagement by transient stimuli. *Spatial Vision*, *17*, 17–50.

Andres, M., Davare, M., Pesenti, M., Olivier, E., & Seron, X. (2004). Number magnitude and grip aperture interaction. *NeuroReport*, *15*, 2773–2777.

Andres, M., Seron, X., & Olivier, E. (2007). Contribution of hand motor circuits to counting. *Journal of Cognitive Neuroscience*, *19*, 563–576.

Atkinson, J., Campbell, F. W., & Francis, M. R. (1976). The magic number 4±0: a new look at visual numerosity judgments. *Perception*, *5*, 327–334.

Balakrishnan, J. D., & Ashby, F. G. (1991). Is subitizing a unique numerical ability? *Perception and Psychophysics*, *50*, 555–564.

Balakrishnan, J. D., & Ashby, F. G. (1992). Subitizing: Magical numbers or mere superstition? *Psychological Research*, *54*, 80–90.

Beran, M. J., Taglialatela, L. A., Flemming, T. M., James, F. M., & Washburn, D. A. (2006). Nonverbal estimation during numerosity judgements by adult humans. *Quarterly Journal of Experimental Psychology*, *59*, 2065–2082.

Bergmann Tiest, W. M., & Kappers, A. M. L. (2007). Haptic and visual perception of roughness. *Acta Psychologica*, *124*, 177–189.

Camos, V., & Tillmann, B. (2008). Discontinuity in the enumeration of sequentially presented auditory and visual stimuli. *Cognition*, *107*, 1135–1143.

Carlson, R. A., Avraamides, M. N., Cary, M., & Strasberg, S. (2007). What do the hands externalize in simple arithmetic? *Journal of Experimental Psychology: Learning Memory and Cognition*, *33*, 747–756.

Cashdan, S. (1968). Visual and haptic form discrimination under conditions of successive stimulation. *Journal of Experimental Psychology*, *76*, 215–218.

Cave, K. R., & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, *22*, 225–271.

Cordes, S., Gelman, R., Gallistel, C. R., & Whalen, J. (2001). Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychonomic Bulletin and Review*, *8*, 698–707.

Coren, S. (1993). *The left-hander syndrome: The causes and consequences of left-handedness*. New York: Vintage Books.

Craig, J. C., & Lyle, K. B. (2001). A comparison of tactile spatial sensitivity on the palm and fingerpad. *Perception & Psychophysics*, *63*, 337–347.

Dehaene, S., Dehaene-Lambertz, G., & Cohen, L. (1998). Abstract representations of numbers in the animal and human brain. *Trends in Neurosciences*, *21*, 355–361.

DeLucia, P. R., Mather, R. D., Griswold, J. A., & Mitra, S. (2006). Toward the improvement of image-guided interventions for minimally invasive surgery: Three factors that affect performance. *Human Factors*, *48*, 23–38.

Dodson, M. J., Goodwin, A. W., Browning, A. S., & Gehring, H. M. (1998). Peripheral neural mechanisms determining the orientation of cylinders grasped by the digits. *Journal of Neuroscience*, *18*, 521–530.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological review*, *96*, 433–458.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429–433.

Ernst, M. O., Lange, C., & Newell, F. N. (2007). Multisensory recognition of actively explored objects. *Canadian Journal of Experimental Psychology*, *61*, 242–253.

Gallace, A., Tan, H. Z., & Spence, C. (2006). Numerosity judgments for tactile stimuli distributed over the body surface. *Perception*, *35*, 247–266.

Gallace, A., Tan, H. Z., & Spence, C. (2008). Can tactile stimuli be subitised? an unresolved controversy within the literature on numerosity judgments. *Perception*, *37*, 782–800.

Gallistel, C. R., & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition*, *44*, 43–74.

Gallistel, C. R., & Gelman, R. (2000). Non-verbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences*, *4*, 59–65.

Goodwin, A. W., John, K. T., & Marceglia, A. H. (1991). Tactile discrimination of curvature by humans using only cutaneous information from the fingerpads. *Experimental Brain Research*, *86*, 663–672.

Goodwin, A. W., & Wheat, H. E. (2004). Sensory signals in neural populations underlying tactile perception and manipulation. *Annual Review of Neuroscience*, *27*, 53–77.

Gordon, I. E., & Morison, V. (1982). The haptic perception of curvature. *Perception and Psychophysics*, *31*, 446–450.

Grierson, L. E. M., & Carnahan, H. (2006). Manual exploration and the perception of slipperiness. *Perception & Psychophysics*, *68*, 1070–1081.

Hollins, M., & Risner, S. R. (2000). Evidence for the duplex theory of tactile texture perception. *Perception and Psychophysics*, *62*, 695–705.

