

VU Research Portal

Multimodal Perception and Action

Nelson, J.S.

2020

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Nelson, J. S. (2020). *Multimodal Perception and Action*. [PhD-Thesis - Research and graduation internal, Vrije Universiteit Amsterdam].

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl


Multimodal perception and action

Jacob Nelson

The research presented in this thesis was supported by the Department of Human Movement Sciences, Faculty of Behavioural and Movement Sciences, Vrije Universiteit Amsterdam, Amsterdam Movement Sciences, the Netherlands, and the Institute of Brain and Behaviour Amsterdam, the Netherlands. This work was supported by the Horizon 2020 grant 642961: "PACE."



PERCEPTION
AND ACTION
IN COMPLEX
ENVIRONMENTS



Copyright © Jacob Nelson 2020. All rights reserved.

VRIJE UNIVERSITEIT

MULTIMODAL PERCEPTION AND ACTION

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor of Philosophy aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. V. Subramaniam,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de Faculteit der Gedrags- en Bewegingswetenschappen
op maandag 23 november 2020 om 15.45 uur
in de online bijeenkomst van de universiteit,
De Boelelaan 1105

Door

Jacob Solomon Nelson

geboren te Colorado, Verenigde Staten

Promotor: prof. dr. J.B.J. Smeets

Copromotor: dr. E. Brenner

Table of Contents

<u>GENERAL INTRODUCTION</u>	<u>6</u>
<u>SPATIAL REPRESENTATION OF THE WORKSPACE IN BLIND, LOW VISION AND SIGHTED HUMAN PARTICIPANTS</u>	<u>13</u>
<u>HOW WELL DO WE KNOW WHEN WE ACT?</u>	<u>30</u>
<u>HOW THE MODALITIES OF A SENSORY RHYTHMIC GUIDE AND OF FEEDBACK INFLUENCE RHYTHMIC TAPPING</u>	<u>39</u>
<u>THE ACCURACY OF INTERCEPTING MOVING TACTILE TARGETS</u>	<u>48</u>
<u>EPILOGUE</u>	<u>67</u>
<u>SUMMARY</u>	<u>74</u>
<u>REFERENCES</u>	<u>77</u>
<u>ACKNOWLEDGEMENTS</u>	<u>90</u>

1

General Introduction

The word ‘sensation’ refers to “a perception associated with stimulation of a sense organ or with a specific body condition” (American Heritage Dictionary, 2020). Colloquially, human perception is divided into five ‘senses’: sight, hearing, touch, smell, and taste; in reality, there are many more than five types of human sense organs, and each type plays a very specific role in human perception. For example, the task of determining the approximate position and orientation of our own body in the dark or with our eyes closed is relatively easy for most of us, but it requires input from a number of different kinds of receptors located in our skin, joints, muscles, and vestibular organs (Lackner & DiZio, 2005). Another example would be our sense of vision, which relies on input from both rod and cone cells in the eye (Gouras & Zrenner, 1981). It should come as no surprise that each sense, even in the broadest meaning of the word, is useful for accurately and precisely relaying certain types of information while being completely oblivious to others. Normally sighted individuals can easily spot the Moon on a clear night from a distance of hundreds of thousands of kilometers away, though no human sense organ besides the eye can detect such a distant object from here on Earth; meanwhile, loudspeakers producing music at even an uncomfortably high volume may not vibrate enough for us to see the movement. The normal human perception of the world therefore relies on the combination of many different sense organs all relaying signals to the brain, so that we can form what we perceive to be a complete picture of the present state of the world around us. This picture allows us to act in a way consistent with our current environment and our longer-term goals.

Our actions concern not only the locations of stimuli and our own body, but also the past and future states of these stimuli. A baseball batter needs to be able to predict when the incoming ball will pass within striking distance of her bat; the time of the ball’s passing is quintessential to the success of this action. Musicians and dancers must similarly be aware of the speed and rhythm at which they are performing, as well as the speed and rhythm of their fellow performers. Failure of the performers to synchronize with one another results in what most people would consider a sloppy or amateurish performance. Even something as simple as pouring coffee requires that the pourer keep track of the amount of coffee currently in the cup and the rate at which it is rising, the position of the stream relative to the cup, as well as the gradually decreasing weight of the carafe. Missing or misjudging even one of these sources of information can result in a spill (Vaina, 1989).

1.1. Space

One fundamental task of our various sensory systems is to help us determine the location of objects, creatures, and other people. Human perception of space can be described both in egocentric (“one meter to my right”) and in allocentric (“below the table”) terms (Burgess, 2006). We use both egocentric and allocentric reference frames to build our perception of the world around us, and we rely on a number of sources of sensory information—vision and proprioception, among others—to accomplish this task. Manipulating an object in our hand relies on egocentric reference frames, as we are only concerned with the spatial relation between the hand and object, regardless of where in the universe the hand and the object are. However, if we want to place the object down, walk to another room, and later return to retrieve it, allocentric reference frames provide far more useful information for recalling where it is (“the keys are in the bowl by the door”) than considering the angle and distance of the object from your hand. Spatial perception can also

differ greatly when using purely haptic information versus when using purely visual information; consider the last time you felt around in the dark for an object that you couldn't find, only to turn on the light and see the object right in front of you. Previous work has shown that people's spatial perception can fall between what one would expect from a strictly allocentric reference frame and that from a strictly egocentric one: when blindfolded individuals were asked to feel a bar resting on a table and subsequently rotate a second bar such that it was parallel to the first, the orientation they set the bar to differed substantially from parallelity (Kappers, 2004). Furthermore, the difference in angle between the two bars' orientations was greater when the two bars were set farther apart from each other.

Given that spatial perception relies in part on visual information, it should come as no surprise that blind (or blindfolded sighted) individuals perceive space differently from sighted individuals in some cases. People who became blind early in life, for example, are sometimes better than blindfolded sighted individuals at locating the source of a nearby sound (Voss et al., 2004), or at detecting differences in similar tactile stimuli (B. Jones, 1972). These same people might have difficulty pointing to an object they have placed on a table, though they can typically navigate back to the same spot and correctly locate the object on the table (Hollins & Kelley, 1988). The complicated relationship between spatial perception and visual experience demonstrates that spatial perception does not depend on one modality alone, but on a complex interaction between modalities.

1.2. Time

Just as important as knowing where things are around us is knowing when things are happening around us. Temporal perception allows us to estimate the speed of moving objects and predict future events, for example to estimate whether we have time to safely cross the street in front of a distant but approaching car. Temporal perception is also crucial for rhythmic activities such as singing, dancing, playing a musical instrument, or even speaking (for a review concerning both of these kinds of tasks, see Grondin, 2010). In order to execute any of these tasks, we need access to information about the tempo of the actions we are performing, possibly in relation to the tempo of an external stimulus, such as a metronome or a dance partner. As with spatial perception, temporal perception relies on input from a number of sensory modalities to accurately model the passage of time. Information gained from these various sensory modalities can provide vastly different indications of the timing of an event for both external and internal reasons. As light travels far faster than sound waves through air, distant events (such as a firework exploding overhead) can be seen long before they are heard, leading to a muddled perception of when exactly the witnessed event actually occurred. Internally, differing neural mechanisms for processing visual, auditory, and haptic information result in differences of precision for determining when an event occurred, even if said event is a self-executed action.

This thesis focuses on the perceived timing of brief, nearby events, such as a nearby beep or flash, or a self-generated finger tap. Even a discrepancy of a few tens of milliseconds between a visual and auditory event can be detected by the human brain (Zampini et al., 2003); it is for this reason that some people do not enjoy watching videos with wireless headphones, whose audio output tends to lag a phone or computer screen output by roughly a hundred milliseconds. One useful method of quantifying perceivable discrepancies between two similar stimuli is to calculate the so-called just noticeable difference (JND), or the

approximate magnitude of difference between two stimuli necessary for people to reliably tell them apart. In a temporal perception context, this usually refers to the number of milliseconds necessary to separate the onset of two separate events before someone can reliably determine which event occurred first. In addition to differences in temporal precision between sensory modalities, there is some evidence that self-generated motor actions may be more precisely timed than external events of the same duration. Understanding the relative precisions of each sensory modality, including when the stimulus is self-generated, is crucial to understanding how we perceive time and how accurately and precisely we are able to do so.

Another method commonly used to assess people's sense of rhythm and timing is the so-called finger tapping paradigm. In this task, participants are asked to tap their finger rhythmically along with an external guide stimulus, such as a metronome. In most cases, people are able to synchronize their taps to the rhythm of the guide reasonably well, though their exact accuracy and precision depends on a number of variables (for a review of the tapping literature, see Repp, 2005). For example, it is well documented that people tend to tap slightly too early on average. This so-called negative mean asynchrony is not yet well understood, despite being a nearly universal finding. People's performance also depends, unsurprisingly, on the so-called stimulus inter-onset interval, or how rapidly the hypothetical metronome ticks. Up to a point, the faster the metronome ticks, the smaller that negative mean asynchrony becomes, and the less variable their tap timing becomes. It should also come as no surprise that musicians tend to be better at this task than non-musicians, and drummers are particularly exceptional at rhythmic tapping tasks. A common method in experiments to assess learning or adaptation is to gradually alter the feedback a participant receives on their performance (see, for example, Kooij et al., 2015). Imagine that your computer introduced a slight but gradually increasing angular offset in your cursor motions, such that every time you moved the mouse forward, the cursor traveled upward as normal, as well as a pixel to the left. After a week, perhaps the difference between how the cursor "should" move and how it actually moves could be as much as 30 degrees and you might not notice. This is a useful tool in experiments because it demonstrates that even when one is not consciously aware of a change in their environment, they can alter their actions to compensate for that change. These subtle perturbations can be designed to occur not only in space but in time, for example to alter the timing of feedback given on a finger tap task in an attempt to shift participants' subsequent tap times.

1.3. Interception

Many tasks in everyday life require a sense of both space and time. A large number of popular sports primarily involve some version of catching or hitting a quickly moving ball. This kind of behavior is generally referred to as interception, and can be found across the animal kingdom, typically in terms of hunting. Interception in humans is usually thought of as a visual task, but plenty of animals rely on olfaction, echolocation, and electroreception to capture fleeing or unsuspecting prey (M. E. Nelson & MacIver, 1999; Surlykke & Kalko, 2008; Tricas, 1982). Humans typically rely on a (visually relayed) combination of position and velocity information about their target to estimate when they should move and where they need to be to hit the target, and they continue to adapt their own interceptive motion in response to incoming information, up until about one hundred milliseconds before they hit the target (Eli

Brenner & Smeets, 2015). Put another way, it takes roughly one hundred milliseconds for people to process incoming visual information and update their motor actions accordingly, meaning any changes to the target trajectory that take place right before the time of interception cannot be accounted for in time. Under the right circumstances, people are capable of intercepting a moving target with temporal precision as high as five to six milliseconds, a window far more precise than the temporal information relayed by any one sensory modality (Eli Brenner et al., 2012).

Humans can increase the temporal precision of their motor actions by making faster movements. This comes at the expense of spatial precision, but there are a number of techniques to compensate for this. In experimental interception tasks involving a circular target moving across a screen, people tend to follow the target with their hand before intercepting, matching the target's velocity, which likely optimizes performance with regard to spatial and temporal precision (de la Malla & López-Moliner, 2015). People also tend to follow the target with their eyes, as opposed to looking at their hand or at the space where they want (or need) to hit the target. This technique maximizes the availability of spatial and temporal information which can be used for interception. When instructed to look at the location where they must hit the target rather than at the moving target itself, people have a harder time hitting the target (de la Malla et al., 2017). Given that humans perform this specific oculomotor action to maximize their performance, one might assume that human interception is not just commonly a visual task, but in fact necessarily a visual task. Other sensory modalities may not convey the spatial and temporal information necessary, or we may simply be unable to meaningfully incorporate the information into our motor actions. In this thesis we explored whether this was the case by means of a tactile interception task.

1.4. Outline of the thesis

Much is still not known regarding the ways in which our perceptions, and the differences and limitations thereof, determine our actions. Of particular interest to us were questions regarding the precision and accuracy of vision, audition, and haptics, and how, in some cases, these senses work together when we perform basic motor actions. In our multisensory world, it can be difficult to determine exactly what sources of information influence a given action and to what degree; we therefore also sought to tease apart the influences of these different senses to understand how they individually inform us of our environment. This thesis comprises four diverse studies, which can broadly be said to explore the nature of spatial and temporal perception across the tactile, proprioceptive, visual, and auditory sensory modalities. The thesis will be arranged in such a way as to first discuss spatial perception, move on to temporal perception, and finish with a discussion of an interception task, which requires both spatial and temporal information.

Spatial perception is quintessential to our ability to navigate in, and interact with our environment. Whether reaching for a cup of tea, shaking someone's hand, or scratching an itch on our back, we need access to information both about the location of the target stimulus and about that of our own body. For most people, these kinds of abilities come naturally and are already easy to execute from a young age. However, as mentioned earlier, our sense of space relies on a combination of both egocentric and allocentric reference frames. For people who were born blind or who became blind later in life, an allocentric reference frame may not be as developed, or indeed even as useful in terms of orienting oneself relative to the

portion of the environment within physical reach of the body. Given that even normally sighted individuals make spatial miscalculations when deprived of their sense of sight, we wondered whether these miscalculations would be more pronounced in people who have never seen their environment, or whose sense of vision was lost at some point in their lives. **Chapter 2** explores the effect of experience of visual information on one's spatial perception and ability to execute simple actions within the space directly in front of their own body. In collaboration with researchers at Utrecht University and the Italian Institute of Technology, we tested whether blindfolded sighted adults would perform differently from people who were born blind and/or people who became blind later in life in a line drawing task. If such a difference in performance were to occur, it would mean that years of visual input lead to a gradual buildup of experience that meaningfully affects how we interact with the space around us, even if we currently cannot see. As we found no differences in performance between our early-blind, late-blind, and sighted participants, we concluded that the spatial motor skills required for our task actually develop independently from vision and do not benefit from past visual experience.

Our sense of time is also inherently dependent on our perception of the world around us. Generally speaking, events are multisensory in nature, meaning there is more than one source of information (e.g. vision and audition) that we can rely on to determine exactly when something occurred. Given the multisensory nature of most stimuli in our daily lives, it can be difficult to determine whether our perception is substantially dominated by one particular sense. One of the many differences between our various sensory modalities is the degree of temporal precision of the information available to the brain. Just as senses like vision and touch provide far more acute information about the location of a given stimulus than our senses of smell or taste could do, our ability to time a given stimulus also varies across sensory modalities. For sources within our reach, the time it takes for the stimuli (light and sound) to reach our body is nearly identical, but this does not necessarily mean that we can determine that exact moment with equal precision. Previous research has shown that auditory stimuli can be detected with greater temporal precision than visual stimuli (see, for example, Sugano et al., 2010), but the role of motor actions in timing is less well understood. As mentioned previously, people are able to achieve extremely high temporal precision when executing certain motor actions, but the nature of this enhanced temporal precision and how it relates to the precision of vision and audition is not yet fully understood. **Chapter 3** explores how well we perceive the timing of our own motor actions and of external events. In this study we demonstrate that the act of moving actually improves one's ability to determine the temporal order of two brief events.

Chapter 4 considers our ability to maintain a rhythm by tapping our finger. When tapping along with a guide rhythm (such as a metronome), people tend to tap a bit too early on average. They are, however, quite consistent in their performance. To date, most studies consider only one sensory modality (typically auditory) for the guide, and in cases where an explicit feedback stimulus is provided, it generally occurs in the same modality as the guide. It should be noted that as we also receive tactile and proprioceptive information about our actions and their influence on the world, haptic feedback is nearly always present in rhythmic tapping studies, or indeed any study wherein a participant performs a motor action. While rhythmic tapping paradigms are often thought of in terms of music and are therefore primarily auditory in nature, there are reasons one might want to consider a cross-sensory experimental paradigm. Dancers, for example, often practice their routines in front of a

mirror, in which case they receive visual (and haptic) feedback about the timing of their own movements in relation to the music they are attempting to follow. Members of a choir or orchestra rely on the visual signals of a conductor to remain on beat, with the added challenge that their fellow performers will also inevitably be off beat sometimes and may disrupt each other. From these examples we can easily conclude that it should be possible to perform cross-modality sensorimotor synchronization actions, but of course the precision with which we can do this, and our temporal sensitivity to these different sensory modalities, remains to be seen. To investigate this, we devised a study wherein participants tapped along with a guide rhythm (either a flash or a beep) and also received sensory feedback (again, either a flash or a beep) regarding the timing of each of their taps. We subtly altered the timing of the feedback they received such that it could arrive slightly too early or too late relative to the tap, in order to see how this small perturbation would affect our participants' ability to tap in synchrony with the guide. We found that when the modality of the guide and feedback differ from one another, people's ability to synchronize with the guide decreases, and their tap times can even be influenced by the relative offset of the most recent feedback stimuli.

The thesis ends with a discussion on interception. While interception in humans is typically thought of as a visuomotor task (that is, we see the moving target we want to intercept), little prior work has been done to explore whether the necessary information must come from the visual system in order to successfully intercept a target. Ecologically speaking, it may make the most sense to design a visual interception task, as the vast majority of human interception is performed using visual information. After all, our sense of vision is far better suited than our other senses for purposes of locating and detecting distant, fast-moving objects. Objects that are stationary or within our reach for longer periods of time typically do not warrant an interception motion when we could more easily grasp for them at our leisure. However, the fact that a task does not have obvious ecological validity does not mean it is not worth considering. If interception in humans is so often visual, one could imagine that our ability to intercept is also limited to targets in the visual domain, and that attempts to intercept a target based on another sensory modality would be utterly hopeless. However, if we are indeed capable of performing this task, it would mean that the neural mechanisms involved in calculating target trajectory are independent from the visual system. **Chapter 5** discusses a collaboration with the Italian Institute of Technology, wherein we blindfolded participants and examined their sense of the position and velocity of an object touching their left arm. We then had the participants use their right index finger to tap the object as it moved along the arm. We found that people's performance in the so-called 'position' and 'velocity' experiments could be used to predict performance in the 'interception' experiment, and that this performance could be modelled in a way similar to that in a visual interception task.

2

Spatial representation of the workspace in blind, low vision and sighted human participants

Adapted from:

Nelson, J. S., Kuling, I. A., Gori, M., Postma, A., Brenner, E., & Smeets, J. B. J. (2018). Spatial representation of the workspace in blind, low vision, and sighted human participants. *i-Perception*, 9(3), 1–15. doi: 10.1177/2041669518781877

Abstract

It has been proposed that haptic spatial perception depends on one's visual abilities. We tested spatial perception in the workspace using a combination of haptic matching and line drawing tasks. There were 132 participants with varying degrees of visual ability ranging from congenitally blind to normally sighted. Each participant was blindfolded and asked to match a haptic target position felt under a table with their non-dominant hand using a pen in their dominant hand. Once the pen was in position on the tabletop, they had to draw a line of equal length to a previously felt reference object by moving the pen laterally. We used targets at three different locations in order to evaluate whether different starting positions relative to the body give rise to different matching errors, drawn line lengths, or drawn line angles. We found no influence of visual ability on matching error, drawn line length or line angle, but we found that early-blind participants are slightly less consistent in their matching errors across space. We conclude that the elementary haptic abilities tested in these tasks do not depend on visual experience.