Hubbard, E. M., Piazza, M., Pinel, P., & Dehaene, S. (2005). Interactions between number and space in parietal cortex. *Nature Reviews Neuroscience*, *6*, 435–448.

Immersion Corporation: CyberGrasp. (n.d.).

Izard, V., & Dehaene, S. (2008). Calibrating the mental number line. *Cognition*, *106*, 1221–1247.

Jenmalm, P., Birznieks, I., Goodwin, A. W., & Johansson, R. S. (2003). Influence of object shape on responses of human tactile afferents under conditions characteristic of manipulation. *European Journal of Neuroscience*, *18*, 164–176.

Jensen, E. M., Reese, E. P., & Reese, T. W. (1950). The subitizing and counting of visually presented fields of dots. *Journal of Psychology*, *30*, 363–392.

Johansson, R. S., & Vallbo, A. B. (1979). Tactile sensibility in the human hand: Relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *Journal of Physiology*, *Vol. 286*, 283–300.

Johnson, K. O., & Hsiao, S. S. (1992). Neural mechanisms of tactual form and texture perception. *Annual Review of Neuroscience*, *15*, 227–250.

Jones, L. A. (1986). Perception of force and weight. Theory and research. *Psychological Bulletin*, *100*, 29–42.

Jones, L. A., & Ho, H. N. (2008). Warm or cool, large or small? The challenge of thermal displays. *IEEE Transactions on Haptics*, *1*, 53–70.

Kahrimanovic, M., Bergmann Tiest, W. M., & Kappers, A. M. L. (in press). Haptic perception of volume and surface area of 3-D objects. *Attention, Perception & Psychophysics.*

Kappers, A. M., Koenderink, J. J., & Lichtenegger, I. (1994). Haptic identification of curved surfaces. *Perception & Psychophysics*, *56*, 53–61.

Kaufman, E., Lord, M., Reese, T., & Volkmann, J. (1949). The discrimination of visual number. *American Journal of Psychology*, *62*, 498–525.

Klatzky, R. L., & Lederman, S. J. (1995). Identifying objects from a haptic glance. *Perception & Psychophysics*, *57*, 1111–1123.

Klatzky, R. L., & Lederman, S. J. (1999). Tactile roughness perception with a rigid link interposed between skin and surface. *Perception and Psychophysics*, *61*, 591–607.

Klatzky, R. L., & Lederman, S. J. (2003). Representing spatial location and layout from sparse kinesthetic contacts. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 310–325.

Klatzky, R. L., Lederman, S. J., & Metzger, V. A. (1985). Identifying objects by touch: an "expert system". *Perception & Psychophysics*, *37*, 299–302.

Lakatos, S., & Marks, L. E. (1999). Haptic form perception: Relative salience of local and global features. *Perception and Psychophysics*, *61*, 895–908.

Lederman, S. J. (1981). The perception of surface roughness by active and passive touch. *Bulletin of the Psychonomic Society*, *18*, 253–255.

Lederman, S. J., Browse, R. A., & Klatzky, R. L. (1988). Haptic processing of spatially distributed information. *Perception & Psychophysics*, *44*, 222–232.

Lederman, S. J., & Klatzky, R. L. (1987). Hand movements: a window into haptic object recognition. *Cognitive Psychology*, *19*, 342–368.

Lederman, S. J., & Klatzky, R. L. (1990). Haptic classification of common objects: Knowledge-driven exploration. *Cognitive Psychology*, *22*, 421–459.

Lederman, S. J., & Klatzky, R. L. (1997). Relative availability of surface and object properties during early haptic processing. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 1680–1707.

Lederman, S. J., & Klatzky, R. L. (2004). Haptic identification of common objects: Effects of constraining the manual exploration process. *Perception & Psychophysics*, *66*, 618–628.

Lederman, S. J., & Taylor, M. M. (1972). Fingertip force, surface geometry, and the perception of roughness by active touch. *Perception and Psychophysics*, *12*(5), 401–408.

Logan, G. D., & Zbrodoff, N. J. (2003). Subitizing and similarity: Toward a pattern-matching theory of enumeration. *Psychonomic Bulletin and Review*, *10*, 676–682.

Loomis, J. M., Klatzky, R. L., & Lederman, S. J. (1991). Similarity of tactual and visual picture recognition with limited field of view. *Perception*, *20*, 167–177.

Luwel, K., Beem, A. L., Onghena, P., & Verschaffel, L. (2001). Using segmented linear regression models with unknown change points to analyze strategy shifts in cognitive tasks. *Behavior Research Methods, Instruments, and Computers*, *33*, 470–478.

Magee, L. E., & Kennedy, J. M. (1980). Exploring pictures tactually. *Nature*, *283*, 287–288.