2.1 Introduction

The environment is filled with rich, multisensory stimuli. In humans, the visual system plays a crucial role in information processing, especially at distances that cannot be physically reached. For nearby stimuli, touch and proprioception help to form a complete picture of the world. Visual information plays an important role in the development of various spatial abilities, as it provides information about the position and arrangement of the surrounding environment that the nervous system would not otherwise necessarily have access to.

2.1.1. Effects of early-blindness

Several studies suggest that the brain may respond to visual deprivation by improving the sensitivity of the remaining sensory systems. Early-blind individuals show enhanced skills for some auditory tasks such as localization of a single sound on the horizontal plane (Lessard et al., 1998; Röder et al., 1999; Voss et al., 2004). Early-blind individuals have also been shown to have lower spatial thresholds for tactile discrimination of stimuli (Brown & Stratton, 1925; B. Jones, 1972), contrary to an earlier report (Seashore & Ling, 1918). Postma *et al.* (2007) found that early-blind and late-blind adults were significantly faster than blindfolded sighted adults at placing wooden shapes into a board with corresponding cutouts, as well as in placing the same shapes in the correct positions on a flat surface from memory. However, some studies did not find superior performance in haptics due to absence of vision. Jones and Vierck (1973) did not find a difference in threshold between blind and sighted individuals in detecting which of two stimuli pressed against their arm was longer. Gori *et al.* (2010) found that visually impaired children (one late-blind and the rest early-blind) were slightly better than sighted children at tactile object size discrimination, but that they were drastically worse at object orientation discrimination, perhaps suggesting that orientation perception relies more heavily on the visual system.

On the other hand, blind individuals unsurprisingly suffer from deficits in a number of skills ranging from auditory and spatial localization to navigation (see Cuturi *et al.* (2016) for a review on this topic). Recent studies show that early blind individuals fail in localizing sounds under particular auditory settings (Gori et al., 2010; Vercillo et al., 2015). Hollins and Kelley (1988) observed an interesting discrepancy between early-blind adults' ability to point to an object on a table after having walked partway around the table, and their ability to place the object in its original position after having walked around the table; whereas blindfolded sighted people performed well in both of these tasks, early-blind adults performed well at replacement but not at pointing. It has also been shown that while early-blind and blindfolded sighted adults are both better at remembering symmetrical configurations of blocks on a table than at remembering asymmetrical ones, the early-blind group was worse at remembering vertically symmetrical configurations than horizontal ones (Cattaneo et al., 2010). The authors proposed that the blindfolded-sighted adults were able to benefit from enhanced attention to vertical symmetry granted by the visual system.

In some cases, it does not appear to matter when a person first became blind in order to have reduced performance at a skill. Gentaz and Hatwell (1998) observed no differences between early-blind and late-blind adults in their ability to match the orientation of a rod to

that of a reference rod, but concluded that their reliance on gravitational cues for information differed from that of sighted participants (Gentaz & Hatwell, 1996).

2.1.2. Haptic space perception

When describing locations of objects, people can rely on two kinds of frames of reference. Using an egocentric reference frame, one might describe objects as “one meter to my left” or as “ten degrees to my right”. Using an allocentric reference frame one relies on external references, such as “one meter to the left of the laptop” or “below the table”. To manipulate an object, one needs to know its egocentric location, but when remembering or describing its position it can be more useful to rely on the allocentric location because it does not depend on one’s own spatial orientation. There is evidence that eye-centered coordinates are fundamental for the association of sensory signals (Cohen & Andersen, 2002; Jay & Sparks, 1987; Pouget et al., 2002), and the visual modality might offer a spatial background for remapping sensory information to obtain stable externally defined coordinates when one’s spatial orientation (including that of the eyes) changes. When the visual signal is missing, such spatial remapping may not occur. In agreement with this idea, congenitally blind individuals are not subject to the “crossed hand illusion” (Röder et al., 2004), presumably because they do not have the perceptual conflict with an externally anchored reference system for tactile stimuli. There is also experimental evidence that congenitally blind individuals do not remap auditory stimuli onto externally defined coordinates (Röder et al., 2007). It has been suggested that it is specifically *early-blind* people whose spatial experience of the world depends largely on egocentric reference frames rather than allocentric reference frames (Iachini et al., 2014; Pasqualotto et al., 2013). Therefore, one might generally expect the early-blind to rely more on egocentric reference frames for judging spatial relations.

It has been shown that two bars on a tabletop that are felt to be parallel can actually have dramatically different angles from one another, and that this difference scales almost linearly with the angular distance between them (Kappers, 2003; Kappers & Koenderink, 1999). The researchers concluded that participants’ perception of haptic space relied on a combination of allocentric and egocentric reference frames, the latter of which was centered on the hand. We expect that similar deviations will be found for drawing lines in a specified direction.

Using a similar parallel-bars paradigm, Postma et al. (2008) showed that when instructed to wait ten seconds between exploring a reference bar and rotating a test bar, early blind participants relied more on an egocentric reference frame than late-blind and blindfolded sighted adults. Based on this increased reliance on an egocentric reference frame, we expect early-blind adults who are instructed to draw straight fronto-parallel lines starting at different distances from the body midline to draw lines with more strongly deviating angles than sighted people. Similarly, when instructed to draw lines of a given length at different distances from the body, early-blind adults’ drawn lines may be longer at larger distances from the body due to relying more heavily on an egocentric reference frame, at least if such a reference frame is in polar coordinates (direction and distance).

Kuling et al. (2016) have shown that sighted adults are not always accurate at placing their hand on a table above the location of their other, unseen hand under the table. Kuling et al. (2014) showed that matching a target using a handheld pointing tool does not result in differences in magnitude or variability of error compared to using one’s own fingertip in a

visuo-proprioceptive matching task. Furthermore, while it has been shown that using proprioception alone in a matching task is less effective than combining sight and proprioception (van Beers et al., 1999), it also remains to be tested whether early-blind adults' proprioceptive accuracy and precision in a position matching task are worse than that of blindfolded sighted adults. We speculate that this matching error is already ameliorated by way of the visual system, and that early-blind adults may show greater matching errors than blindfolded sighted adults due to a lack of visual experience.

2.2. Methods

2.2.1. Participants

All participants in the experiment ($n=132$, mean age 48.5 years, $SD = 17.0$, 22 left-handed, 86 females) were attendees of the "ZieZo Beurs", a convention for blind and visually impaired people, who visited the Utrecht University booth that was set up for the event. All participants were naïve to the purpose of the experiment, with the exception of one participant with normal vision who was also an author. Participants were given an explanation of the task and were asked for their informed consent. Once verbal consent was given, they were asked several general questions about their vision, age and handedness.

For the analyses, the participants were grouped based on level of vision. The five groups were referred to as early-blind, late-blind, low vision, high vision, and sighted. Early-blind participants were all born completely blind and have never experienced any level of vision. Late-blind participants all became completely blind either due to a congenital condition leading to gradual loss of vision or due to an accident. The age of total vision loss in the late-blind group ranged from 2 to 47 years. Participants in the low vision group reported having some degree of visual perception, but no more than 10% in either eye. Participants in the high vision group reported having more than 10%, but no more than 80% vision in either eye. Participants in the sighted group reported having normal or corrected-to-normal vision. Level of vision was determined by self-report. The low vision and high vision groups consisted of participants who possessed congenital visual deficits as well as deficits acquired later in life by various means.

Due to the nature of the convention, participants were not actively matched between groups, and it was not possible to perform an in-depth intake session with each participant. We therefore lacked information regarding our participants' specific visual capacities, as well as haptic processing abilities such as the ability to read braille, which could potentially have an impact on their ability to perform the tasks described here. Details of the participant groups can be found in **Table 2.1**.

<i>Group</i>	<i>n</i>	<i>Age (mean ± SD)</i>	<i>Left-handed</i>	<i>Female</i>
<i>Early-blind</i>	7	43 ± 17	1	5
<i>Late-blind</i>	9	64 ± 11	2	5
<i>Low vision</i>	30	51 ± 12	8	15
<i>High vision</i>	31	55 ± 17	6	18
<i>Sighted</i>	55	42 ± 18	5	43
<i>Total</i>	132	49 ± 17	22	86

Table 2.1: Group-level characteristics of participants.

2.2.2. Experimental set-up & procedure

The set-up consisted of a 100 cm x 65 cm wooden board on which flip chart paper sheets could be mounted. The board rested on two supporting trestles, resulting in a table (Fig. 1). Four metal washers were placed under the board as haptic targets for the non-dominant hand. The washers had a diameter of 3 cm with a 1 cm gap in the middle, such that participants could comfortably guide their fingertip to the center of the washer. Two of these washers were placed along the center of the board, at 15 and 35 cm away from the long edge (henceforth referred to as the near and far targets, respectively). The remaining two washers were placed 20 cm to the right and 26.6 cm to the left of the near target, respectively (henceforth referred to as the side targets). The different lateral distances of the two side targets was an error that we compensated for in the analysis. The right side target was only used by right-handed participants, and the left side target only by left-handed participants. An aluminum bar of 10 x 1 x 1 cm served as a reference object, which participants held briefly at the beginning of the experiment.

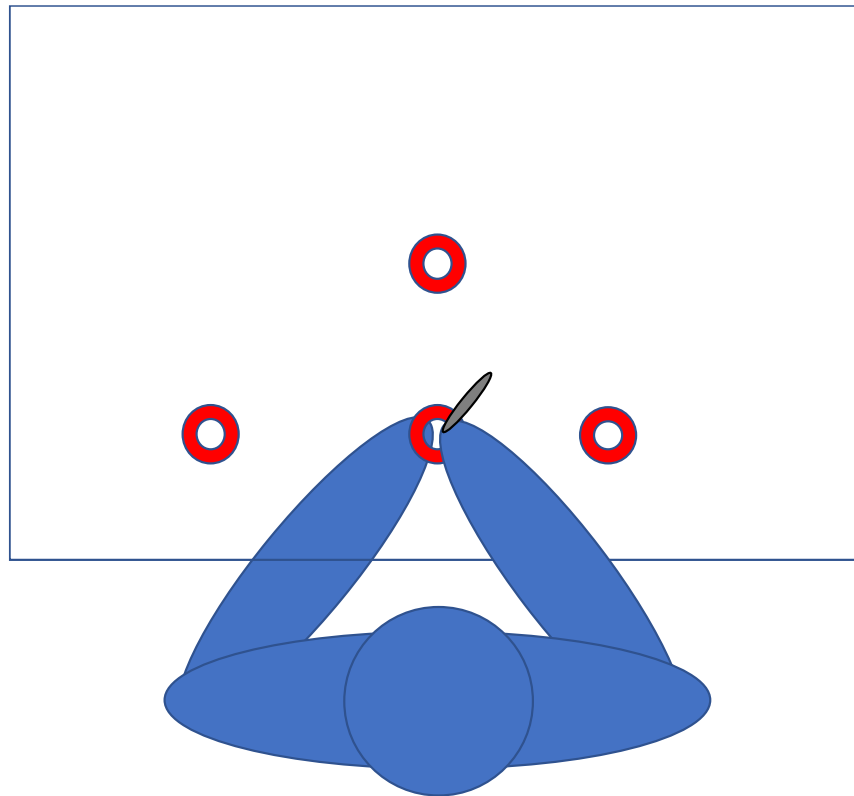


Figure 2.1: Top view of a participant at the experimental table. The haptic targets, placed under the table, are indicated by red rings, contacted by the index finger of the non-dominant hand. The dominant hand was holding a pen above the table. Target size is misrepresented here for the sake of visibility.

Participants sat in a chair in front of the setup and were blindfolded. The participants sat such that their mid-sagittal body plane was aligned with the near and far haptic targets. The task was explained, and participants were invited to explore the table with their hands to get a feel for the positions of the haptic targets. Participants were asked to confirm that they could reach all three targets before continuing. The experimenter handed the reference object to the participants and instructed them to get a feel for the length of the object. Participants were allowed to hold it for several seconds in their dominant hand and manipulate it. They were not allowed to set it down on the table, run their finger along the length of it, or touch it with their non-dominant hand. Most participants held the object between their thumb and pinkie finger to get a feel for the object's length; they were not restricted to holding it in their palm. The experimenter explained that the participants' task would be to draw lines of this length. The experimenter took the reference object back once the participants said they had an idea of the length (typically after five to ten seconds) and confirmed that they understood the task.

The experimenter then placed an ink marker in the participants' dominant hand and guided the index finger of their non-dominant hand toward one of the three haptic targets.

Participants were asked to use the marker to indicate on the paper where their non-dominant hand was placed, and from this point draw a line with the same length as the reference object. Three targets were used: the near and far targets, as well as the side target that corresponded with their dominant hand. Right-handed participants were asked to draw all lines from left to right, and left-handed participants were asked to draw all lines from right to left. All participants thus drew three such lines, one from each of the three targets in a randomized order, without any practice or familiarization. Because we did not know in advance how many participants we would be able to recruit, we did not counterbalance the order of the targets across participants. After a participant was finished, the sheet of paper was removed, and a fresh sheet was placed on top of the setup for the next participant.

2.2.3. Analysis

For each participant and each target, we determined the position of the dot at the start of the line, the matching error (i.e. the two-dimensional vector between the haptic target and the dot at the start of the line), the length of the drawn line (i.e. the shortest distance between the start and the end of the drawn line) and the angle of the drawn line relative to the coronal plane. For each of our analyses, we examined whether the early-blind performed differently than the sighted group. We used one-tailed independent samples t-tests to do so, as we expected early-blind participants to show greater matching errors, length ratios between near and far target lines, and angular differences between near and side target lines than sighted participants. In cases where a significant difference was found, we also performed one-tailed t-tests to determine whether the early-blind group differed from the other groups.

In this paper, four questions are investigated:

- 1) Do early-blind participants show greater matching errors than sighted participants?
- 2) Do early-blind participants show less matching error consistency than sighted participants?
- 3) Do participants draw longer lines from far targets than from near targets and is this difference larger for early-blind than for sighted participants?
- 4) Do participants draw differently oriented lines from the near and side targets and is the angular difference larger for early-blind participants than for sighted participants?

2.2.3.1. *Matching errors*

We first analyzed the magnitude of the matching errors of each participant group, ignoring the direction of the error. To determine whether participants' matching errors (both magnitude and direction) were consistent across targets, we calculated a consistency value as described in previous studies (Kuling et al., 2016, 2017). The consistency value was determined for all three combinations of the three target positions (i.e. ME_{near} compared with ME_{far} , ME_{far} compared with ME_{side} , and ME_{side} compared with ME_{near}), after which an average value for each participant was calculated. In order to determine a meaningful baseline, we also calculated a chance consistency error by comparing every individual matching error with every other matching error; a value above this baseline indicates that participants are at least somewhat consistent in their errors across targets (Kuling et al., 2016, 2017). We analyzed

whether early-blind participants had significantly larger matching error magnitudes and lower consistency values than sighted participants.

2.2.3.2. *Line analysis*

For the comparison of the drawn lines, we looked at drawn line lengths starting from the near and far targets, as well as drawn line angles starting from the near and side targets. For the lengths, we analyzed whether the ratio of the far target line length to the near target line length was larger than 1, and whether it was larger for the early-blind than for the sighted participants. For the angle, we examined whether the angular differences between drawn lines at the near and side targets were greater for early-blind participants than for sighted participants. To correct for the difference in relative position of the left and right side targets, we multiplied left-handed participants' values by 20 cm / 26.6 cm, as it is known that the angle at which a line is felt as fronto-parallel scales linearly with distance from the body midline (Kappers, 2002).

2.3. Results

The data of three participants (two low vision, one sighted) were completely excluded from the analysis because they drew lines from arbitrary starting points and reported that they did not realize they needed to match the haptic targets or pay attention to the reference object length. Data from a further ten participants (three late-blind, two low vision, three high vision, two sighted) were excluded from analyses of the drawn lines, but not from that of the matching task, due to the participants drawing one or more lines in the wrong direction or, in two cases, starting their drawn lines from somewhere other than the (correctly located) haptic targets.

The raw data of an example representative participant (**Figure 2.2**) show matching errors that are quite consistent over the three targets: the drawn lines all started about 3 cm from the target position (further away and a bit to the left). Furthermore, the lengths of this participant's drawn lines are about 6 cm, considerably shorter than that of the 10 cm reference object. The line drawn from the side target is rotated a few degrees clockwise relative to the one drawn from the near target.

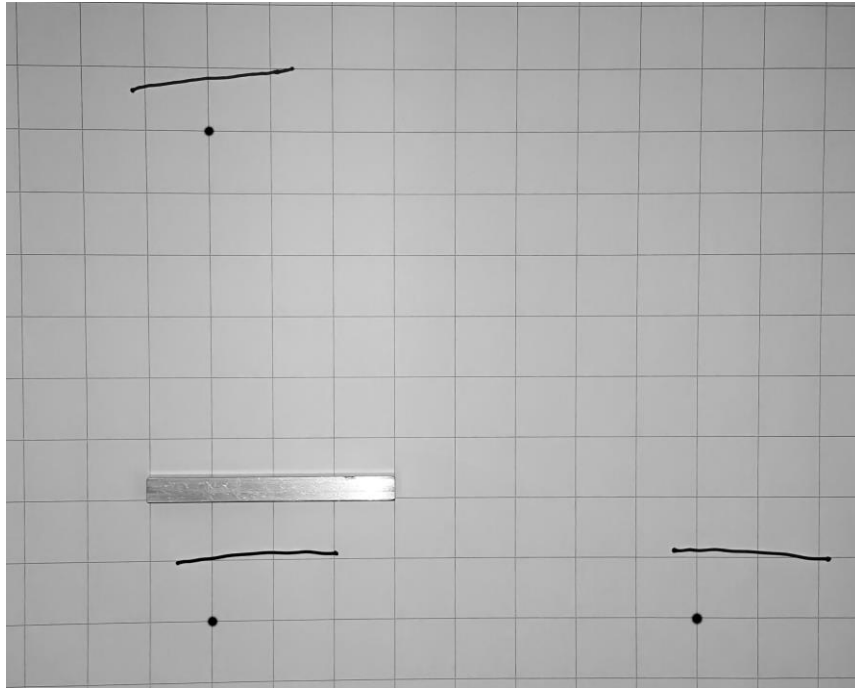


Figure 2.2: Example experiment sheet from a 20-year-old right-handed male in the sighted group with representative performance. Black dots indicate the location of the haptic targets under the table, and the reference object is displayed above the near target for comparison. The target dots and reference object were not present on the paper during the experiment. The grid squares have 2.5 cm sides.