Mandler, G., & Shebo, B. J. (1982). Subitizing: an analysis of its component processes. *Journal of Experimental Psychology: General*, *111*, 1–22.

Millar, S., & Al-Attar, Z. (2004). External and body-centered frames of reference in spatial memory: Evidence from touch. *Perception and Psychophysics*, *66*, 51–59.

Minamizawa, K., Kamura, S., Kawakami, N., & Tachi, S. (2008). A palm-worn haptic display for bimanual operations in virtual environments. In M. Ferre (Ed.), *Haptics: Perception, devices and scenarios* (Vol. 5024 of Lecture Notes on Computer Science, pp. 458–463). Berlin/Heidelberg: Springer.

Moretto, G., & Di Pellegrino, G. (2008). Grasping numbers. *Experimental Brain Research*, *188*, 505–515.

Moyer, R. S., & Landauer, T. K. (1967). Time required for judgements of numerical inequality [47]. *Nature*, *215*, 1519–1520.

Nabeta, T., Ono, F., & Kawahara, J. I. (2003). Transfer of spatial context from visual to haptic search. *Perception*, *32*, 1351–1358.

Newell, F. N., Ernst, M. O., Tjan, B. S., & Bülthoff, H. H. (2001). Viewpoint dependence in visual and haptic object recognition. *Psychological Science*, *12*, 37–42.

Norman, J. F., Norman, H. F., Clayton, A. M., Lianekhammy, J., & Zielke, G. (2004). The visual and haptic perception of natural object shape. *Perception and Psychophysics*, *66*, 342–351.

Nothdurft, H. C. (2000). Salience from feature contrast: additivity across dimensions. *Vision Research*, *40*, 1183–1201.

Overvliet, K. E., Mayer, K. M., Smeets, J. B. J., & Brenner, E. (2008). Haptic search is more efficient when the stimulus can be interpreted as consisting of fewer items. *Acta Psychologica*, *127*, 51–56.

Overvliet, K. E., Smeets, J. B. J., & Brenner, E. (2007a). Haptic search with finger movements: Using more fingers does not necessarily reduce search times. *Experimental Brain Research*, *182*, 427–434.

Overvliet, K. E., Smeets, J. B. J., & Brenner, E. (2007b). Parallel and serial search in haptics. *Perception & Psychophysics*, *69*, 1059–1069.

Overvliet, K. E., Smeets, J. B. J., & Brenner, E. (2008). The use of proprioception and tactile information in haptic search. *Acta Psychologica*, *129*, 83–90.

Oyama, T., Kikuchi, T., & Ichihara, S. (1981). Span of attention, backward masking, and reaction time. *Perception & Psychophysics*, *29*, 106–112.

Pasini, M., & Tessari, A. (2001). Hemispheric specialization in quantification processes. *Psychological Research*, *65*, 57–63.

Pesenti, M., Thioux, M., Seron, X., & De Volder, A. (2000). Neuroanatomical substrates of arabic number processing, numerical comparison, and simple addition: A pet study. *Journal of Cognitive Neuroscience*, *12*, 461–479.

Phillips, J. R., & Johnson, K. O. (1981). Tactile spatial resolution. II. neural representation of bars, edges, and gratings in monkey primary afferents. *Journal of Neurophysiology*, *46*, 1192–1203.

Piazza, M., Mechelli, A., Butterworth, B., & Price, C. J. (2002). Are subitizing and counting implemented as separate or functionally overlapping processes? *NeuroImage*, *15*, 435–446.

Piazza, M., Mechelli, A., Price, C. J., & Butterworth, B. (2006). Exact and approximate judgements of visual and auditory numerosity: An fMRI study. *Brain Research*, *1106*, 177–188.

Pinel, P., Piazza, M., Le Bihan, D., & Dehaene, S. (2004). Distributed and overlapping cerebral representations of number, size, and luminance during comparative judgments. *Neuron*, *41*, 983–993.

Plaisier, M. A., Bergmann Tiest, W. M., & Kappers, A. M. L. (2008a). Haptic pop-out in a hand sweep. *Acta Psychologica*, *128*, 368–377.

Plaisier, M. A., Bergmann Tiest, W. M., & Kappers, A. M. L. (2008b). Haptic search for spheres and cubes. In M. Ferre (Ed.), *Haptics: Perception, devices and scenarios* (Vol. 5024 of Lecture Notes on Computer Science, pp. 275–282). Berlin/Heidelberg: Springer.

Plaisier, M. A., Bergmann Tiest, W. M., & Kappers, A. M. L. (2009a). One, two, three, many - subitizing in active touch. *Acta Psychologica*, *131*, 163–170.

Plaisier, M. A., Bergmann Tiest, W. M., & Kappers, A. M. L. (2009b). Salient features in three-dimensional haptic shape perception. *Attention, Perception & Psychophysics*, *71*(2), 421–430.

Plaisier, M. A., Kuling, I. A., Bergmann Tiest, W. M., & Kappers, A. M. L. (2009). The role of item fixation in haptic search. In *Third joint eurohaptics conference and symposium on haptic interfaces for virtual* *environment and teleoperator systems* (pp. 417–421).

Pont, S. C., Kappers, A. M. L., & Koenderink, J. J. (1997). Haptic curvature discrimination at several regions of the hand. *Perception and Psychophysics*, *59*, 1225–1240.

Purdy, K. A., Lederman, S. J., & Klatzky, R. L. (2004). Haptic processing of the location of a known property: Does knowing what you've touched tell you where it is? *Canadian Journal of Experimental Psychology*, *58*, 32–45.

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, *80*, 127–158.

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: evidence for a parallel tracking mechanism. *Spatial Vision*, *3*, 179–197.

Revkin, S. K., Piazza, M., Izard, V., Cohen, L., & Dehaene, S. (2008). Does subitizing reflect numerical estimation? *Psychological Science*, *19*, 607–614.

Riggs, K. J., Ferrand, L., Lancelin, D., Fryziel, L., Dumur, G., & Simpson, A. (2006). Subitizing in tactile perception. *Psychological Science*, *17*, 271–272.

Rosenholtz, R. (2001). Search asymmetries? what search asymmetries? *Perception and Psychophysics*, *63*, 476–489.

Rosenholtz, R., Nagy, A. L., & Bell, N. R. (2004). The effect of background color on asymmetries in color search. *Journal of Vision*, *4*, 224–240.

Ross, J. (2003). Visual discrimination of number without counting. *Perception*, *32*, 867–870.

Ten Hoopen, G., & Vos, J. (1979). Effect or numerosity judgement of grouping of tones by auditory channels. *Perception & Psychophysics*, *26*, 374–380.

Theeuwes, J. (1993). Visual selective attention: A theoretical analysis. *Acta Psychologica*, *83*, 93–154.

Treisman, A. (1985). Preattentive processing in vision. *Computer Vision, Graphics and Image Processing*, *31*, 156–177.

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

Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. *Psychological Review*, *95*, 15–48.

Treisman, A., & Souther, J. (1985). Search asymmetry: a diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, *114*, 285–310.

Trick, L. M. (2008). More than superstition: Differential effects of featural heterogeneity and change on subitizing and counting. *Perception & Psychophysics*, *70*, 743–760.

Trick, L. M., & Pylyshyn, Z. W. (1993). What enumeration studies can show us about spatial attention: evidence for limited capacity preattentive processing. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 331–351.

Trick, L. M., & Pylyshyn, Z. W. (1994). Why are small and large numbers enumerated differently? A limited-capacity preattentive stage in vision. *Psychological Review*, *101*, 80–102.

Van der Horst, B. J., & Kappers, A. M. L. (2007). Curvature discrimination in various finger conditions. *Experimental Brain Research*, *177*, 304–311.

Van Oeffelen, M. P., & Vos, P. G. (1982a). Configurational effects on the enumeration of dots: counting by groups. *Memory and Cognition*, *10*, 396–404.

Van Oeffelen, M. P., & Vos, P. G. (1982b). A probabilistic model for the discrimination of visual number. *Perception and Psychophysics*, *32*, 163–170.

Van Oeffelen, M. P., & Vos, P. G. (1984). Enumeration of dots: an eye movement analysis. *Memory and Cognition*, *12*, 607–612.

Walsh, V. (2003). A theory of magnitude: Common cortical metrics of time, space and quantity. *Trends in Cognitive Sciences*, *7*, 483–488.

Watson, D. G., Maylor, E. A., & Bruce, L. A. M. (2007). The role of eye movements in subitizing and counting. *Journal of Experimental Psychology: Human Perception and Performance*, *33*, 1389–1399.

Whalen, J., Gallistel, C. R., & Gelman, R. (1999). Nonverbal counting in humans: The psychophysics of number representation. *Psychological Science*, *10*, 130–137.

Wijntjes, M. W. A., Van Lienen, T., Verstijnen, I. M., & Kappers, A. M. L. (2008). Look what I have felt: Unidentified haptic line drawings are identified after sketching. *Acta Psychologica*, *128*, 255–263.

Wolfe, J. M. (1993). Guided search 2.0: the upgrade. *Proceedings of the Human Factors and Ergonomics Society*, *2*, 1295–1299.

Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science*, *9*, 33–39.

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 419–433.

Worchel, P. (1951). Space perception and orientation in the blind. *Psychological Monographs*, *65*(15, (Whole No. 332)), 1–28.