The individual matching errors for right-handed participants show quite some variability, but there was no evident difference between groups (**Figure 2.3a**). We focus our analysis on the magnitudes of the matching errors (**Figure 2.3b**). We found that early-blind participants did not have greater mismatch magnitudes than sighted participants (one-tailed independent samples t-test; $t(59) = -0.184$; $p = 0.573$).

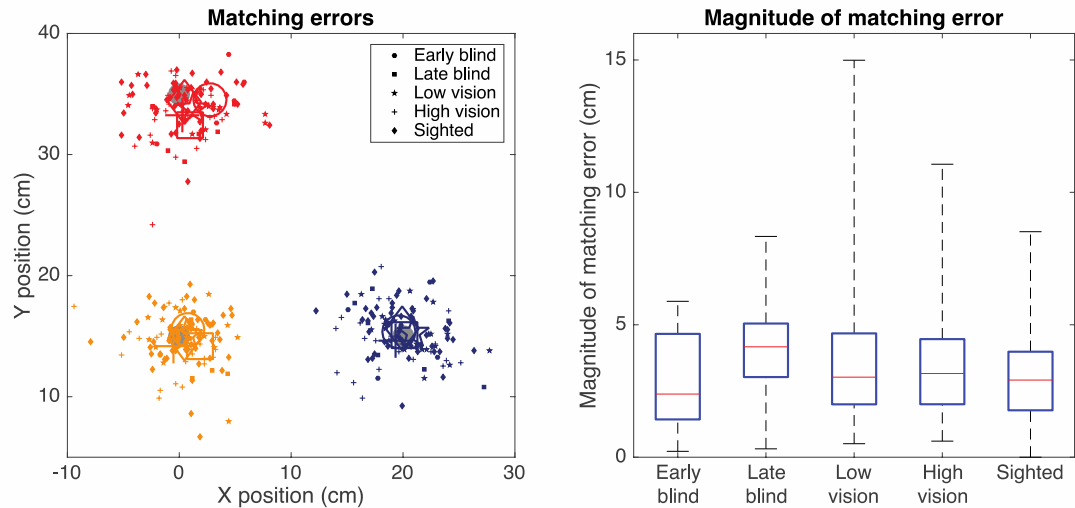


Figure 2.3: Matching errors. **a:** Matching errors of all right-handed participants. Left-handed participants' data are not shown here due to the positioning of the left side target. Target positions are denoted by large gray disks at (0,15), (0,35), and (20,15) cm. The large colored symbols denote mean matching errors for each group. **b:** magnitudes of matching error for all participants. Box plots show median (red line), interquartile range (blue box) and full range (black marks).

The consistency of the early-blind participants' errors across target positions is close to chance level, whereas the other groups show more consistency (**Figure 2.4**). Early-blind participants showed a significantly lower consistency value than sighted participants (one-tailed independent samples t-test; $t(59) = 2.264$; $p = 0.013$). The early blind also showed a significantly lower consistency value than the other three groups (late-blind: $t(14) = 1.189$, $p = 0.040$; low vision: $t(33) = 1.183$, $p = 0.038$; high vision: $t(36) = 2.894$, $p = 0.003$).

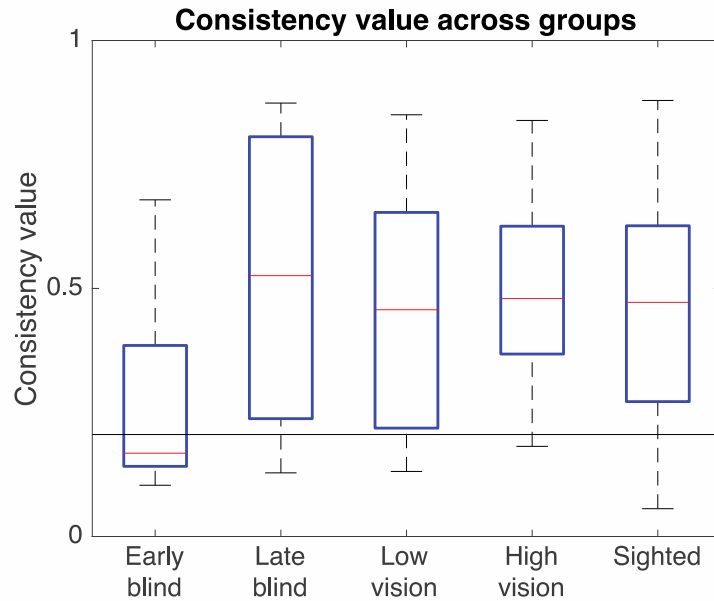


Figure 2.4: Matching error consistency values for all participants. The horizontal line indicates the median chance consistency value, 0.207. For an explanation of the calculated consistency values, see (Kuling et al., 2016; Kuling et al., 2017).

Most participants drew lines that were considerably shorter than the 10 cm length of the felt object (**Figure 2.5a**): on average only 6.59 cm (66% of the reference object length). They drew lines of similar lengths at the different distances: the average ratio of the far and near line lengths was 1.08 (**Figure 2.5b**). Early blind participants' line length ratios were not significantly larger than those of sighted participants (one-tailed independent samples t-test, $t(57) = 1.50$, $p = 0.07$). Furthermore, neither early-blind nor sighted participants' line length ratios were significantly larger than 1.0 (one-tailed one sample t-test, early-blind: $t(6) = 1.25$, $p = 0.13$; sighted: $t(51) = -0.04$, $p = 0.52$), indicating that there was no tendency to draw lines of a certain angular length, rather than ones that match the reference object's actual length.

Lines from the side target were drawn in a systematically different direction than those from the near target (**Figure 2.6**). The 5.3° angular difference was significant across groups (clockwise for right-handed participants and counter-clockwise for left-handed participants; one-tailed paired-samples t-test, $t(118) = 6.26$, $p < 0.001$). The angular difference between lines drawn from the near and side targets was not significantly larger for early-blind than for sighted participants (one-tailed independent samples t-test, $t(57) = -1.00$, $p = 0.84$). These findings indicate that while most participants draw lines at different angles when beginning in different directions from their body, in accordance with them relying to some extent on a (polar) egocentric representation, the angle was not particularly large in early-blind participants, suggesting that the early-blind do not rely more strongly on such a representation.

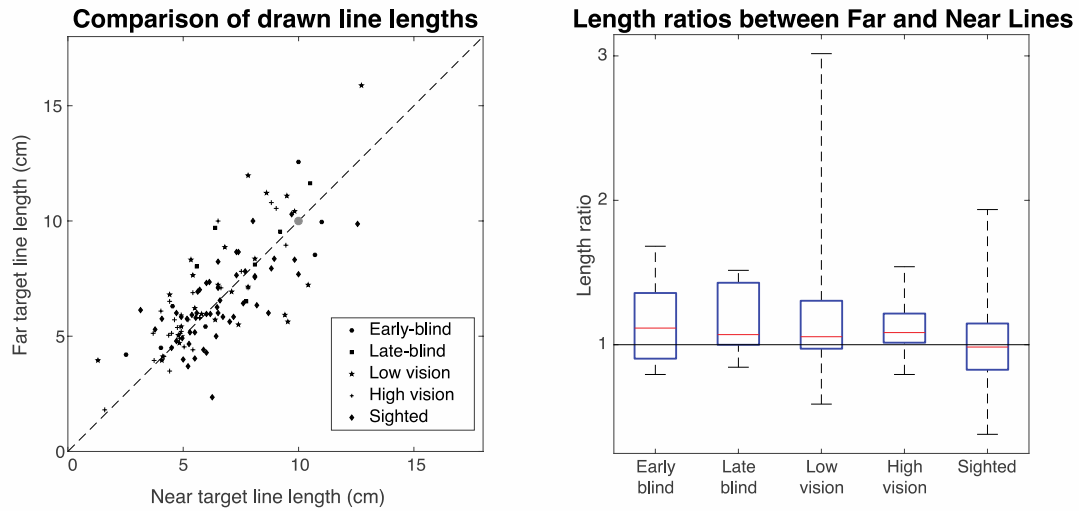


Figure 2.5: Length reproduction. **a:** Comparison of drawn line lengths from the near and far targets for all participants. The dashed line indicates unity, and the gray dot indicates perfect reproduction of line length from both near and far targets (10 cm). **b:** Box-plot of length ratios between the lines drawn from the far and near targets.

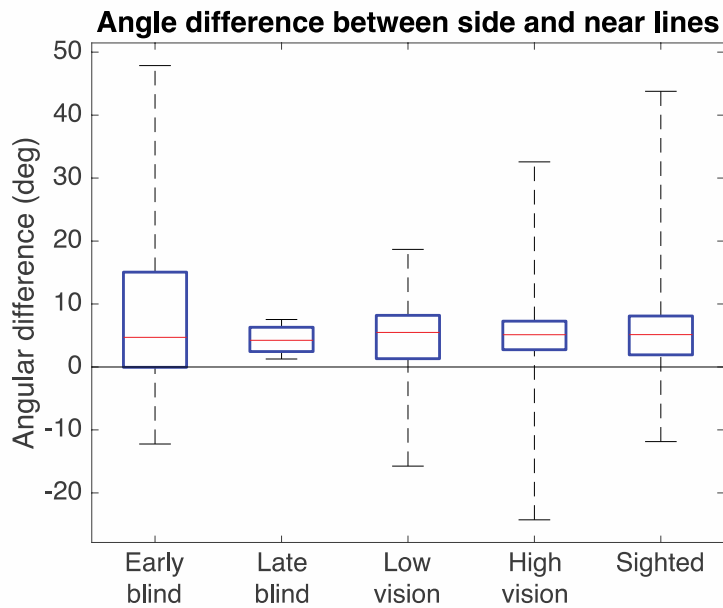


Figure 2.6: Difference in angle between near and side target lines for all participants. A positive value indicates a more clockwise orientation for a more rightward target position. Left-handed participant data has been normalized to correspond to a 20 cm difference between the near and side targets.

2.4. Discussion

We did not observe the differences that we anticipated between early-blind and blindfolded sighted participants. Both early-blind and sighted participants drew lines that are much shorter than the felt bar and were drawn at angles that similarly depended on eccentricity. As these effects were similar across groups, they can better be attributed to a general property of haptic perception than to an effect that depends on visual experience. The sizes of our participant groups were relatively small, so we may have failed to detect small but real effects. Our interest was in clear-cut effects of lack of visual experience, for which we investigated differences between early-blind and sighted participants. The expected systematic effects of lack of visual experience were not evident in our sample.

We are not aware of a theory that makes specific predictions for how the performance in our task would depend on details of the impairment such as the age of onset of blindness. Given the time-constraints inherent to data-collection at a convention, we did not collect data on the precise history of our participants' visual impairments. Therefore, we cannot perform any exploratory analysis on this. Similarly, we did not intend to investigate biases caused by other factors than being blind, so we did not examine whether participant height, hand size, or arm length played a role in task performance.

2.4.1. Matching errors

We observed average matching errors of three centimeters, which is in line with previous reports (Haggard et al., 2000; Kuling et al., 2016; von Hofsten & Rösblad, 1988). What is particularly interesting about this is that participants who have never experienced vision are just as accurate at matching their hand positions as blindfolded sighted participants, indicating that vision does not play a role, even indirectly, in this task. The low consistency value in early-blind participants (**Figure 2.4**) is remarkable. While the actual size of the matching error is well in line with that of other groups (**Figure 2.3b**), this error was not consistent across target positions in the early-blind participants, in contrast to the other groups. It is possible that growing up with visual input plays a role in developing this consistency, which early-blind people would therefore lack. However, given that the observed range of consistency values among early-blind participants falls within that of nearly all other groups, and given that the early-blind group consisted of only seven participants, it is not clear whether this difference is indicative of an actual matching deficit in early-blind individuals, or is simply due to chance.

2.4.2. Length reproduction

It is worth pointing out that our participants fall well short of the desired length when drawing lines, regardless of level of visual experience or starting position of the line (**Figure 2.5a**). This appears to be in contrast with the findings of Hermelin and O'Connor (1975), who observed that when congenitally blind and blindfolded sighted children were asked to reproduce vertical movements traveled by their hand, they were accurately able to do this for path lengths of 10-30 centimeters. Their task relied on kinesthetic movement information rather than on static tactile information. Our task involves the transformation of static tactile information (a felt bar in one's hand) to dynamic proprioceptive information. It is possible that this transformation is responsible for the systematic undershoot. They noticed a

difference in performance between blind and sighted participants when reproducing the longer movement distances (25-30 cm): the blind participants undershot these more than the sighted. As a hand cannot statically feel a length of 25-30 cm, this difference for longer movement distances cannot be tested with our paradigm of a hand-held length.

Previous work has shown that the horizontal-vertical illusion exists in the haptic domain as well as the visual domain. Heller and Joyner (1993) demonstrated that early-blind adults show evidence for a vertical compression of haptic space, but the size of the effect was small compared to the undershoot that we observe in our data. As such, we do not expect that the lengths of our participants' drawn lines would be substantially longer if we altered the orientation in which they were instructed to hold the reference object.

During the review process it was suggested that hand size may have influenced participants' perception (a smaller hand might lead to the percept of a longer bar) and thus on the length of the drawn lines. We did not measure hand size during the experiment, but analysed our data in light of gender (on average women have smaller hands). Indeed, it appears that women drew somewhat longer lines than men: the median line length (calculated over all groups) was 6.4 cm for women and 5.8 cm for men.

Lederman, Klatzky, and Barber (1985) tested participants' ability to reproduce a distance between two points by moving their index finger after traveling along an indirect path between the two points with that finger. They observed that early-blind participants tended to make larger reproductions of the distance than late-blind participants. They also observed an overshoot in reproduced length for all participants for lengths up to 25 cm, and an undershoot for greater distances. Since Lederman et al.'s participants traced a path with their finger while attending to the distance between the starting and ending positions, whereas our participants acquired length information without any tracing motion, our participants' systematic undershoot of the target distance cannot be related to this overshoot.

Moscatelli et al. (2014) tested whether participants could accurately reproduce lengths of a triangle path felt by a tactile stimulation device. Their participants could much more accurately reproduce the perceived displacement than our participants reproduced perceived length: their reported 98% accuracy is far higher than our 66% accuracy. A reason for this difference could be that the two experiments differed in how participants perceived the distances they had to reproduce. Their participants rested their index finger on a spherical device that rotated while the finger remained stationary on top of it, providing the tactile experience of moving a finger over a surface. Our participants were allowed to manipulate a 3D reference object in their dominant hand, but they were not permitted to run a finger along the length of the object.

The findings of Moscatelli et al. (2014), Lederman et al. (1985) and Hermelin and O'Connor (1975) all show that people are able to reproduce a given length reasonably accurately. In all these experiments, the participants could judge the length based on haptic motion information. As our participants made systematic errors and lacked this information, it may be necessary to have haptic motion information in order to accurately reproduce lengths in a drawing task. Irrespective of the reason for the distances being underestimated in our study, the fact that we observed no differences between early-blind and sighted participants in the ratio of lengths of lines drawn from the near and far targets (**Figure 2.5b**)

suggests that for this task, the early-blind do not rely more on an egocentric coordinate system than the sighted.

2.4.3. Parallel reproduction

Kappers (2004) demonstrated that haptic space perception relies on a combination of egocentric and allocentric reference frames, explaining the tendency of the angle of the test bar to lie somewhere between what would be considered parallel in purely allocentric or purely egocentric reference frames. Her experiments to map out haptic space were far more extensive than ours, but the differences that we see in drawn line orientation (**Figure 2.6**) are of similar magnitude to what she observes. Van Mier (2013) tested sighted adults' ability to draw a line parallel to a felt but unseen bar. Despite the fact that her participants were allowed to see the line they were drawing, there was still an angular difference between the orientation of the reference bar and that of the drawn line. Although this difference is smaller than what she observed in her purely haptic parallelity task, it still suggests that the reference frame experienced by each hand is important. This similarity is particularly interesting because we never explicitly told our participants to draw parallel lines; we simply always instructed them to draw the lines "from left to right" (or from right to left, if the participant was left-handed).

We observed no differences between early-blind and sighted participants in angular differences between the drawn lines from the near and side targets. This confirms that for this task, the early-blind do not rely more on an egocentric coordinate system than the sighted. The small but significant difference in angle between the drawn lines from the near and side targets can therefore be said not to rely on visual experience. Previous work has shown that women make substantially larger errors than men when matching the orientations of two haptically felt bars (Kappers, 2003). We did not observe a similar effect in the orientation of the drawn lines. Our female participants even showed a slightly smaller median angular difference than our male participants (5.0° and 5.5°, respectively).

Postma et al. (2008) observed differences between early-blind, late-blind, and sighted participants in a delayed bar parallelity task. When asked to wait ten seconds between feeling a reference bar and setting a test bar parallel to it, blindfolded-sighted participants performed better than late-blind participants, who in turn performed better than early-blind participants. The fact that we do not observe a similar effect, despite a delay between feeling the reference object and drawing the lines, likely stems from the fact that we do not explicitly ask our participants to consider the egocentric relation of the object to the participant. Instead, we only asked them to consider the length of the object, regardless of orientation, so the angle of the participants' drawn lines should have no relation to the angle at which they felt the object.

2.5. Conclusion

We expected that early-blind participants would rely more on information processed within an egocentric reference frame, but found no indication of this in our results.

2.6. Data

All data pertaining to this study, both analog and digital, are freely available for review. Interested parties are invited to contact the first author for specific questions or access to the data and materials.

2.7. Acknowledgements

This work was supported by the Horizon 2020 grant 642961: “PACE” and by the Dutch Technology Foundation STW grant 12160. The authors would like to thank Frans-Jozef Halkes for constructing the experimental table and reference object on very short notice and Nena van Balen, Lisa Broeders, and Malou Proper for their help with data collection. We would further like to express our appreciation to the organizers and participants of the ZieZo Beurs for making this experiment possible.

3

How well do we know when we act?

Adapted from:

Nelson, J. S., Brenner, E., & Smeets, J. B. J. (2020). How well do we know when we act? Submitted on 11 December, 2019 to *Human Movement Science*, currently under review.

Abstract

People can intercept moving objects with a much better temporal precision than that with which they can judge the relative timing of two external events. Is this because they are better at judging the timing of their own actions than at judging when external events happened? We developed a method to reliably present stimuli at various moments within 50ms of when participants tapped on a surface. Participants then tapped and judged whether a flash or a beep occurred before or after their tap. We also presented a flash within 50ms of a beep and asked participants to judge which was presented first. Participants detected whether the beep was before or after the tap reasonably precisely (standard deviation of about 70ms). Their precision was lower for the flash (about 170ms). They were very poor at judging whether a flash occurred before or after a beep, much poorer than one would expect from their precision in comparing the timing of a presented stimulus to that of a tap. Our results cannot be explained by a combination of three independent judgments of time, one each for the flash, beep, and tap. This poses a fundamental problem for explaining the precision of interception.

3.1. Introduction

Actions such as hitting a baseball can only be successful if they are timed to an external stimulus (the ball's motion) with high precision. Such actions can be timed to within 6 ms (Eli Brenner et al., 2012), which is as precise as the timing of different aspects of the movement in dart throwing (Smeets et al., 2002) or of the two hands in synchronous tapping (Doumas et al., 2008). Judgments of the relative timing of two events that are not related to one's own actions are generally very much poorer (Brenner & Smeets, 2010; Matthews et al., 2016; Virsu et al., 2008). In simple interception tasks the action is not always timed to within 6 ms, because it is influenced by many factors (Eli Brenner & Smeets, 2015). Importantly, when asked to make judgments about the timing of their performance on a slightly less precisely timed interceptive action, the judgments were more precise than the performance. In that study the judgments had a temporal precision (when we mention a value for precision we are always referring to a standard deviations) of less than 15 ms (de la Malla, López-Moliner, & Brenner, 2012). Does this mean that people are particularly good at judging timing with respect to their own actions? Is this because they are extremely good at judging when they acted?

An essential component in the precise timing of motor tasks is adjusting one's movement to the continuous stream of visual information, both about the target and about one's own hand (Eli Brenner & Smeets, 2018). In this way, the hand's position can continuously be adjusted to judgments of a target's motion to ensure successful interception (Eli Brenner & Smeets, 2015). In other timing tasks, such as tapping rhythmically, such continuous information about the 'target' is absent. Indeed, when producing a single tap in synchrony with a rhythmic stimulus, the precision is worse than 30 ms (Bartlett & Bartlett, 1959). In this case additional variability is presumably introduced by having to judge the rhythm itself and having to judge the time that elapsed since the last component of the rhythm (Doumas et al., 2008; for reviews see Repp, 2005 and Repp & Su, 2013). The goal of the current study was to determine how well we can judge when external events occurred with respect to a simple action. Is it close to the 6 ms mentioned above, or can such high precision only be achieved through continuous adjustments?

One way to evaluate timing precision is to ask participants to judge the temporal order of two events. One can for instance ask participants to tap on a surface, present a short visual stimulus (flash) around the time of the tap, and ask the participants whether the tap occurred before or after the flash. One can do so for various values of the tap-flash asynchrony, and use this to determine the precision of judging temporal synchrony for this particular task. The value that one obtains represents a combination of the precision of judging the moment of the tap (σ_T) and the precision of judging when the flash occurred (σ_F). Assuming that these are the only relevant factors and that they are independent, precision on this task (σ_{FT}) will be given by:

$$\sigma_{FT}^2 = \sigma_F^2 + \sigma_T^2 \quad (1)$$

If one repeats the experiment, but now presents an auditory stimulus (beep) instead of a visual one (the flash), one can express the precision in terms of the precision of the tap and that of the beep (σ_B):

$$\sigma_{BT}^2 = \sigma_B^2 + \sigma_T^2 \quad (2)$$

In order to get an estimate for the precision of the perceived timing of the flash and the beep, we can also perform an experiment in which we present the flash and the beep, and participants have to judge which one was presented first. The precision in this task is given by:

$$\sigma_{BF}^2 = \sigma_B^2 + \sigma_F^2 \quad (3)$$

Combining these three equations yields a value for the precision of perceiving the moment of one's tap:

$$\sigma_T = \sqrt{\frac{\sigma_{FT}^2 + \sigma_{BT}^2 - \sigma_{BF}^2}{2}} \quad (4)$$

It is known that the precision in timing judgments differs between participants and depends on the exact parameters of the stimulus and on other conditions (Eagleman, 2008; Sowiński & Dalla Bella, 2013). We therefore let each participant perform all three tasks using the same range of temporal asynchronies (conditions).

Despite the many studies examining judgments of temporal synchrony between auditory, visual, and tactile stimuli, very few studies have examined the precision with which people can estimate the moment of their own motor action in relation to an auditory or visual stimulus (Rohde & Ernst, 2013; Sugano et al., 2010; Vercillo et al., 2017). These studies were directed at studying adjustments to temporal offsets rather than at studying precision. In the present study, we devised a method to reliably present stimuli before or after a participant's finger tap (as in Rohde & Ernst, 2013) and explored the precision with which they could determine the temporal order of these stimuli.

3.2. Methods

3.2.1. Participants

Sixteen participants took part in the experiment. All participants provided their written informed consent before the start of the experiment and after all questions had been answered. For two of these participants, the initial data we obtained were unusable due to unnoticed technical difficulties. Given that no information had to be withheld from the participants, we let them take part a second time. Three participants' data were excluded because the participants were unable to keep their finger steady enough for our method to reliably present the stimuli. Another two participants were excluded from analysis because they did not follow instructions during the experiment, and a final two opted to end the experiment early. This resulted in a total of nine usable participant datasets.

3.2.2. Setup

Participants sat at a table in a dimly lit room and rested their right hand on a wooden platform. They held their left hand on a laptop keyboard. An infrared marker was attached to the nail of the participants' right index finger. The position of this marker was recorded using an Optotrak 3020 movement registration system (NDI, Waterloo, Ontario, Canada) at 500 Hz. A cover prevented participants from seeing their hand (**Figure 3.1**). They were instructed to fixate on a green LED which was placed on the cover above the hand and delivered the visual stimulus (a 50 ms flash). Participants wore headphones for delivery of the auditory stimulus (a 50 ms 2100 Hz tone, henceforth referred to as a beep). A second LED was placed out of

sight of the participant. This LED flashed (50 ms) whenever our experimental software assumed the participant tapped their finger (see below).

The experiment was managed using PsychToolbox version 3.0.12 (Brainard, 1997; Pelli, 1997). The LEDs were operated and powered using a LabJack U3-LV (LabJack Corporation, Lakewood, Colorado, USA) and a custom Matlab LabJack toolbox (available at <https://preview.tinyurl.com/y65nqahe>). Voltages of the LED and headphones were recorded using an Optotrak Data Acquisition Unit (ODAU). This allowed us to determine the precise timing of the visual and auditory stimuli in relation to the position of the finger.



Figure 3.1. Experimental setup. Room and laptop screen are well lit for clarity. Participants placed their hand inside the pipe on the right. The green LED atop the pipe (the LED on the right within the area indicated by the green circle; the other, red LED was not used in this study) provided the visual stimuli.

3.2.3. Conditions and Procedure

Our aim was to determine the temporal precision of judging when the auditory, visual and motor events had occurred. We consider the onset of the beep, the onset of the flash, and the moment the finger first touches the surface as the moments of the auditory, visual and motor events. If the timing judgments are independent of each other, and comparisons only depend on the precision of the judgments involved, one can determine the precision of all three events by testing three comparisons: flash-tap, beep-tap, and beep-flash. We did so using a temporal order judgment task with the method of constant stimuli.

Presentations of the flash and the beep could be timed perfectly relative to each other at each asynchrony. Since the taps were generated by the participant, the time of the tap had to be predicted in order to present the flash or beep at the planned asynchrony. We did so on the basis of the movement of the finger, by performing a cubic fit in real-time to the most recent 70 ms of finger data. During the experiment, the height of the finger was constantly monitored and the cubic fit was performed to extrapolate the finger position and estimate

how far in the future the finger would reach the platform. Once the prediction returned a value of no more than 50 ms, we used that time as the predicted moment of the tap, and delivered the beep or flash relative to that moment. Since we predicted the time of the tap about 50 ms before it actually happened, we presented asynchronies in a range of ± 50 ms.

Participants took part in two sessions on different days. During the first session, they were first briefly trained on each of the three tasks (flash-tap, beep-tap, and flash-beep). During the training, each task was presented in a separate block of presentations. The blocks were presented in random order. The asynchrony between the two stimuli was set to its extreme, so there were only two options: +50 ms or -50 ms. In blocks involving taps, participants tapped their right index finger once and were presented with a single flash or beep. After each tap-stimulus pair, participants indicated whether the stimulus came before or after the tap by pressing 1 or 2 on a laptop keyboard with their left hand. In the flash-beep task, participants were presented with both a flash and a beep. They then indicated on the keyboard whether the flash or beep came first, again by pressing 1 or 2. Each block of the training session consisted of 50 stimulus pair presentations. Before, during, and after the training, participants were free to ask any questions about the experiment.

After the training was completed, participants were confronted with longer blocks of presentations for each of the three tasks, again in random order. For each task there were 11 conditions, that only differed in the chosen asynchrony between the stimuli (values between -50 and 50 ms in 10 ms increments). For each task there were a total of 275 stimulus pair presentations (25 presentations per condition). The second day was identical to the first, except that the training was omitted, and the participants completed the three blocks in the opposite order. At the completion of the experiment on the second day, any further participant questions were addressed.

3.2.4. Analysis

Tap times were determined from three sources of information (Schot et al., 2010): the position (height) of the finger had to be close to zero, the velocity had to be negative (downward), and the acceleration had to be positive and at a peak. The times of the onset of the flash and the beep were determined from the voltages that we recorded. We combined the data of the two sessions for each participant and analyzed the data per task. In order to determine the precision with which participants could correctly determine the order of stimulus events, we fit a cumulative Gaussian to the data using Matlab 2015b and the `psignifit` toolbox (Schütt et al., 2016). We used a maximum pool length of 0.1 ms, lapse and guess rates of zero, and a stimulus range of -1000 to 1000 ms. For the tasks involving taps, we used the actual asynchrony of each trial, regardless of the condition. Our main interest was in the precision as given by the standard deviation determined by `psignifit`.

3.3. Results

Our tap predictor worked well. The predicted tap time was slightly earlier than the actual tap time (median 4 ms), with a median standard deviation of 12 ms. The participants' precision in judging the temporal order of the stimuli (as determined with individual psychometric fits) varied considerably across tasks and participants (**Figure 3.2D**). In order to obtain an indication of the performance in each task, we plotted overall mean values for each

condition together with cumulative Gaussian curves based on the median point of subjective equality (-17, -8, and -44 ms) and precision (166, 68, and 521 ms) for the flash-tap, beep-tap, and beep-flash tasks (**Figure 3.2A-C**).

As our main interest was precision, we plotted its value for each participant in each of the three tasks (**Figure 3.2D**). The values for the beep-flash task were unexpectedly large for most participants. As our conditions ranged between -50 and 50 ms, it is difficult to interpret values that exceed 300 ms (three times the size of the range of our stimulus asynchronies), but it is clear that most of the values for the beep-flash task are much larger than the values for the other tasks. Based on equation (4), we can predict a maximal value for σ_{BF} (poorest possible precision) by assuming that the value of σ_T is negligible (we use $\sigma_T = 0$). These values are shown by the open symbols in **Figure 3.2D** (BF_0). The median of these predicted poorest possible precision values is 178 ms. For all except two participants, the predicted poorest precision is far better than the actual observed precision in the beep-flash task.

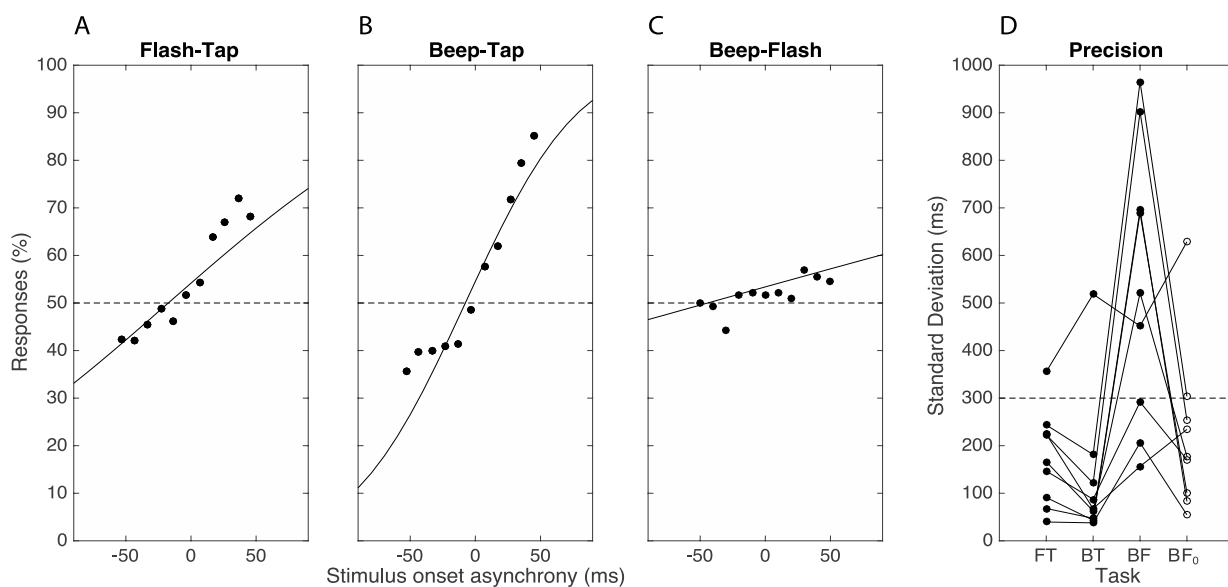


Figure 3.2. Results (A-C) Overall total fraction of responses that corresponded with what we considered to be a positive asynchrony for each condition of the three tasks. We considered the asynchrony to be positive if the event that is mentioned first in the name of the task occurred first. The points indicate the overall mean across participants. Horizontal positions correspond to the actual mean asynchronies. The curves are cumulative Gaussians, each based on the median value of the fit-parameters for the individual participants' data for that task. **(D)** The precision of the individual participants in each task. BF_0 (open disks) indicates the maximal expected value for the BF task (see equation 4). We consider the magnitudes of the precision to be quite unreliable when they are larger than about 300 ms (dashed line) due to the small differences between responses within our limited range of asynchronies.

3.4. Discussion

Our novel method for tap prediction allowed us to present auditory and visual stimuli at chosen moments before the participant's tap. The precision with which we were able to predict the moment of the tap was 12 ms. Since the prediction was based on fitting a curve at each instant it can deal with individual differences in how participants tap. Participants were able to distinguish whether an external stimulus preceded or succeeded their tap. They did so best for a beep, but they could also do so quite reliably when it was a flash. Most surprisingly, most participants performed extremely poorly on judging the relative timing of the flash and the beep. The performance was more variable than we could explain by assuming independent temporal judgements of three events. Consequently, we are unable to determine the precision of the timing of the tap.

The precision in comparing the timing of a beep relative to an active tap in our study is comparable to that found in studies using (passive) tactile stimulation (55-100 ms (Frissen et al., 2012; Harrar & Harris, 2008; Ley et al., 2009)). Most of our participants' precision was poorer when comparing the timing of a flash relative to an active tap than when comparing the timing of a beep relative to a similar tap. The precision also appears to differ more across participants for the flash than for the beep. Previous studies that compared the timing of a flash with (passive) tactile stimulation do not suggest that precision is systematically worse than for auditory stimuli (Harrar & Harris, 2008; Lange et al., 2018; Spence et al., 2001), so this may be specific to the particular visual and auditory stimuli that we used. Some previous studies that compared active tapping with auditory and visual stimuli found somewhat better precision than the studies that used passive tapping (Sugano et al., 2010), but others did not (Rohde & Ernst, 2013; Vercillo et al., 2017).

The precision that we found in the beep-flash task was much poorer than one would expect from judgments of the timing of the same stimuli relative to a tap. Moreover, the precision is poorer than values reported in the literature: for audiovisual temporal order judgements, reported values range from 30 ms (Zampini et al., 2003), or 50-75 ms (Keetels & Vroomen, 2012; Spence et al., 2003) up to 120 ms (Donohue et al., 2010). All our participants produced larger values. When asked to report whether two stimuli occurred at the same moment, rather than which occurred first, the reported precision is also well below 200 ms (Maij, Brenner, & Smeets, 2009; van der Burg, Alais, & Cass, 2013; Van der Burg, Orchard-Mills, & Alais, 2015; Virsu et al., 2008).

Why did we find less precise judgements in the beep-flash task than previous studies? The most conspicuous difference between our study and previous studies is that our study used a much smaller range of asynchronies than any other study. People adjust what they consider to be simultaneous presentation to the asynchronies that they experience, and this adjustment is larger when the asynchrony is not detected (Rohde & Ernst, 2013). This suggests that using a smaller range might give rise to poorer precision because there is more variability in what is considered to be simultaneous. Although it puzzles us why the precision in the beep-flash task was so much poorer than one would expect from the literature, this is less important than the fact that for most of our participants the precision of judgements of the temporal order of beeps and flashes is much more variable than one could maximally expect on the basis of performance in the other two tasks (BF_0 in **Figure 3.2**). As the same beeps and flashes were involved, and the maximal expectation is based on assuming that one knows

precisely when the tap occurred ($\sigma_T = 0$), the assumptions underlying equations 1-4 must be wrong.

One of the assumptions that might be wrong is that we assume that the precision of judging the time of a beep (σ_B) or a flash (σ_F) is independent of the task. Our results could be accounted for by assuming that precision improves when one executes a movement. There is some previous evidence that points in this direction. In an auditory judgment task, participants operating a robotic arm were more precise in reporting the duration of a tone when they were free to move during the tone presentation than when the arm was locked in place (Wiener et al., 2019). Participants also had smaller just noticeable differences in a haptic temporal order judgement task when moving their arm than when their arm was passively moved or when they were holding still (Frissen et al., 2012). A possible reason for having mechanisms that improve temporal precision by relating external events to ongoing actions is that temporal precision is probably most critical when guiding ongoing actions, especially when interacting with moving targets (Brenner & Smeets, 2015; Nelson et al., 2019).

A mechanism that relates external events to ongoing actions could be responsible for judgments as to whether one crossed a moving target's path before or after the target being quite precise (de la Malla, López-Moliner & Brenner, 2012). However, the better temporal precision in that task than in the Flash-Tap and Beep-Tap tasks of the current study might just be because the moving target was constantly visible, or because comparing events within a single modality yields better temporal precision than comparing between modalities (Virsu et al., 2008).

Using a novel approach for delivering stimuli before a participant completes an anticipated movement we showed that participants are able to judge the timing of their action relative to external stimuli better than judging the timing of those stimuli relative to each other. This difference was larger than can be accounted for on the basis of a constant precision for the timing of perceptual and motor events. We are therefore unable to determine the precision with which we time our actions.

3.5. Acknowledgements

The authors would like to thank Frans-Jozef Halkes for his extensive help with nearly all aspects of the complicated experimental setup, as well as Richard Cassius for writing the custom software we used to remotely operate the Optotrak. Finally, we would like to thank Heiko Schütt for his help with using psignifit. This work was supported by the EU Horizon 2020 grant 642961: "PACE."

4

How the modalities of a sensory rhythmic guide and of feedback influence rhythmic tapping

Adapted from:

Nelson, J. S., Brenner, E., & Smeets, J. B. J. (2020). How the modalities of a sensory rhythmic guide and of feedback influence rhythmic tapping. Submitted on 9 January, 2020 to *Journal of Motor Behavior*, currently under review.