Wu, B., Klatzky, R. L., Shelton, D., & Stetten, G. D. (2005). Psychophysical evaluation of in-situ ultrasound visualization. *IEEE Transactions on Visualization and Computer Graphics*, *11*, 684–693.

Yamamoto, N., & Shelton, A. L. (2007). Path information effects in visual and proprioceptive spatial learning. *Acta Psychologica*, *125*, 346–360.

Zelinsky, G. J., & Sheinberg, D. L. (1997). Eye movements during parallel-serial visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 244–262.

Samenvatting

In dit proefschrift werd een serie studies naar de haptische perceptie van meerdere objecten die tegelijkertijd in de hand worden gehouden uiteengezet. Hiertoe werden zowel het visuele zoekparadigma als het aantal-bepalings-paradigma vertaald naar het haptische domein. Bovendien werden twee nieuwe soorten haptische stimuli ontwikkeld. De eerste soort bestond uit een oppervlak waarop items konden worden geplaatst. De tweede soort stimulus bestond uit een verzameling objecten die tegelijkertijd in de hand gehouden konden worden. In beide gevallen bleek analyse van de exploratieve strategie van belang bij het interpreteren van de responsietijden. We hebben een manier geïntroduceerd om deze bewegingen te analyseren. Bovendien hebben we de responsietijden gemodelleerd om conclusies over het de onderliggende processen te kunnen trekken. Hieronder volgt een samenvatting van de resultaten uit de verschillende hoofdstukken en de algemene conclusies uit dit proefschrift.

Haptisch zoeken

In visuele zoektaken wordt over het algemeen een set van items gepresenteerd op een scherm. De taak van de proefpersoon is om te zeggen of een bepaald item (het doel) aanwezig is tussen de overige items (afleiders). Het aantal items wordt gevarieerd en de responsietijden worden gemeten als functie van het aantal items. De helling van deze functie (zoekhelling) is een maat voor hoe efficiënt de zoektaak werd uitgevoerd. Als deze helling klein is wordt er gezegd dat de taak *parallel* over alle items werd uitgevoerd. Dit wordt ook wel pop-out-effect genoemd (Treisman & Gelade, 1980; Treisman & Souther, 1985). Een grote helling geeft aan dat de taak *serieel* werd uitgevoerd. Door bepaalde doel- en afleidercombinaties te kiezen kunnen conclusies over de opvallendheid van bepaalde objecteigenschappen getrokken worden. In Hoofdstuk 2 tot en met 5 werden verschillende haptische varianten van een zoektaak gepresenteerd.

In Hoofdstuk 2 kregen proefpersonen de opdracht om een oppervlak met daarop stukjes schuurpapier te onderzoeken en te rapporteren of het doel, dat bestond uit schuurpapier met een andere ruwheid dan de afleiders, aanwezig was. In dit geval waren de zoekhellingen klein (20 ms/item) als het doel ruwer was dan de afleiders. Wanneer het doel minder ruw was dan de afleiders werden veel grotere hellingen gemeten (260 ms/item). Dit suggereert dat er een zoekasymmetrie is in dit geval. Bovendien lieten we zien dat deze grote verschillen tussen de zoekhellingen gepaard gaan met verschillen in de zoekstrategie die proefpersonen gebruikten. Als het doel ruwer was dan de afleiders kon het doel meestal gevonden worden door een enkele veeg met de hand over het oppervlak. In het omgekeerde geval (doel minder ruw dan afleiders) waren gedetailleerdere handbewegingen noodzakelijk. In deze studie lieten we zien dat analyse van handbewegingen inzicht geeft in de gebruikte zoekstrategie en van belang is bij het interpreteren van zoekhellingen. Door zoekhellingen te relateren aan handbewegingen lieten we voor de eerste keer zien dat pop-out-effecten kunnen optreden in haptische zoektaken met vrije exploratie.

Een andere interessante vondst uit Hoofdstuk 2 was dat proefpersonen sommige delen van het oppervlak een paar keer onderzochten. Dit duidt erop dat proefpersonen onzeker waren over welke delen van het oppervlak ze al onderzocht hadden. In dat geval zou de zoektaak vergemakkelijkt kunnen worden door ruimtelijke visuele informatie over het oppervlak aan te bieden. Dit werd op een gecontroleerde manier onderzocht in Hoofdstuk 3. Er werd in die studie gebruik gemaakt van een kracht-terugkoppelingsapparaat om de haptische stimulus (een oppervlak met gebieden met een andere wrijvingscoëfficiënt dan de achtergrond) te genereren. Deze stimulus kon enkel met de wijsvinger gevoeld worden. Er kon een visuele stimulus op dezelfde positie als de haptische gepresenteerd worden. De proefpersonen kregen de opdracht om een doel met hogere wrijvingscoëfficiënt dan de afleiders te vinden. Het doel was altijd alleen haptisch herkenbaar, maar er kon visuele informatie gegeven worden over de posities van de items en/of de positie van de vinger van de proefpersoon. Analyse van de bewegingen van de vinger over het oppervlak liet zien dat proefpersonen visuele informatie over itemposities konden gebruiken ook als ze geen visuele informatie over hun vingerpositie kregen. Verder bleek dat als er enkel informatie over vingerpositie werd gegeven, de proefpersonen eerst systematisch het hele oppervlak onderzochten maar daarna gerichte bewegingen naar eerder gevoelde items maakten. Dit was niet het geval als er geen visuele informatie over vingerpositie werd gegeven. Dit betekent dat visuele informatie over vingerpositie helpt bij de vorming van een ruimtelijke representatie van de items op het oppervlak.