Abstract

Synchronizing motor actions to an external rhythmic sensory guide is essential for performing actions such as dancing or making music. People's ability to do so depends on the sensory modalities of this guide. Do the sensory modalities that provide feedback about performance interact with the modality of the guide? We investigated this by manipulating the timing of auditory or visual feedback about one's performance. We devised a method to predict the time of a finger tap before it occurred. This allowed us to deliver feedback both before and after the time of the finger tap. We shifted this feedback while participants tapped along with an auditory or visual guide. The responses to such perturbations were analyzed for the four possible combinations of guide and feedback. Our participants most clearly adjusted their tap times to perturbed feedback timing for auditory feedback with a visual guide. Rather than compensating for the perturbation when timing their taps, they followed it, thus increasing the timing error. There was no consistent influence of the perturbation when the guide and the feedback were in the same modality. We discuss the implications of these findings for the way in which people adjust their tapping to performance feedback.

4.1 Introduction

Whether we are singing, dancing, playing a musical instrument, or just walking or speaking, we need to time our actions precisely. We often need to synchronize our actions to what happens in our environment, such as to the motion of an approaching ball, or to the actions of other musicians or of a dance partner. Our ability to do so is commonly studied in rhythmic tasks by asking participants to tap their finger (Delignières & Marmelat, 2014; Hove, Iversen, Zhang, & Repp, 2013; Repp, 2010a), press a button (Aschersleben et al., 2001; Blais et al., 2015; Repp, 2008), or tap with drumsticks (Fujii et al., 2011; Manning & Schutz, 2015) in synchrony with an external guide. The external stimulus can be auditory (Torre et al., 2013), visual (Blais et al., 2015), or haptic (Wing et al., 2010).

Across paradigms and studies, some findings prove to be very consistent. For instance, people tend to tap slightly earlier than they should (for a review of the literature, see Repp, 2005; Repp & Su, 2013), and both the extent to which people tap too soon and the variability in tap timing are larger when the interval between the stimuli is larger (Fujii et al., 2011; Mates et al., 1994; Yang et al., 2018; Zendel et al., 2011). Perhaps unsurprisingly, trained musicians tap closer to the correct time than non-musicians in such synchronization tasks (Repp, 2010; Repp & Doggett), with drummers synchronizing particularly well (Krause et al., 2010). A common way to try to determine what information is used for synchronizing tapping with a sensory guide is by manipulating either the guiding rhythm (Repp, 2000, 2002, 2011; Thaut et al., 1998; Torre et al., 2013) or the feedback that the participant receives (by generating an auditory or visual stimulus that supposedly occurs at the time of the participant's tap but actually occurs slightly later; Mates & Aschersleben, 2000; Pfordresher & Palmer, 2006).

When tapping in synchrony with a simple guide, one presumably adjusts one's movements so that the taps appear to occur at the same time as the guide. The timing of the taps is therefore likely to be influenced by any factor that influences judgments of simultaneity. When tapping with one's finger one has haptic information about the movement and the moment of the tap. In addition, one might see and hear the tap. It is known that judgments about simultaneity with an external stimulus can be influenced by adding visual or auditory feedback about the tap (Keetels & Vroomen, 2012b; Sugano et al., 2016). There is reason to expect auditory feedback to have a stronger influence on tapping times than visual feedback because auditory stimuli are judged more precisely (Welch & Warren, 1980; Nelson et al., submitted), so that a tendency towards optimally combining all cues to judge the time of the tap (Alais & Burr, 2004; Murai & Yotsumoto, 2018) will result in auditory signals being given more weight than visual ones.

It is likely that not only the modality of the feedback is relevant, but also its relation with the modality of the guide. Since detecting asynchronies is more precise within a single modality than between modalities (Virsu et al., 2008), we hypothesized that we would find a stronger influence of feedback in the modality that is used for the guide. To test this hypothesis, we conducted two experiments on rhythmic tapping, both of which explore the effects of small (<30 ms) temporal perturbations of feedback on tapping performance. In Experiment 1, the temporal perturbation of the feedback varied gradually across successive taps. In Experiment 2, the perturbations were introduced instantaneously. Each experiment consisted of four sessions, one for each of the four combinations of modalities (auditory and visual) for the guide and feedback.

4.2 Methods

4.2.1 Participants

Ten participants took part in Experiment 1. Seven different participants (including the first author) took part in Experiment 2. All participants were right handed and reported normal or corrected to normal vision and hearing. All participants provided written informed consent before the start of the experiment. Participants (except for the first author) were naïve to the purpose of the experiment and were debriefed after the experiment was finished.

4.2.2 Setup




Figure 4.1. Experimental setup. Room and laptop screen are well lit for clarity. Participants placed their hand inside the pipe on the right. The LEDs within the area indicated by the green circle atop the pipe could provide the visual guide (red LED) and visual feedback (green LED). The headphones could deliver an auditory guide and auditory feedback.

Participants sat at a table in a dimly lit room and rested their right hand on a wooden platform (**Figure 4.1**). They held their left hand on a laptop keyboard. An infrared marker was attached to the nail of the participants' right index finger. The position of this marker was recorded at 500 Hz using an Optotrak 3020 movement registration system (NDI, Waterloo, Ontario, Canada). A cover prevented participants from seeing their hand. They were instructed to look at two LEDs (one red, one green) which were mounted at the top of this cover (indicated by the green circle in the figure), directly above the tapping finger. These LEDs delivered the visual stimuli (50 ms flashes, red for the guide, green for the feedback).

Participants wore headphones for delivery of the auditory stimuli (50 ms beeps at 490 Hz presented to the left ear for the guide or at 2100 Hz presented to the right ear for the feedback). A third LED was placed out of sight of the participant and informed the experimenter about the moment that our software had anticipated that the participant would tap their finger (as will be explained below). The experiment was managed using PsychToolbox version 3.0.12 (Brainard, 1997; Pelli, 1997). The LEDs were operated and powered using a LabJack U3-LV (LabJack Corporation, Lakewood, Colorado, USA) and a custom Matlab LabJack toolbox (available at <https://preview.tinyurl.com/y65nqahe>). Voltages applied to the LEDs and headphones were recorded using an Optotrak Data Acquisition Unit (ODAU). This allowed us to determine the precise timing of the visual and auditory stimuli in relation to the position of the finger.

4.2.3 Conditions and Procedure

The task in both experiments was to tap in synchrony with a rhythmic guide. Depending on the condition, the guide was either presented as flashes or as beeps. In both cases the rhythm was a simple regular beat with an inter-onset interval of 650 ms. The feedback stimulus that was presented near the moment that participants tapped their finger could also be a flash or a beep. As such, each experiment consisted of four conditions: visual-visual, visual-auditory, auditory-visual, and auditory-auditory (henceforth referred to as VV, VA, AV, and AA, respectively). The first term denotes the modality of the guide, and the second term denotes that of the feedback. The conditions were presented in separate blocks, in an independent random order for each participant.

We intended to present the feedback stimuli at various delays relative to the time at which the participant tapped, including negative delays. Since the taps were generated by the participant, the time of the tap had to be predicted in order to present the flash or beep with the intended delay, especially when the flash or beep was to be presented before the tap. We used the same tap prediction method as in our previous work (Nelson et al., submitted). The height of the finger was measured every 2 ms. With every new measurement we used the most recent 70 ms of measurements of the finger's height to predict when the finger would hit the platform (zero height). We fit a third-order polynomial to the height measurements as a function of time, and determined the future zero-crossing of this polynomial. Once the time at which the polynomial crosses zero was less than 50 ms in the future, we used that time as the predicted moment of the tap for that tap, and delivered the beep or flash relative to that time.

In both experiments, the timing of the feedback was perturbed with respect to the predicted tap time. As a result, the feedback stimuli occurred shortly before or shortly after the participants' taps. In Experiment 1, the delay with respect to the predicted tap time varied gradually with each tap (sinusoidally, with an amplitude of 30 ms and a period of 20 stimuli). We opted for a sinusoidal variation as a function of stimulus number because this allows for an easy detection of motor adaptation (Hudson & Landy, 2012). The stimulus number at the beginning of each condition was determined at random from the 20 possible values. In Experiment 2, the delay remained constant for 20 taps (either +30 ms or -30 ms) and then switched to the other value for the next 20 taps. The stimulus number was either 1 or 21 at the beginning of each condition.

Each experiment was performed in a single session lasting roughly an hour. Participants began with a training block in which the feedback was veridical to get

accustomed to the stimuli and paradigm. The guide was presented first visually and then (once the participant was ready) auditorily during the training phase. The feedback was always auditory. They indicated verbally when they were comfortable with the task and were ready to begin the experiment, at which time the training block ended. Each of the four conditions was then presented in a separate block, which ended when the participant had tapped 550 times. Participants were informed that they were free to briefly pause their tapping during the experiment if they felt they had lost the rhythm or just wanted to rest.

4.2.4 Analysis

Tap times were determined from the finger position data using three sources of information (Schot et al., 2010). The probability of a moment being the moment of the tap increased linearly from 0 when the finger was 7 mm above the table to 1 when the finger was on the table. Values higher than 7 mm were given a probability of zero. Those lower than 0 mm due to skin compression were given a probability of 1. The probability of a moment being the moment of the tap also depended on the velocity of the finger. The probability was zero if the finger was moving upwards and increased linearly from 0.5 to 1 with the speed of the finger as a fraction of its maximal downward speed during that condition. Finally, the probability of a moment being the moment of the tap depended on the acceleration of the finger. It was zero except at peaks in upward acceleration, in which case it was the peak value as a fraction of the maximal upward acceleration during that condition. The moment of the tap was when the product of these three probabilities was maximal. The times of the onset of the flashes and the beeps were determined from the voltages that we recorded. As we are interested in the use of feedback, we binned taps based on the delay of the feedback received on the previous taps (20 bins for experiment 1, 40 bins for Experiment 2). The first tap from each block was excluded, as it had no preceding feedback. If the participant paused briefly during the experiment, the first tap after resuming was also excluded. For Experiment 1, we assessed the size of the influence by fitting a sinusoidal function to the binned and averaged tap data. For Experiment 2, we assessed the size of the influence by averaging the tap data bins for each of the two feedback asynchronies.

4.3 Results

The difference between tap and feedback time does not precisely match the intended delay, as the tap time calculated by the tap predictor does not always perfectly match the participants' actual tap time as determined in the manner described in the previous section. However, our tap predictor worked well. On average, the actual moment of a tap coincided with the predicted time (the predictor had a median error of being 0.1 ms early), with a median standard deviation of 8 ms. The performance of our predictor was much better than in our earlier study (Nelson et al., submitted), presumably because the single taps in that study were more variable than the rhythmic taps in the present study.

As was to be expected, participants tapped too early in all conditions. On average, the participants of Experiment 1 tapped 33 ± 13 , 77 ± 40 , 70 ± 25 , and 40 ± 14 ms too early (mean \pm standard deviation across participants) for the VV, VA, AV, and AA conditions (**Figure 4.2**). Thus, the tendency to tap too early was smallest (and most consistent across participants) when the guide and feedback were provided in the same modality (VV and AA). This is

consistent with the idea that alignment is most accurate within a single modality, so that adapting one’s actions to align the feedback with the guide is least biased when the guide and feedback are both beeps (AA) or both flashes (VV).

In line with the reasoning mentioned in the previous paragraph, we expected that the tapping would be adjusted to align the feedback with the guide, especially in the VV and AA conditions. Such adjustments would be visible as sinusoidally varying tap-errors that are in anti-phase with the delay in the feedback. Surprisingly, however, we see very little evidence of such adjustment. Only in the VV condition (top left panel of **Figure 4.2**) is there a slight tendency to tap later (black circles closer to zero) when the feedback is too early (most negative values of red diamonds), but the effect is not particularly compelling. The condition in which the response to the perturbation is by far the clearest is the VA condition (top right panel of **Figure 4.2**). In this condition the asynchrony in tapping is in phase with the perturbation, as if the delay of the feedback relative to a given tap drags the following tap in the same direction. The amplitudes (in ms) of the fits support this impression: -4 ± 12 , 31 ± 19 , 6 ± 9 , and 1 ± 10 , for the VV, VA, AV and AA conditions, respectively, with negative values corresponding to the predicted direction of the effect.

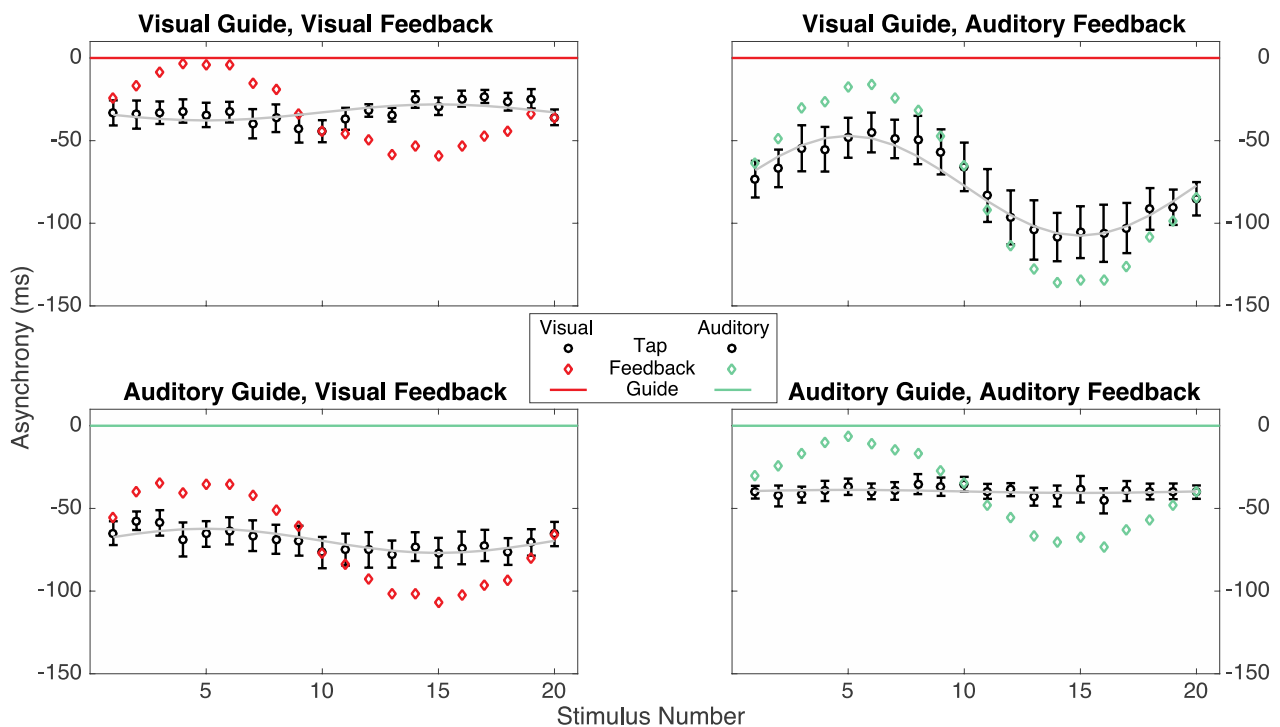


Figure 4.2. Mean asynchronies for the participants’ taps and times at which feedback about the taps were presented in Experiment 1 as a function of stimulus number (modulo 20). Feedback delay varied sinusoidally with stimulus number (period length 20 stimuli). Tap and feedback times are both presented relative to the timing of the guide (a tap error of 0 indicates that the participant tapped at the same time as the onset of the guide stimulus). Negative values indicate that the tap or feedback stimulus occurred earlier than the associated guide stimulus. As we expect the tap error to be influenced by the timing of the feedback of the preceding tap, tapping errors are aligned with the feedback on the previous tap (rather than with those of the current tap itself). The gray curve represents a sinusoidal fit to the data. Error bars indicate the standard error of the mean across participants’ mean values.

On average, the participants of Experiment 2 had similar biases as in Experiment 1, but with more variability across participants. They tapped 32 ± 52 , 53 ± 77 , 73 ± 37 , and $53 \pm$

37 ms too early (mean \pm standard deviation across participants) for the VV, VA, AV, and AA conditions (**Figure 4.3**). The effect of the perturbation was also similar to that in experiment 1: the amplitudes (in ms) of the effect were: 6 ± 16 , 29 ± 18 , 14 ± 11 , and 1 ± 20 , for the VV, VA, AV and AA conditions, respectively. Thus, again the feedback influences the tapping asynchrony most in the VA condition, to a lesser degree in the AV condition, and hardly at all in the VV and AA conditions. Moreover, again tap times mainly follow the feedback rather than correcting for it. As in Experiment 1, in the VA condition the shift in tapping asynchrony was substantial: it was again almost half of the perturbation magnitude. The overall tendency to tap too early was smallest for the VV condition and largest for the VA condition, but was nearly identical for the VA and AA conditions. The tendency to tap too early was therefore still smallest when the guide and feedback were provided in the same modality.




Figure 4.3. Mean times at which participants tapped and times at which feedback about the taps were presented in Experiment 2. Details as in Figure 4.2 except that the gray lines simply show the average tap error following each of the two feedback asynchronies.

4.4 Discussion

We anticipated that participants would combine the perturbed visual or auditory feedback that we provided with the haptic information about the timing of the taps that is always present. Doing so would provide a more precise estimate of the moment of the tap, which might make it possible to detect smaller asynchronies and therefore to tap closer to the guide. We hypothesized that this would be particularly effective when the feedback and the guide were presented in the same modality. We found that the tendency to tap too early was indeed smallest when feedback and rhythm were presented in the same modality.

However, on a tap by tap basis, the asynchrony was mainly influenced by the feedback when feedback and guide were presented in different modalities, and in particular when the guide consisted of flashes and the feedback was a beep (the VA condition). Moreover, the response to perturbing the feedback on the previous tap was to tap later if the feedback was particularly late, and earlier if the feedback was particularly early, rather than doing the opposite in order to correct for the apparent timing error.

Why do participants (in the AV and VA conditions) tap even earlier when the feedback on the previous tap suggested that they tapped earlier than usual, rather than tapping later to return to what they presumably consider to be the correct time? This might be because the visual or auditory feedback contributes to the judged onset of the time interval for the next tap, rather than being used as additional feedback about the tapping time. One can think of this as the feedback contributing to the judged timing of the guide, or at least to the judged timing of the onset of the time interval for the next tap. This would explain why the effect is strongest when auditory feedback (for which the timing can be judged quite precisely) is combined with a visual guide (for which temporal precision is lower; (Welch & Warren, 1980), because that is the case in which the feedback is likely to have the strongest influence. This reasoning also explains why the overall tendency to tap too early was larger in conditions in which the offsets in the feedback influenced the timing of the taps. The tendency to tap too early led to the feedback being early, so if the feedback contributes to the judged onset of the time interval for the next tap, that tap will be even earlier. Obviously, this will not continue indefinitely, because there is always also haptic feedback about the moment of the tap. The jagged pattern in the tapping asymmetries in the VA condition of Experiment 2 suggests that participants might start to detect the discrepancy between their tap time and the rhythm when the asymmetry reaches a certain value, and respond to this by ‘resetting’ their rhythm (Fitch & Rosenfeld, 2007; Mayville et al., 1999), only to be dragged along again in a somewhat systematic manner by the perturbed feedback.

The influence of feedback was clearly not strongest when the feedback was presented in the same modality as was used for the guide. Presenting both guide and feedback in the same modality does reduce the tendency to tap too early, presumably because the temporal error is judged most precisely when comparing signals within the same modality. However, the main influence of perturbing the feedback appears not to be mediated by detecting the error between the current tap and its guide, but by redefining the moment of the previous tap. Maintaining the same interval with respect to that moment gives rise to errors that follow rather than cancelling the asynchrony.

4.5 Acknowledgements

The authors would like to thank Frans-Jozef Halkes for his extensive help with nearly all aspects of the complicated experimental setup, as well as Richard Cassius for writing the custom software we used to remotely operate the Optotrak. This work was supported by the EU Horizon 2020 grant 642961: “PACE.”

5

The accuracy of intercepting moving tactile targets

Adapted from:

Nelson, J. S., Baud-Bovy, G., Smeets, J. B. J., & Brenner, E. (2019). Accuracy of Intercepting Moving Tactile Targets. *Perception*, 48(8), 685–701.
<https://doi.org/10.1177/0301006619861876>

Abstract

When intercepting a moving target, we typically rely on vision to determine where the target is and where it will soon be. The accuracy of visually guided interception can be represented by a model that combines the perceived position and velocity of the target to estimate when and where to hit it and guides the finger accordingly with a short delay. We might expect the accuracy of interception to similarly depend on haptic judgments of position and velocity. To test this, we conducted separate experiments to measure the precision and any biases in tactile perception of position and velocity and used our findings to predict the precision and biases that would be present in an interception task if it were performed according to the principle described above. We then performed a tactile interception task to test our predictions. We found that interception of tactile targets is guided by similar principles as interception of visual targets.

5.1 Introduction

Interception is a common action in daily activities such as trying to hit a moving ball when taking part in various sports. It can be defined as attempting to reach a specific place at a specific time in the future, coinciding with the arrival of another object at that point so that one makes contact with the object. Previous work has shown that the movements that we make when trying to intercept targets are constantly updated (Brenner & Smeets, 2018; Brenner, Smeets, & de Lussanet, 1998; de la Malla & López-Moliner, 2015; Smeets & Brenner, 1995). It has been suggested that doing so results in the final error mainly being determined by the visual information that is available about 100 ms before (Brenner & Smeets, 1997; Franklin & Wolpert, 2008; Oostwoud Wijdenes, Brenner, & Smeets, 2011; Smeets, Wijdenes, & Brenner, 2016). While much of the available research on interception is about using visual information to guide motor actions, it stands to reason that interception using information from other sensory modalities should also be possible (e.g. Shaffer et al., 2013). A common interception task within the animal kingdom is catching prey. Various species of animals employ a variety of senses to catch prey, including the use of electroreception, olfaction and judgment of hydrodynamic flow in some species of fish (Tricas, 1982; Nelson & MacIver, 1999) and echolocation in bats and whales (Johnson et al., 2004; Surlykke & Kalko, 2008). Shooing an insect that is moving along our arm on a dark night also requires that we have some sense of where the insect currently is, and possibly where it will be in the near future. We here confirm that tactile information can be used to intercept a target that one feels moving across one's arm, and examine whether such interception is controlled in a similar manner to interception of visual targets. Contrary to many studies that try to determine how movements are controlled by evaluating trajectories (e.g. Arzamarski et al., 2007; Fajen & Warren, 2007; Slupinski, de Lussanet & Wagner, 2018), we do so by evaluating the errors that people make.

When evaluating whether people intercept haptic targets according to the same principles as they do visual targets we must consider both precision and systematic errors. The precision of interception is limited by the resolution of the sensory and motor signals (Körding & Wolpert, 2006). For visual targets that people are tracking with their eyes, the sensory resolution depends on the resolution of the fovea and on how well people know the orientation of their eyes. When intercepting haptic targets, the sensory resolution depends on the part of the skin that is stimulated. In addition to the random errors due to limited sensory resolution, systematic perceptual errors are also common both when using virtual targets (Rossi et al., 2018) and in real life (Troscianko, Wright & Wright, 1999). They may be particularly evident under unnatural circumstances (La Scaleia et al., 2014). Misjudging acceleration (Senot et al., 2012) influences interception, as does relying on otherwise biased visual information (de la Malla et al., 2018). When judging positions from touch alone, judgments of the position on the arm are likely to be biased (Brooks et al., 2019; Green, 1982). When relying on proprioception to guide the hand that is used to intercept the target, the position of the moving hand is likely to be biased in a subject-specific manner that is highly stable over time (Kuling, Brenner, & Smeets, 2016; Smeets et al., 2006). The haptic percept of target speed might be biased as well, because we know that the reproduction of a haptic percept of length is biased (Nelson et al., 2018) and that the size of holes is judged differently with different effectors (Drewing, 2018).

To evaluate whether errors in the judgments that we consider to be essential for interception can account for the accuracy of interception we must consider how the judgments are used to control interception. Based on our knowledge of the accuracy of the

perceptual judgments of putative underlying information, we can make model predictions for the accuracy of interception, and compare these with that of actual performance. For visual targets we have done so both for the precision and the bias (Brenner & Smeets, 2015; de la Malla, Smeets, & Brenner, 2018). The simple model for interception that we use is based on continuously updating the perceived position, and extrapolating this position over the remaining time to interception on the basis of the perceived velocity. The position at which people aim to intercept a moving target is based on the sensed position \tilde{x} of the target (which might differ from the actual position x), extrapolated by the sensed velocity \tilde{v} (which might differ from the actual velocity v) for the remaining time until the tap. This aimed position is continuously updated, but the last update that can influence the outcome is one visuomotor delay τ before the tap. The parameter τ combines the dynamical properties of the eye, brain and arm into a single value that represents the last moment at which newly acquired information can influence the outcome. We justify the use of a constant value by assuming that judgments improve gradually enough for differences in the time needed to make the adjustments to be neglected. The tapping error E will therefore be equal to the error in the prediction:

$$E = \tilde{x}_\tau + \tilde{v}\tau - x \quad (1)$$

The bias in interception follows directly from systematic differences between the perceived and actual position and velocity. For the precision of interception, we base our prediction on the assumption of normal distributions of uncorrelated random errors in judging position and velocity (Brenner & Smeets, 2015):

$$\sigma_E^2 = \sigma_{\tilde{x}_\tau}^2 + \tau^2 \sigma_{\tilde{v}}^2 \quad (2)$$

We devised a series of three experiments to examine whether this interception principle that was developed to account for the interception of visual targets also applies to haptic interception. We probed biases and the precision of judgments about the position and velocity of a tactile stimulus on the arm, and measured people's ability to intercept a target moving along the arm. The first experiment presented participants with a moving stimulus on the arm, whose speed participants were instructed to reproduce by moving their finger through the air. The second experiment required participants to use their right fingertip to touch the location of a stationary tactile stimulus on the left arm. The third experiment required participants to intercept the target as it moved along their left arm. We examined whether performance on the first two experiments was consistent with performance on the third, as it should be if interception were controlled similarly to how it is done visually.

5.2 Methods

5.2.1 Participants

We performed three experiments: Position, Velocity, and Interception. Fifteen young healthy adults were recruited to take part in the Position and Velocity experiments. One participant stopped after a few trials because the participant had sensory deficits that made it impossible to feel the stimulus. The remaining fourteen participants completed the Position and Velocity experiments in random order in a single sitting. Seventeen young healthy adults took part in the Interception experiment. Three participants were excluded from analysis because the LED marker of the finger was hidden from the measurement cameras' views in

too many trials (see data analysis), and a fourth was excluded because of technical difficulties with saving data during the experiment. The remaining thirteen participants completed the Interception experiment in one sitting. All participants provided their written informed consent before the start of the respective experiments.

5.2.2 Set-up

We used a tapping task to measure how well people could match the position of a static tactile stimulus on their arm and intercept it when it was moving along the arm. We used a reproduction task to measure how well they could match the target's speed. Participants sat at a table in a room illuminated by a single lamp and rested their left arm palm-down on a platform on the table (**Figure 5.1**). The height of the far end of the platform was adjusted for each participant such that the top of the forearm was approximately parallel to the table. The tactile stimulus, delivered to the top of the forearm, was a bearing ball (4 mm diameter) which was free to roll within a fixture attached to a haptic device (an omega.3, Force Dimension, Nyon, Switzerland). The haptic device, capable of performing precise closed-loop movements, rolled the target along the participant's forearm. As each person's forearm has a different size and shape, the starting height of the fixture was calibrated for each participant individually. The position of the participant's forearm was such that the middle of the forearm corresponded to the middle of the haptic device's range of motion; the target never traveled farther than 10 cm in either direction from this central position.

Infrared LED markers were attached to the fixture and to the participant's right index finger to measure the interception error. In order to relate this error to the position on the left arm, we attached markers to the left elbow and major knuckle of the left index finger. The positions of these four markers were recorded at 400 Hz using an Optotrak Certus movement registration system (Northern Digital International, Waterloo, Ontario, Canada). Participants wore shutter goggles for the duration of the experiments. The experiments each consisted of multiple short trials. At the start of a trial the shutter goggles turned opaque and the lamp that illuminated the room switched off. This ensured that during the trial participants would have no visual information about either the target or their own body. Shortly after the shutter goggles turned opaque a stimulus was applied and the participant responded. Eight seconds after the start of the trial the lights turned back on and the shutter goggles cleared. Eight seconds was long enough to ensure that the participants had always finished their response before they could see again.



Figure 5.1. The experimental setup. Red dots indicate the position of the Optotrak markers, and white circles indicate the two starting points for the finger, one for the tapping movements of the Position and Interception experiments and the other for the movements in the Velocity experiment.

5.2.3 Procedure

5.2.3.1 Position experiment

Our first step was to measure how accurately and precisely participants could move their right fingertip to the location of a static tactile stimulus on their left arm. They started with their right index finger at the starting point on the table (**Figure 5.1**). The target moved to one of 21 positions above the arm (between -10 and 10 cm from the center, in increments of 1 cm) and then descended on the arm. The target then remained at this position for the duration of the trial. Participants were instructed to tap the target as soon as they felt it. Participants were instructed that corrective movements were not allowed. Because the target remained in place on the arm, participants touched the target in the case of a successful tap, and the arm otherwise. In this way they received feedback on their performance. For each position on the arm there were two trials, resulting in 42 trials per participant. Additionally, there was a calibration ‘trial’ at the very beginning and end of a session. During the

calibrations the participants had full vision and placed their finger directly on the visible target to ensure alignment between the finger and target markers. We considered the resulting separation between the marker positions as our definition of an error of zero in the data-analysis.

5.2.3.2 *Velocity experiment*

Our second step was to estimate how well people could judge and reproduce the speed of a target moving across their arm. Participants were instructed to pay attention to the velocity of the target moving across their forearm while resting their right index finger at a starting point in front of and to the left of the arm platform (**Figure 5.1**). At the start of each trial, the target would move into position above the participant's forearm, then descend and immediately begin traveling along the arm. The target started moving at one of six positions, 6, 8 or 10 cm from the center of the forearm in either direction. If the target started closer to the wrist it always traveled in the direction of the elbow, and vice versa. The target trajectory always ended 10 cm from the center of the forearm. The target could travel at one of three speeds: 10, 15, or 20 cm/s (henceforth referred to as slow, medium, and fast). It moved for 0.8-2 s, depending on the speed. For participants whose forearm was shorter than 20 cm, a scaling factor was applied to both the distances traveled and the durations such that the target never moved past the elbow or wrist, but still achieved the same velocities. Once the target completed its movement, participants were to move their right index finger rightward at a speed matching that of the target they just felt. They moved their finger to the right (orthogonal to the target's motion) and were instructed to keep moving as far as they comfortably could. In this way we ensured that they did not reproduce the target's motion rather than its velocity (for instance by matching the distance that it traveled). They moved in the same direction for all trials to minimize possible motor-related differences in reproduction across trials. Once the participant had completed the movement, they returned their finger to the starting point. There were five trials for each condition (categorized by one of six starting points and one of three speeds), resulting in a total of 90 trials per participant.

5.2.3.3 *Interception experiment*

Finally, we tested participants' ability to intercept a moving tactile target. As in the Velocity experiment, the target moved along the participant's arm. Again, the target could start 6, 8, or 10 cm from the center of the forearm in either direction, and traveled at 10, 15, or 20 cm/s, resulting in 18 unique conditions. The task for the participants was the same as that in the Position experiment: they started with their right index finger at the starting point on the table (**Figure 5.1**) and were instructed to tap the target as soon as they felt it. As a form of feedback, the haptic device stopped moving if it detected a sudden increase in force, such that a successful interception would halt the target's movement for the trial. If the participant tapped behind the target, the target continued its trajectory. If the participant tapped ahead of the target so that the target bumped into the finger, the collision triggered the device to stop. Corrective movements of the hand were not allowed. Trials were presented as ten blocks, each consisting of one trial from each condition in a random order. The first block served as a familiarization phase and was excluded from analysis. This resulted in a total of 162 trials that could be used for the analysis. Participants were not explicitly informed of the blocking or of the familiarization phase, but they were informed when they had made it halfway through the experiment and were free to take a break whenever they wished. As in the Position experiment, a calibration 'trial' took place at the very beginning

and end of each participant session, and the offset that was considered to represent perfect interception was calculated from these trials.

5.2.4 Analysis

5.2.4.1 Position experiment

We started the analysis by correcting the finger position for the difference between the finger LED marker and the fixture LED marker, averaged across the two calibration trials. The time of tap was defined as the point in time at which a sharp upward acceleration peak was observed for the finger while the finger was not moving upward (see Schot, Brenner, & Smeets, 2010). We considered the location of participants' initial tap on the arm as their tap location. Any subsequent movements or corrections were ignored, in line with the instructions to the participants. Positions were defined relative to the midpoint of the arm, which corresponded to the midpoint of the positions at which the targets appeared. Positive values were towards the wrist. We defined the participants' tap errors as the one-dimensional offset (along the path of the target) between the locations of the participants' finger and the target at the time of the tap. Errors were positive if the tap was too close to the wrist, in accordance with the coding of positions. In order to determine the bias and precision of the tap for each participant we determined the slope and residuals of a linear regression of their tap errors as a function of target position (see subsection 'Model predictions' for further details). Movement times were also determined for each trial.

5.2.4.2 Velocity experiment

Because finger velocity is not constant over time, we had to pick the segment of the finger velocity profile that we believed would best represent the participant's attempt to reproduce the perceived target velocity. We defined the velocity used to match the target as the average rightward speed (that is, only considering movement in one dimension) during the time that the participant's finger passed through the middle ten centimeters of their trajectory. This segment often, but not always, included the participant's maximum velocity. To check whether the actual velocity of the target corresponded to the programmed velocity, we determined the velocity of the infrared marker on the target by differentiating the marker's position with respect to time, again over the middle ten cm. The actual velocity of the target always matched the programmed velocity (mean \pm standard deviation across trials: 10.0 ± 0.02 cm/s, 15.1 ± 0.04 cm/s, and 20.1 ± 0.1 cm/s for slow, medium, and fast targets, respectively).

We determined this measure of perceived velocity for each trial. We then calculated a Weber fraction (the standard deviation divided by the mean) for each participant's velocity reproductions for slow, medium, and fast targets in order to get a measure of the precision of their velocity reproduction that is independent of the velocity of the target (Westheimer, 1999). Finally, for each participant we calculated the slope of velocity reproductions as a function of target speed to quantify any bias in reproducing felt velocities. We focused on target speed, rather than velocity, as we found that the target's movement direction had no influence on participants' reproduction speeds.

5.2.4.3 Interception experiment

As in the Position experiment, we considered the initial tap location and defined participant's tap errors as the offset between the locations of the participant's finger and the

target at the time of the tap. We calculated the slope and residuals of the participants' tap errors, now taking the target's movement direction into account. The slope was constrained to be the same for both directions of target motion, but the intercept was allowed to differ. As in the position experiment, movement times were determined for each trial.

5.2.4.4 Model predictions

We will compare our results from the Interception experiment to predictions based on the Velocity and Position experiments with the help of the model of equations (1)-(2). However, before doing so we must introduce a commonly observed bias in the responses: a contraction bias (Poulton, 1979). We assume that the perceived position (\hat{x}) is biased because tactile information about target position (\tilde{x}) is combined with a prior (p) that the target is at the center of the forearm (Angelaki et al., 2009). If w is the weight given to the prior:

$$\hat{x} = w p + (1 - w) \tilde{x} \quad (3)$$

so that if \tilde{x} is unbiased (on average, $\tilde{x} = x$), the systematic tapping error in the Position experiment ($E = \hat{x} - x$) will be :

$$E = w p + (1 - w) \tilde{x} - x = w (p - x) \quad (4)$$

We can determine w experimentally from the slope of the tapping error as a function of the target position.

Assuming that p and x are combined optimally and that their errors are independent, so that:

$$w = \frac{\sigma_{\tilde{x}}^2}{\sigma_p^2 + \sigma_{\tilde{x}}^2} \quad (5)$$

The variance of the participants' tapping errors (assuming that motor noise is negligible so that the variability corresponds with the variability in the perceived position \hat{x}) must be:

$$\sigma_E^2 = w^2 \sigma_p^2 + (1 - w)^2 \sigma_{\tilde{x}}^2 \quad (6)$$

Equations (5) and (6) can be combined to give

$$\sigma_{\tilde{x}} = \frac{\sigma_E}{\sqrt{1-w}} \quad (7)$$

and

$$\sigma_p = \frac{\sigma_E}{\sqrt{w}} \quad (8)$$

Thus, we can use the measured variability in the participants' tap locations (σ_E ; the standard deviation of the residuals of the fit in the Position experiment) and the weight given to the prior (w ; slope of error as a function of position) to estimate the standard deviation of the position judgments ($\sigma_{\tilde{x}}$) and of the prior (σ_p). We will use the same prior for the interpretation of the interception experiment and assume that it is combined with the estimate of where the target will be at the time of the tap. The anticipated tapping position \hat{x} will be based on the latest tactile information \tilde{x}_τ ,

$$\hat{x} = w p + (1 - w)(\tilde{x}_\tau + \tilde{v} \tau) \quad (9)$$

Assuming that the velocity is misjudged proportionally ($\tilde{v} = k v$), we get:

$$E = w p + (1 - w)(\tilde{x}_\tau + k v \tau) - x \quad (10)$$

$$\sigma_E^2 = w^2 \sigma_p^2 + (1 - w)^2 (\sigma_{\tilde{x}}^2 + \tau^2 \sigma_{\tilde{v}}^2) \quad (11)$$

In equation (11) we have neglected uncertainty in judging the visuomotor delay. We assume that the additional variability due to the motion of the target is a product of uncertainty in the perceived velocity and the delay itself ($\tau \sigma_v$). The standard deviation in the velocity judgments $\sigma_{\hat{v}}$ can be written as $v W_v$, where W_v is the Weber fraction for judging velocity that we estimated from the Velocity experiment and v is the velocity of the target in question. The weight w given to the prior in the interception experiment is expected to be larger than in the Position experiment because of the additional variability introduced by having to judge the velocity. With all the above-mentioned assumptions the optimal combination (highest precision) is achieved when

$$w = \frac{\sigma_{\hat{x}}^2 + (\tau v W_v)^2}{\sigma_p^2 + \sigma_{\hat{x}}^2 + (\tau v W_v)^2} \quad (12)$$

The higher weight given to the prior when intercepting moving targets, because $\sigma_{x_m} > \sigma_x$, corresponds with a steeper slope of lines fit to the tapping errors as a function of the target position at the time of the tap. When intercepting moving targets, we expect that the variance in participants' tapping errors will be:

$$\sigma_E^2 = w^2 \sigma_p^2 + (1 - w)^2 (\sigma_{\hat{x}}^2 + (\tau v W_v)^2) \quad (13)$$

where w is calculated using Equation (12), and $\sigma_{\hat{x}}$ and σ_p are the values from the Position experiment (Equations (7) and (8)). In our calculations we assume that participants have a similar sensorimotor delay for tactile and visual information (τ about 100ms). Biases in judging velocity or position will become apparent in the Position or Velocity experiment. Their effects on the anticipated systematic errors in interception will emerge directly when relying on equation (9).