Verder laten we in Hoofdstuk 3 zien dat modellen die responsietijden voor visuele zoektaken beschrijven niet zonder meer van toepassing zijn op haptische zoektaken. In visuele modellen wordt meestal aangenomen dat een zoektaak parallel werd uitgevoerd als de zoekhelling klein is. De stimulus die in Hoofdstuk 3 werd gebruikt kon alleen serieel onderzocht worden, maar in de condities zonder visuele informatie over itemposities werden zoekhellingen gevonden die niet significant verschilden van nul. De oorzaak hiervoor is dat proefpersonen het gehele oppervlak moesten onderzoeken om de items te vinden. Hierdoor waren responsietijden vrijwel onafhankelijk van het aantal items. Dit geeft nog eens aan dat het van belang is om haptische zoekhellingen te interpreteren in combinatie met de exploratieve strategie die werd gebruikt.

In Hoofdstukken 4 en 5 bestonden de aangeboden items uit drie-dimensionale vormen waarvan er meerdere tegelijk in de hand gehouden werden. Dit type stimulus werd gebruikt om te onderzoeken welke drie-dimensionale vormeigenschappen het meest opvallend zijn. In Hoofdstuk 4 kon het doel een bol, kubus, tetraëder, ellipsoïde of cilinder zijn. De afleiders waren bollen of kubussen en proefpersonen kregen weer de taak om te bepalen of het doel aanwezig was. Uit de resultaten bleek dat de zoekhelling het kleinst was als het doel randen had en de afleiders niet of omgekeerd. In dat geval kon het doel gevonden worden door de hele aangeboden set vormen in de hand te houden. Waneer zowel doel als afleiders randen hadden of beide niet, waren zoekhellingen groot. Bovendien gebruikten proefpersonen dan een seriële zoekstrategie waarbij ze de vormpjes een voor een uit de hand lieten vallen. Dit geeft aan dat een rand een zeer saillante vormeigenschap is.

Omdat de aangeboden vormen in Hoofdstuk 4 opgehangen werden aan flexibele draden konden ze vrij in de hand worden geroteerd en hun positie kon herschikt worden. In Hoofdstuk 5 werd onderzocht hoe haptische vormperceptie wordt beïnvloed door deze vrijheid om de vormen te manipuleren. Verder werd ook onderzocht of er verschillen zijn in responsietijd voor verschillende posities van het doel in de hand. Hiertoe werden kubussen en bollen zodanig bevestigd dat alleen rotaties en kleine translaties nog mogelijk waren of dusdanig dat ze volledig gefixeerd waren in de ruimte. Uit de resultaten bleek dat er geen verschillen in responsietijden waren voor de verschillende posities van het doel in de hand. Verder bleken responsietijden korter te zijn als de vormpjes gefixeerd waren met een starre verbinding dan waneer rotaties en kleine translaties mogelijk waren. Een verklaring hiervoor is dat het gemakkelijker was om kracht uit te oefenen op de vormen waneer ze volledig waren gefixeerd.

Haptisch aantal bepalen

De tweede soort taak die werd geëxtrapoleerd van het visuele naar het haptische domein was een aantal-bepalings-taak (Hoofdstuk 6 tot en met 9). Net als in de zoektaken werd een variërend aantal items aangeboden, maar de taak voor de proefpersoon was in dit geval om te zeggen hoeveel items er waren. De stimulus bestond altijd uit een aantal drie-dimensionale vormen die gelijktijdig in de hand gehouden konden worden. Uit visuele studies is bekend dat aantal-bepaling tot ongeveer 4 items snel en foutloos is, terwijl voor hogere aantallen de responsietijden en foutpercentages snel toenemen (e.g. Atkinson et al., 1976; Mandler & Shebo, 1982; Trick & Pylyshyn, 1993, 1994). Doorgaans wordt aangenomen dat grote aantallen worden bepaald door te *tellen*, maar dat er een efficiënter proces wordt gebruikt voor kleine aantallen. Dit proces staat bekend onder de term *subitizing* (Kaufman et al., 1949).