5.3 Results

5.3.1 Position experiment

When our example participant attempted to tap the stationary target's position on the arm, the taps were biased in the direction of the center of the forearm (**Figure 5.2**). This pattern corresponds to the expected contraction bias: the farther away the target was from the center of the forearm, the larger this error, reaching values of up to 7 cm. We performed such a regression for each of the 14 participants. The (negative) median value of -0.22 for the slopes of these linear regressions (**Figure 5.3A**) corresponds to a weight w of 0.22. The median standard deviation of the residuals of these regressions σ_T is 1.5 cm (**Figure 5.3B**). On average, participants took 737 ms to tap on the target (from the moment the hand started to move).




Figure 5.2. A single participant's tap errors as a function of the location of the target. The solid line is a fit to the participant's data. The negative slope corresponds to a contraction bias.

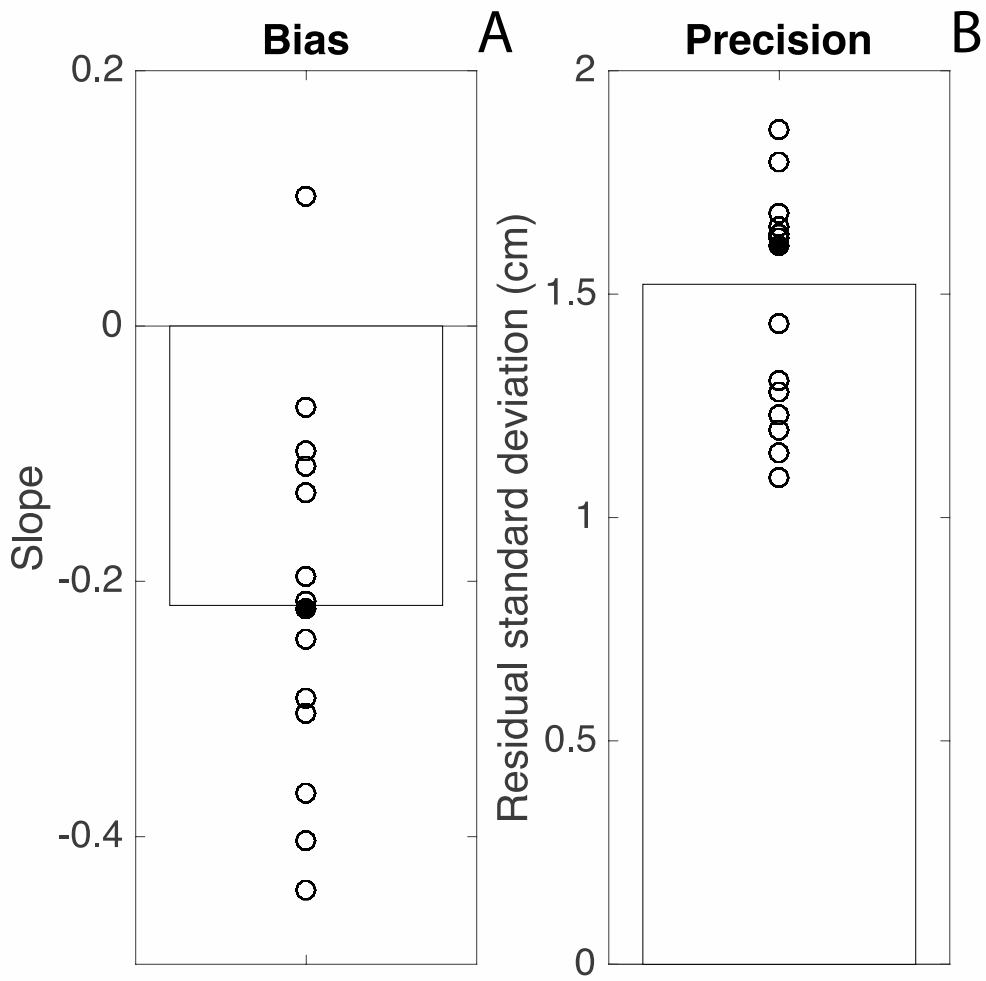


Figure 5.3. The bias (A) and precision (B) in the Position experiment. Circles: values for individual participants (the filled circles indicate the participant whose data are shown in Figure 2). Bars: median of all participants' values.

5.3.2 Velocity experiment

All except one participant clearly moved faster in response to faster targets. The direction of the target's motion had no effect on participants' speed reproductions, so we henceforth group the targets by speed rather than velocity. The data of a single participant (**Figure 5.4**) show that this participant (a different one than for the Position experiment) scaled the speed of the hand movement to the speed of the target. The variability in the reproduction was quite large. Moreover, the reproduction was not veridical (far from the dashed unity line): the reproduced speed for this participant was more than four times the speed of the target. This participant was no exception: all participants produced a speed that was more than twice that of the target (slope larger than 2, **Figure 5.5A**). The medians of the participants' average produced speeds were 35, 61 and 86 cm/s for the three conditions, with respective median standard deviations of 11, 16, and 19 cm/s. Participants differed by up to a factor 5 in their mean reproduced speed for a single condition. Participants' speed reproduction performance showed a median Weber fraction of 0.25 (**Figure 5.5B**). The median Weber fractions were 0.26, 0.25, and 0.21 for the slow, medium, and fast targets, respectively.




Figure 5.4. A single participant's reproductions of the target speed in the Velocity experiment. The solid line is a fit constrained to pass through the origin. The dashed unity line indicates where points would be if the participant reproduced the target speed veridically. This is a different participant from the one whose data are shown in Figure 5.2.

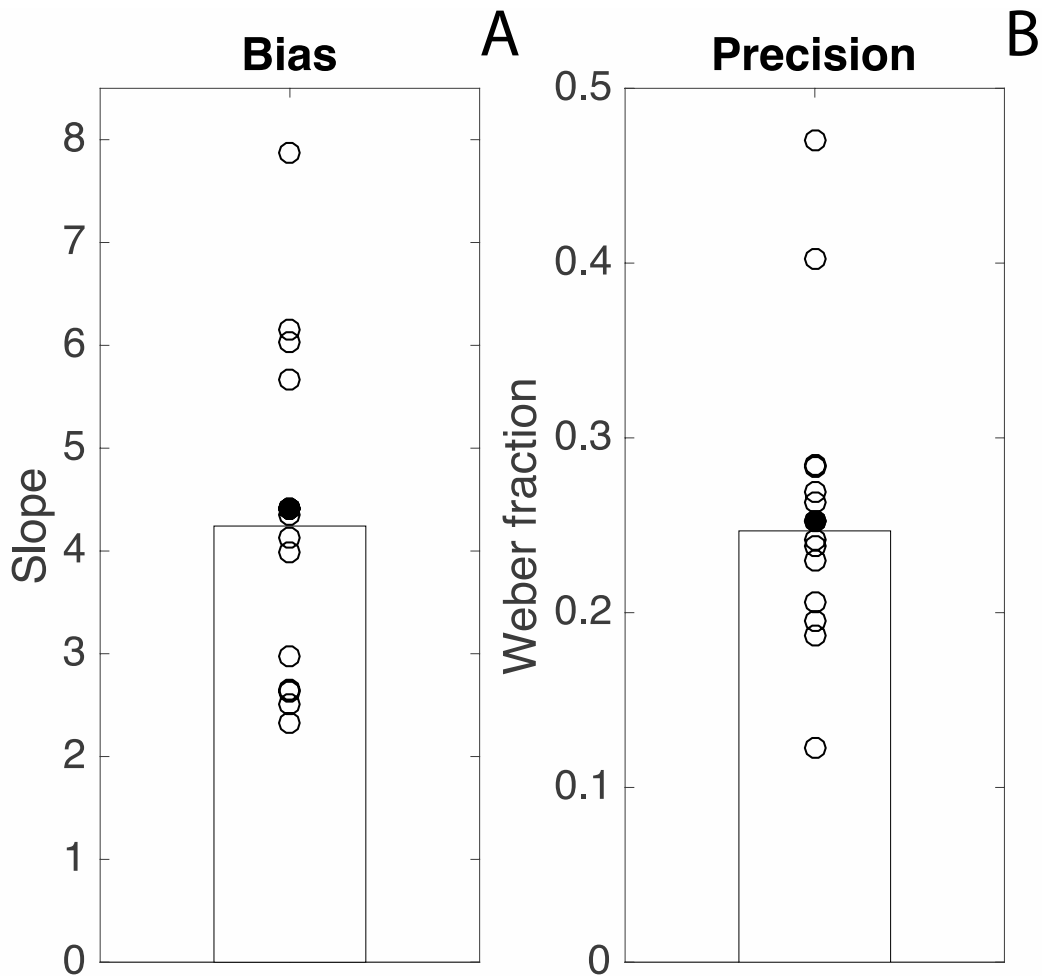


Figure 5.5. The bias (A) and precision (B) in the Velocity experiment. Circles: values for individual participants (the filled circles indicate the participant whose data are shown in Figure 5.4). Bars: median of all participants' values.

5.3.3 Interception: model prediction and experiment

On average, participants took 464 ms to tap on the target (from the moment the hand began to move). Using the median values obtained from the Velocity and Position experiments, we predicted a typical participant's performance in the interception task (**Figure 5.6A**). The predicted behavior resembles that of an example participant (**Figure 6.6B**). Two similarities with the predictions are remarkable. The first is the effect of the direction of target motion (separation between blue and green lines), which is introduced to the model by the overestimation of target velocity. The second is the steeper slope than in the Position

experiment (**Figure 5.2**), which is due to the larger uncertainty about where to tap, and thus a larger reliance on the prior.

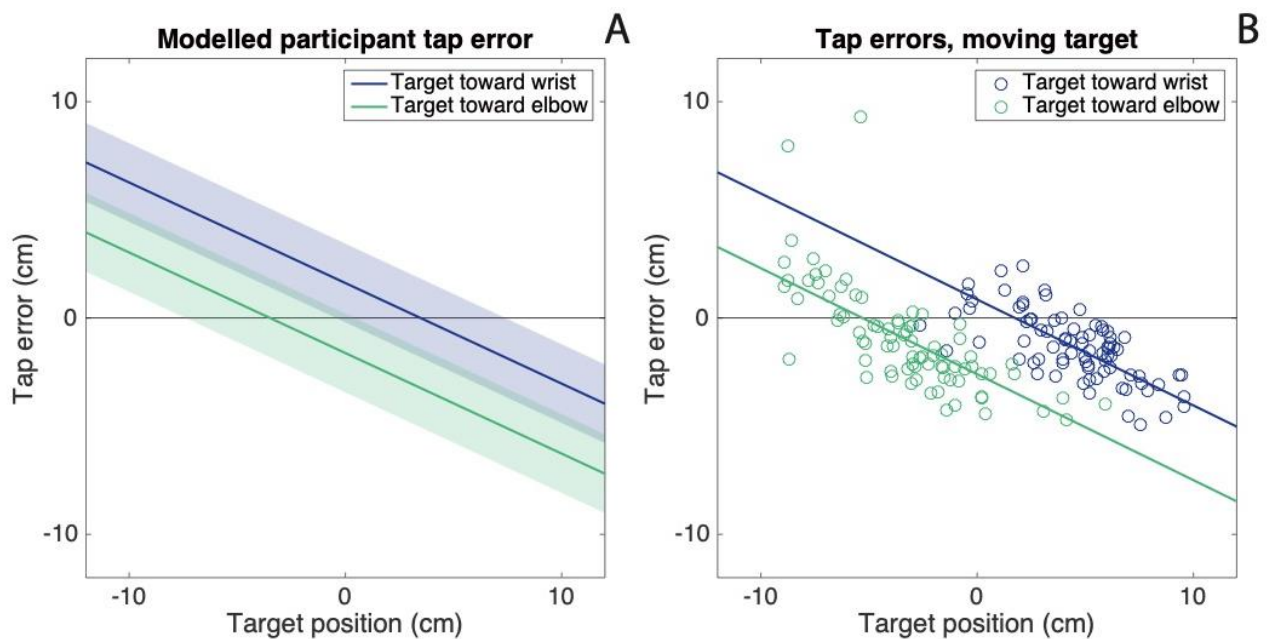


Figure 5.6. **A.** Interception performance as predicted by the model for the medium speed targets based on the median performance in the Velocity and Position experiments. Color indicates the direction of target motion. Shaded areas indicate one standard deviation from the mean. **B.** A single participant’s performance in all trials of the Interception experiment (circles) and the best parallel linear fit (lines).

If we look at the three parameters for which we made predictions, we see that the predictions are very similar to the values observed in the Interception experiment (**Figure 5.7**), despite being based on a different set of participants and the parameters differing considerably across participants. Our model predicts a median slope of -0.47 for the bias towards the center of the forearm, which is considerably steeper than what was observed for static targets. This predicted slope is very similar to what we observed in the Interception experiment (**Figure 5.7A**). The model also correctly predicts the effect of target direction on performance, estimating a median intercept shift of 3.1 cm, essentially equal to the observed shift of 3.0 cm (**Figure 5.7B**). One prediction is incorrect: our model predicted that participants would be less precise in the Interception experiment than in the Position experiment, which was not the case. The model slightly underestimates participants’ precision for Interception, predicting a precision of 1.9 cm, compared to 1.4 cm in the experiment (**Figure 5.7C**).

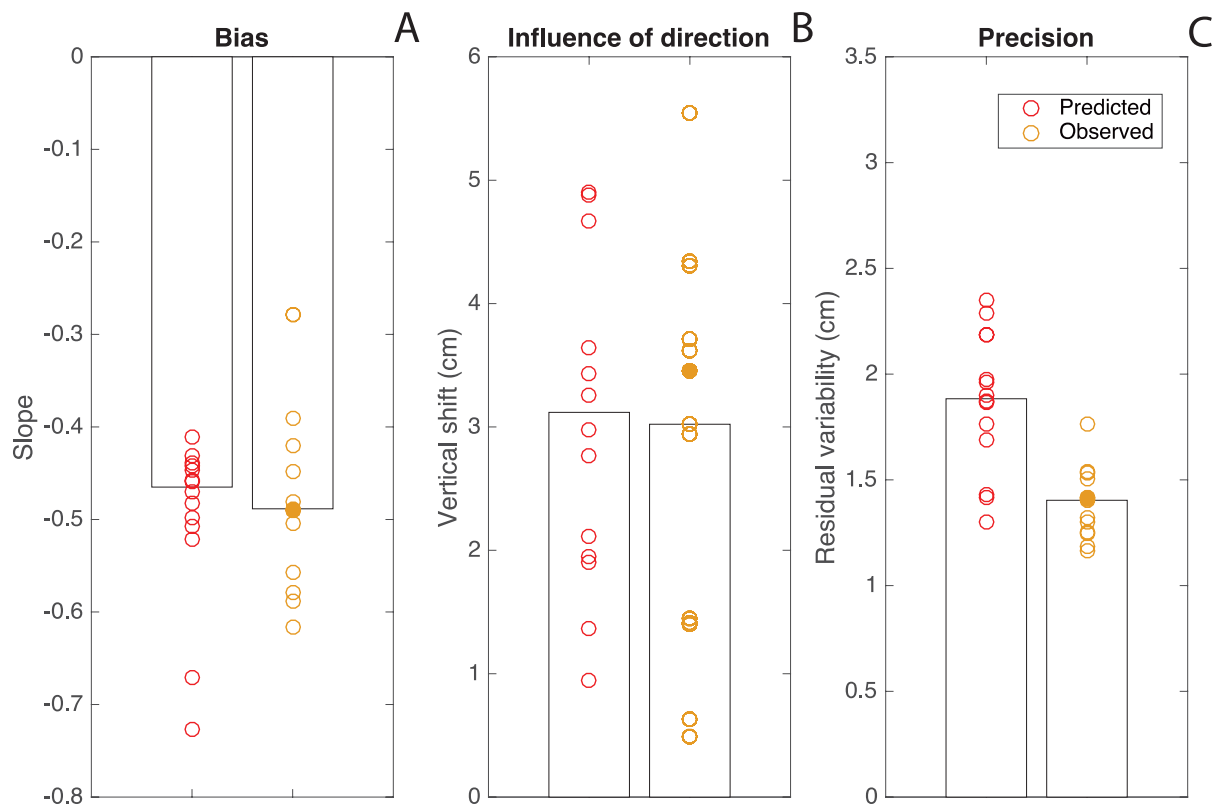


Figure 5.7. Comparison of model predictions of interception performance (red symbols) with the experimental results of the Interception experiment (orange symbols). The bias (A), influence of direction (B) and precision (C) as predicted by the model based on the results of each participant's performance in the Velocity and Position experiments and as determined for each participant in the Interception experiment. Bars indicate median values. The filled orange symbols indicate the values for the participant whose data are shown in Figure 5.6B.

5.4 Discussion

In this paper we demonstrated that interception of a tactile target can be modeled in a similar way to interception of visual targets. The Position experiment demonstrated that participants can locate a static target on their arm using the index finger of the other hand, but that participants exhibit a tendency to aim toward the center of the forearm. The Velocity experiment showed that participants are able to distinguish between different target speeds as perceived by touch, although the reproduction speeds were consistently far too high. The Interception experiment revealed that the biases that were found in the other two experiments influence interception, and showed that participants can intercept a target that is moving along the arm with remarkable precision.

The fact that our prediction for the influence of the direction of motion was so accurate was a surprise to us. When we observed that the target's velocity was overestimated by a factor four, we doubted whether our measure was adequate. In principle, it could be the case that the results did not indicate an overestimation of perceived tactile speed, but an underestimation of the produced hand speed. However, considering our results it is probably

the target's velocity that is misperceived. This might be because the produced hand speed is reasonably well calibrated since it is not unusual to want to move one's hand at a certain speed, whereas people do not often need to judge tactile speed.