In Hoofdstuk 6 werd onderzocht of subitizing ook optreedt voor actieve tast. Proefpersonen kregen de opdracht om zo snel mogelijk te bepalen hoeveel bollen ze in hun hand hadden. Uit de resultaten bleek dat dit foutloos en efficiënter was voor minder dan 4 bollen (167 ms/item) dan voor grotere aantallen (839 ms/item). Tevens werd in dit hoofdstuk aangetoond dat subitizing niet hetzelfde is als nauwkeurig schatten van aantal, volume of gewicht. Er kan geconcludeerd worden dat subitizing ook optreedt in actieve tast.

Om haptische en visuele subitizing te vergelijken werden de haptische experimenten uit Hoofdstuk 6 gereproduceerd in het visuele domein in Hoofdstuk 7. Uit deze experimenten bleek dat ook in het visuele domein geldt dat subitizing niet hetzelfde is als nauwkeurig schatten. Bovendien kon het model dat in Hoofdstuk 6 werd gebruikt om responsietijden voor schatten van aantal te beschrijven ook gebruikt worden om de visuele responsietijden te beschrijven. De overeenkomsten tussen de resultaten van Hoofdstuk 6 en 7 suggereren een vergelijkbaar dan wel hetzelfde onderliggende mechanisme voor aantal-bepaling in de haptische en visuele domeinen.

In visuele studies is gesuggereerd dat subitizing in combinatie met tellen kan optreden. Waneer een veld met stippen wordt opgedeeld in clusters die aantallen uit het subitizing-regime bevatten, dan worden de responsietijden voor het bepalen van het totale aantal stippen kleiner (Van Oeffelen & Vos, 1982a, 1984). In Hoofdstuk 8 werd onderzocht of dit ook haptisch het geval is. In dit hoofdstuk moesten proefpersonen wederom het aantal bollen in hun hand bepalen, maar deze werden onderverdeeld in twee groepen door ze te verdelen over beide handen. De responsietijden voor een bepaald aantal bollen in één hand werden vergeleken met de tijden voor het bepalen van hetzelfde aantal verdeeld over twee handen. Uit de resultaten bleek dat de responsietijden kleiner waren als de bollen over twee handen werden verdeeld en dat er een combinatie van subitizing en tellen gebruikt werd. Dit toont opnieuw aan dat aantallen op vergelijkbare wijzen worden verwerkt in de haptische en de visuele modaliteit. Dit pleit voor een modaliteitsonafhankelijke representatie van aantal-informatie.

In de voorgaande hoofdstukken moesten proefpersonen altijd het aantal identieke objecten bepalen. In het dagelijks leven hebben we interactie met objecten die variëren in allerlei aspecten zoals afmetingen en vorm. Voordat we het aantal objecten in onze hand kunnen bepalen, moeten deze objecten eerst geïndividualiseerd worden, i.e. er moet besloten worden welke delen bij het ene object horen en welke bij een ander. In Hoofdstuk 9 werd onderzocht of de afmetingen en vorm van objecten een rol spelen in dit proces. Eerst werd de rol van de afmetingen van objecten onderzocht door kleinere bollen te gebruiken dan in de voorgaande hoofdstukken. In dit geval werd het subitizing-regime niet uitgebreid naar hogere aantallen dan eerder gevonden werd. In twee andere experimenten werd de invloed van heterogeniteit van de items onderzocht. In dat geval moesten proefpersonen het aantal items bepalen voor een combinatie van bollen met verschillende afmetingen of een combinatie van kubussen en bollen. Er werden geen verschillen gevonden tussen de experimenten in termen van responsietijden, foutpercentages of exploratiestrategie. Dit suggereert dat eigenschappen zoals afmetingen en vorm die van belang zijn bij objectherkenning geen rol spelen bij de individualisatie van objecten.

Algemene conclusies

De resultaten van de gepresenteerde zoektaken laten zien dat er objecteigenschappen zijn zoals materiaal (ruwheid) en vorm (randen) die zeer efficiënt en met een parallelle exploratiestrategie geëxtraheerd kunnen worden. Dit betekent dat haptisch zoeken parallel kan verlopen. Uit de gepresenteerde studies naar haptische aantal-perceptie blijkt dat bepaling van kleine aantallen (< 4) efficiënter en nauwkeuriger is dan van grote aantallen. Als we aannemen dat tellen van grote aantallen serieel verloopt, dan moet het subitizen van kleine aantallen gedeeltelijk parallel verlopen. Voordat het aantal objecten kan worden bepaald moeten deze van elkaar onderscheiden worden. Het feit dat subitizing optreedt suggereert daarom dat het proces van objecten onderscheiden tenminste voor kleine aantallen parallel kan verlopen. Zowel het onderscheiden van objecten als het herkennen van objecten kan dus zeer efficiënt verlopen. Objecteigenschappen die belangrijk zijn voor objectherkenning lijken echter geen rol te spelen bij de individualisatie van objecten. Daarom kan er geconcludeerd worden dat haptisch herkennen en individualiseren van objecten snelle maar onafhankelijke processen zijn.

Publications and award

Journal articles

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. (2008). Haptic pop-out in a hand sweep. *Acta Psychologica, 128*, 368–282.

Kahrimanović, M., Bergmann Tiest, W.M., Plaisier, M.A., Sanders, A.F.J., van der Horst, B.J. & Kappers, A.M.L. (2008). Haptische waarneming van materiaal, vorm en ruimte. *Nederlands Tijdschrift voor Ergonomie, 31(6)*, 14–19.

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. (2009). Salient features in three-dimensional haptic shape perception. *Attention, Perception & Psychophysics, 71(2)*, 421–430.

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. (2009). One, two, three, many – subitizing in active touch. *Acta Psychologica, 131,* 163–170.

Plaisier, M.A., Kappers, A.M.L., Bergmann Tiest, W.M., & Marc O. Ernst. Visually guided haptic search. *Transactions on Haptics*, in press.

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Grabbing subitizing with both hands – Bimanual number processing, *Experimental Brain Research*, 2010, DOI 10.1007/s00221-009-2146-1.

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Haptic object individuation, *IEEE Transactions on Haptics*, in press.

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Similar processing in visual and haptic numerosity judgment (submitted).

Conference proceedings

Vossen, D.L.J., Plaisier M.A., Van Blaaderen, A. (2004). Colloidal crystallization induced by optical gradient forces exerted by optical tweezers. *Proceedings of SPIE – The International Society for Optical Engineering*, 5514, art. no. 102, pp. 755–762.

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. (2008). Haptic search for spheres and cubes. In M. Ferre (Ed.), *Haptics: Perception, devices and scenarios* (Vol. 5024 of Lecture Notes on Computer Science, pp. 275–282). Berlin/Heidelberg: Springer.

Plaisier, M.A., Kuling, I.A., Bergmann Tiest, W.M., & Kappers, A.M.L. (2009). The role of item fixation in haptic search. In *Third joint eurohaptics conference and symposium on haptic interfaces for virtual environments and teleoperator systems* (pp. 417–421).

Refereed abstracts

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Haptic pop-out in a hand sweep, Nederlandse vereniging voor psychonomie (NVP) Wintercongres (2007).

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Haptic and visual subitizing, *Perception, 37*, ECVP abstract Supplement (2008).

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. The effect of clustering on haptic numerosity judgment, *Perception, 38*, ECVP abstract Supplement (2009).

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Efficient bimanual processing of numerosity information, NWO Cognition autumn school: From stimulus to understanding in perception and language (2009).

Plaisier, M.A., Bergmann Tiest, W.M., & Kappers, A.M.L. Haptic numerosity judgment – Do all objects count equally?, Nederlandse vereniging voor psychonomie (NVP) Wintercongres (2009).

Award

Grant from IEEE's Technical Committee on Haptics: 'Student exchange program for cross-disciplinary fertilization' (2008). This grant enabled a six weeks stay at the Max Planck Institute for Biological Cybernetics in Tuebingen, Germany.

Acknowledgements

I would like to thank my promotor Astrid Kappers and co-promotor Wouter Bergmann Tiest for their support throughout my PhD research.

Experimental research would be much more difficult without technical assistance. Therefore, special thanks to Hans Kolijn and Pieter Schiphorst for various types of technical assistance, Jody Wisman for building the response time measuring device (Tast2) and Instrumentatie betawetenschappen for milling so many spheres, cubes, tetrahedrons, ellipsoids and cylinders.

I am grateful for Marc Ernst giving me the opportunity to visit his lab and use his equipment.

Finally, I would like to thank all members and former members of the Physics of Man group for useful discussions throughout my four years there.

Curriculum Vitae

Myrthe Amethyst Plaisier was born in Rotterdam, The Netherlands, on April 13, 1982. After receiving her VWO diploma (high school) from Udens College in 2000, she moved to Utrecht to study experimental physics at Universiteit Utrecht. In 2006 she graduated on a thesis about colloidal crystallisation.

After her graduation she started a four year PhD research in the Physics of Man group at Universiteit Utrecht under supervision of professor Astrid Kappers. During this time she also worked for a short period at the Max Planck Institute for Biological Cybernetics in Tübingen, Germany, in collaboration with dr. Marc Ernst.

Currently, she is finishing ongoing projects related to her PhD research but not included in this thesis.