In our model, we assumed that people have a prior for the target's position at the time of the hit that is combined with the extrapolated position based on the sensed position and velocity some time before the tap to determine where to tap. We therefore determined a weighted average of the anticipated target position and this prior. However, one could argue that the prior does not affect the extrapolated position but affects the sensed position. If so, we would have to replace equation (9) by:

$$T = w p + (1 - w)(\hat{x}_\tau - x_\tau) + (\hat{v} - v)\tau \quad (14)$$

This corresponds with determining a weighted average of the sensed position and the prior, as in the Position experiment, and extrapolating from the resulting position with the perceived velocity. This alternative model differs from our original model in that the weight given to the prior (and thus the dependence of the tapping error on position) is the same in the Interception experiment as in the Position experiment, which is not consistent with our data (different slopes in **Figures 5.2** and **5.6**). Equation (14) also predicts a larger influence of the direction of target motion, but this is less convincing evidence against this option because some of the overestimation of the speed in the Velocity experiment may be the result of underestimating the arm movement velocity after all.

One uncertainty in all our comparisons is that the value that we use for the sensorimotor delay might not be correct for tactile information. The fact that both our prediction for the change in the bias towards the center of the arm and our prediction for the systematic influence of the target's motion on participants' tap locations were correct when using this value supports the value of 100 ms for the delay. We do not have a direct independent estimate of how quickly people can react to tactile location information, but it takes people about 90 ms to adjust their grip strength in response to the sensation of finger slip (Brandes & Heed, 2015; Pruszynski, Johansson, & Flanagan, 2016) so a value of 100 ms is not unrealistic. If the delay is indeed about 100 ms it is unlikely that the correspondence between the bias in reproducing the target's speed and the systematic influence of the target's motion is a coincidence, so the velocity indicated with the hand in the Velocity experiment is probably the perceived velocity.

The results of the Interception experiment showed that participants' performance in a tactile interception task can be modeled on the basis of performance in a speed reproduction task and a static target task. Our model fits the biases in the data best when we consider participants' speed reproductions rather than the target speed. This confirms that participants take the perceived speed of the target into account rather than simply paying attention to an updating percept of target position. Interestingly, our model predicts poorer tap precision than actually exhibited by participants. In its current form, our model cannot produce precisions of interception that are better than those of the Position experiment, because the added variability introduced by the target moving can only cause the precision to decrease. A possible explanation for having found better precision for moving targets is that participants receive more feedback in the Interception experiment than in the Position experiment: they do not only know whether they missed, but also whether they tapped in front of or behind the target. Another possibility is that as tactile receptive fields on the forearm are large, the sequential motion over several adjacent receptive fields allows for

better localization of the moving target. The fact that participants showed considerably smaller movement times is also interesting. All participants were instructed in both the Position and Interception experiments to move as quickly and accurately as possible, but participants still moved faster in response to a moving target. It is possible that participants prioritized speed over accuracy when hitting a moving target so as to ensure that they arrive in time. This makes it even more surprising that participants' precision is better for moving targets.

One conspicuous difference between the results of the current study on intercepting tactile targets and those of the earlier studies on intercepting visual targets is that we find substantial biases for tactile targets whereas studies with visual targets found modest biases, mainly arising from illusions (de la Malla et al., 2018; de la Malla, Smeets, & Brenner, 2017) or from ignoring acceleration (Brenner & Smeets, 2015; Brenner et al., 2016). This difference is easy to understand. When moving a visible hand to a visible target one can simultaneously see both the hand and the target, so any visual bias will affect the localization of both hand and target (Smeets et al., 2006). Visual biases will therefore not affect interception. When moving to a tactile target without seeing one's hand, any tactile bias will affect the perceived location of the target, but not that of the hand. Tactile biases will therefore affect interception. There are many potential origins of biases when dealing with tactile targets, for instance because tactile localization depends on posture (Sadibolova, Tamè, & Longo, 2018).

Both the bias and the precision in the Position experiment were quite similar to those reported for manually indicating short touches with a von Frey filament (Brooks, Seizova-Cajic & Taylor, 2019). The precision values are a combination of perceptual and motor errors. Previous work examining how well people can detect the direction in which a stimulus is displaced along the back of their forearm suggests that the spatial precision of touch corresponds with a standard deviation of about 0.5 cm under those circumstances (Cody et al., 2008). Our participants' precision of about 1.5 cm when tapping a static target can therefore probably largely be attributed to motor errors. This would be consistent with the fact that we did not observe a systematic difference in tap precision along the arm.

Considering the persistent bias in the direction of motion that we find in the Interception experiment it is evident that the feedback that we provided is not enough to quickly learn to avoid biases such as the bias towards the center of the arm or the bias that is introduced by misjudging the velocity. Similar lack of adaptation was found for temporal errors ('spatial feedback' condition in de la Malla et al., 2014).

For the model in the present paper, we made many simplifying assumptions in comparison to the model that we previously proposed to account for the errors when hitting visual targets (Eli Brenner & Smeets, 2015). For instance, we considered motor variability and uncertainty in the remaining duration to have negligible influences on tactile interception. Considering how much poorer position and velocity judgments are for tactile targets (standard deviation of about 15 mm and Weber fraction of about 25%; **Figures 5.3B** and **5.5B**) than for visual targets (standard deviation of about 3mm and Weber fraction of less than 9%; Brenner & Smeets, 2015), these simplifying assumptions seem to be justified. There are of course many more differences between the visual and tactile modality. Despite these differences, stimulus position and motion can be perceived by both the skin and the eyes, and our experiment shows that interception relies similarly on such percepts for visual and tactile targets. Thus, the current findings provide support for the previously proposed mechanism of

guiding interception (Brenner & Smeets, 2015; Brenner & Smeets, 2018), extending its validity to include tactile targets.

5.5 Data

All data pertaining to the study are freely available for review. Interested parties are invited to download the data via the Open Science Framework at https://osf.io/65nuq/?view_only=2abc02237980454d8272d734c4a0fb6a

5.6 Acknowledgements

The authors would like to thank Qinqi Xu for her help with recruiting participants, as well as Fabio Tatti for his help with setting up the experimental system.

6

Epilogue

This thesis aimed to shed light on the contributions of vision, audition, and haptics to our sense of, and interaction with the world around us. Through four varied studies, we discovered that spatial and temporal perception are robust under unusual conditions and function even in settings we are not typically exposed to. In **Chapter 2**, we determined that spatial perception in the peripersonal space functions just as well for people who have been blind their whole lives as it does for sighted people who have only just recently been blindfolded. We expected that drawn lines that started away from the body midline would deviate from fronto-parallelity, and that this deviation would be exaggerated in blind participants if they relied more heavily on egocentric reference frames than sighted participants. We also expected that such a reliance on egocentric reference frames would also cause early-blind participants to draw longer lines when beginning said lines farther away from the body. As these effects were not present in our study, we concluded that spatial perception does not depend on visual ability for our tasks, and that any differences in reliance on egocentric and allocentric reference frames did not translate to difference in performance. **Chapter 3** explored people's ability to determine which of two brief, nearly simultaneous events occurred first. We discovered that the act of tapping greatly improved participants' ability to determine the relative timing of a flash or beep, and we reproduced the well-established finding that auditory perception is more temporally precise than visual perception. This work was extended in **Chapter 4**, which explored people's ability to tap their finger along with a guide rhythm while receiving perturbed feedback about the timing of their own finger taps. We found that people are largely unaffected by perturbed feedback when the guide and feedback share the same modality. However, when the guide and feedback differed in modality—and particularly when the guide was visual and the feedback was auditory—participants' performance was dragged along with the perturbation of the feedback. This suggests that participants are more sensitive to perturbed feedback when dealing with guide and feedback stimuli of differing modalities, which would explain their hindered performance in the cross-sensory conditions. Finally, **Chapter 5** considered people's ability to determine the position and velocity of a target felt on the arm, and ended with a task in which participants had to intercept this target. We found that people could determine the position of the target on the arm with reasonable precision, and that they were also able to distinguish between various target speeds. We were able to devise an empirical model taking into account people's biases and precision with regard to their perception of the target, which correctly predicted performance in the interception task. Our results showed not only that interception is possible when the input stimuli are tactile rather than visual, but also that tactile interception is similar to visual interception in terms of how available information is used. This suggests that the ability to intercept a moving target, while primarily a visual task in humans, is not inherently tied to or dependent on vision.

6.1. Space

Our results in **Chapter 2** showed no difference in performance between blind and blindfolded sighted adults, suggesting that our participants did not depend on prior visual experience for the task at hand. Past work has shown that early-blind individuals tend to rely more on egocentric reference frames than sighted individuals for certain tasks (Iachini et al., 2014; Pasqualotto et al., 2013; Postma et al., 2008), so our finding of no such difference in

our experiment was particularly striking. The relationship between spatial perception and visual experience is complicated, but our findings are perhaps unsurprising in the light of previous work. For example, a previous study had early-blind and blindfolded sighted individuals place an object on a table and then walk away from the table (Hollins & Kelley, 1988). When asked to point to the location of the object after having walked around to another part of the table, early-blind participants showed significantly worse performance at this task than blindfolded sighted participants. This would make sense, as pointing is typically a visually guided task. However, when asked to walk back to the table and place the object in its original location, both groups were equally capable of completing this task. This confirms our conclusion that visual experience is not necessary to interact with the environment directly within our reach.

Participants' consistent tendency to undershoot the length of the felt object is particularly interesting, given other work that shows people are capable of distance reproduction under the right circumstances. Numerous studies with both blind and blindfolded sighted participants have shown that participants are able to reproduce a passively felt traveled distance with the hand (Hermelin & O'Connor, 1975; Lederman et al., 1985), including when the motion is simulated and the hand actually does not move at all (Moscatelli et al., 2014). It has been demonstrated that people rely on a sense of vibration in order to determine how far their finger has traveled across a surface or object (Dallmann et al., 2015; Fardo et al., 2018). We deliberately deprived our participants of this crucial information in an attempt to tease out differences between early-blind and blindfolded sighted participants; the lack of any such difference in our data is interesting and shows that visual experience plays no role here.

6.2. Time

Our results showed that vision, audition, and haptics each play a unique role in temporal perception. To our knowledge, **Chapters 3** and **4** document some of the first studies performed wherein participants receive anticipatory feedback on their own motor actions without restricting the position of the finger by means of an external device.

There is evidence for the hypothesis that one's own motor actions play an important role in temporal perception. For example, in a tactile stimulation task with temporal adaptation, participants were only influenced by the adaptation when their arm was passively moved. When they performed a similar motor action on their own, the temporal adaptation had no effect on participant performance (Tomassini et al., 2012). However, another study has shown that participants are susceptible to temporal adaptation in a visuomotor temporal order judgment task even when performing motor actions, though this study did not include a condition with passive haptic stimuli for comparison (Vercillo et al., 2017). This contradiction of findings demonstrates that the exact nature of how motor actions affect our sense of time is still not fully understood. A third study demonstrated that the effect of movement on temporal perception is not limited to actions that relate to the perceived stimulus; freely moving while listening to an external tone has been shown to improve participants' ability to determine the duration of the tone (Wiener et al., 2019). Our initial goal in **Chapter 3** was to establish temporal precision values for the visual, auditory, and haptic domains to better understand the work presented in **Chapter 4**, and for this reason we did not include a passive tactile condition. While our experiment did not capture the precision with which participants

could distinguish the temporal order of a beep-flash stimulus pair, we did successfully show that participants were far more precise when asked to compare the time of their tap to that of an external stimulus, adding to the evidence that a self-generated motor action improves temporal precision. An interesting future experiment building off of the work in **Chapter 3** could consider a “beep-flash-tap” condition, wherein participants tap their finger but must make a temporal order judgment between the beep and the flash. Improved performance in this hypothetical task compared to that in the beep-flash condition would lend additional weight to the argument that motor actions enhance time perception.

Chapter 4 presents an interesting difference in rhythmic tapping performance between conditions wherein the guide and feedback share a sensory modality and conditions wherein they differ. Specifically, both tap time precision and synchrony were greater in the unimodal conditions than they were in the cross-modal conditions. Furthermore, in our VA condition in both experiments, participants were clearly and consistently affected by the temporal perturbation of the feedback stimuli. Our findings agree with those of previous works which explored the difference in asynchrony between unimodal and cross-modal tapping conditions (for example, see Castro-Meneses & Sowman, 2018; Müller et al., 2008). Taken together, it seems likely that the negative mean asynchrony is indeed worsened by the presence of multiple sensory modalities competing for attention. However, this does not explain why participants appear unaffected by feedback perturbation in the unimodal conditions, or why the VA condition should elicit such a different response from the AV condition. Previous work has shown that participants can better synchronize to an auditory guide rhythm than to a visual one (for reviews, see Repp, 2005; Repp & Su, 2013). Interestingly, we observed no difference in performance between the VV and AA conditions in Experiment 1, and indeed a slightly smaller asynchrony in the VV condition in Experiment 2, contrary to much of the literature. It is possible that people naturally weight auditory information as more important than visual information with regard to temporal perception, and that our participants inadvertently confused the guide and feedback to some extent in the VA condition. This alone would not explain the results in the AV condition, though the effect there is in any case much smaller and may simply be a product of attending to both visual and auditory external stimuli simultaneously.

6.3. Interception

This thesis concludes with a demonstration, laid out in **Chapter 5**, that interception can be executed using only tactile information. Despite the presence of biases and worse precision than generally shown for visual targets, participants were able to determine the position and velocity of the target on the arm, and were able to intercept the target as well. Perhaps our most surprising finding was the result of the Velocity experiment. While nearly all participants were able to distinguish between the three speeds of the target object, they also reproduced speeds well over twice as fast as that of the target object. The relatively poor precision with which our participants could distinguish the different speeds (median Weber fraction of 0.25) is similar to what was found in previous work (Essick et al., 1988), though we observed much greater inter-subject variability with regard to precision in our study. While much research has been done regarding velocity discrimination on the skin (for example, Johnson et al., 1980; Merz et al., 2019; Nguyen et al., 2015), we were unable to find any studies wherein participants had to reproduce a felt speed. As we discussed earlier, the biases

found in **Chapter 2** may be due to a lack of vibrational input. Here we also deliberately deprived participants of vibrational input when reproducing a felt speed. We did this to emulate the kinds of sensory input that participants would receive during their tapping motion in the Position and Interception experiments (where the only tactile feedback occurs at the very end of the motion), and to ensure that participants were not accidentally reproducing felt distance instead of felt speed. In this way, all three experiments required participants to perform a motor action with proprioceptive, but no tactile feedback about their performance during the motor action itself.

Another interesting finding we encountered was the presence of a spatial bias. In both the Position and the Interception experiments, participants showed a slight tendency to aim toward the center of the arm. This bias was roughly twice as large in the Interception experiment as it was in the Position experiment (though individual biases varied quite a lot in the Position experiment), indicating a larger weighting of this bias when trying to hit a moving target. A recent study found that participants underestimated the endpoint of a tactile stimulus traveling along the forearm (Macauda et al., 2018); while we did not observe an effect of stimulus direction as they did, it is possible that the discrepancy in bias size between the Position and Interception experiments could at least in part be due to the presence of a similar phenomenon in the Interception experiment. Length contraction has been explored extensively in the tactile domain, though very often this refers to a contraction of the perceived distance between two static tactile stimuli (Goldreich & Tong, 2013; Tong et al., 2016; Trojan et al., 2014). Some previous work has shown that localization of static stimuli should be better near the wrist or elbow, as opposed to the middle of the forearm (Cholewiak & Collins, 2003). We did not observe a difference in localization performance across the forearm, and we believe that the bias we observed instead corresponds to a perceived “best guess” at where an object on the forearm is most likely to be.

6.4. Methodological issues

A recurring issue in the studies presented herein is a lack of statistical power. The data collection sessions for the studies presented in **Chapters 2** and **5** were each conducted in the span of a few days, due to scheduling constraints. In the case of **Chapter 2**, we recruited every willing participant who visited our booth at the 2017 ZieZo Beurs (a two-day event) in Utrecht, the Netherlands, resulting in over a hundred usable datasets for analysis. Unfortunately, this still resulted in very few fully blind participants (seven early-blind and nine late-blind), compared to fifty-five normally sighted participants. While we acknowledge the possibility that our small and uneven sample sizes may have obscured a small but real effect, we were expecting a relatively large effect size which should have shown itself even in our limited data. Data collection for the study in **Chapter 5** was similarly constrained to two one-week periods in late 2016 and early 2017. Due to participant availability, it was difficult to recruit the same participants for the Interception experiment as we had in the Position and Velocity experiments. Ideally, we would have liked to test individual participant performance across the three experiments as a further validation of our model, but we ultimately opted not to do this in light of the small number of overlapping participants. The fact that participants were nonetheless reasonably consistent was encouraging to us.

The experiment in **Chapter 2** had participants match a pen in one hand to the position of their other hand at three target locations, referred to as the ‘near,’ ‘far,’ and ‘side’ targets.

A measurement error resulted in the side target for left-handed participants being placed 6.6 cm farther away from the body midline than intended. Participants did not report any issues reaching this target. We attempted to correct for this error in analysis using a scaling factor to normalize the values of the line angles for left- and right-handed participants. While we do not suspect that this ultimately affected our results, we ideally would have placed the target at the correct location when setting up the experimental apparatus. This experiment also coarsely grouped participants into categories of visual impairment without controlling for cause of blindness or the nature of the gradual decline of many of the participants' visual capacity. While this may not matter as much for the early-blind group who had effectively no visual experience, one could argue that the participants in the late-blind, low vision, and high vision groups should have been more finely categorized. Finally, given that we had participants grasp the reference object between their thumb and pinkie finger, it may have been useful for us to measure participants' hand size in case it had an influence on the drawn line lengths.

The tap predictor in **Chapters 3** and **4** worked reasonably well, but it was not without flaws. In particular, while all our participants reported no neurological abnormalities, we discovered that some people's hands tended to shake more than others, resulting in false positives when these participants held their index finger in the air. In these cases, the predictor detected the downward motion of the shake as the beginning of a finger tap, resulting in the production of beeps or flashes at times when the participant had not intended to tap. This occurred regularly for some participants, resulting in their exclusion from analysis. Participants also varied in the way that they tapped, despite our consistent instructions to make large, fast movements with the finger. Even after such instructions, some participants moved their finger quite slowly or barely lifted their finger from the table. After discovering during analysis that the predictor could not cope with these kinds of motions, we opted to exclude these participants as well.

The narrow time window in which the predictor could estimate tap times also meant that we had to constrain the time range in our experiment in **Chapter 3**. We suspect that considering a larger time range for the BF condition would provide us with more accurate information concerning people's temporal precision in this task, but we chose to keep the temporal range consistent between the three conditions. Unfortunately, we cannot reliably deliver stimuli more than 50 ms before the actual tap time with our predictor, so we chose to constrain the range for all three conditions to -50 to 50 ms. Researchers have previously predicted the timing of finger taps using a force-feedback device strapped to the finger (Rohde & Ernst, 2013). While the researchers were able to generate stimuli as early as 200 ms before button press, their predictor's precision was also found to be relatively low. Their method would be very useful in exploring the difference between active and passive motor actions by including a condition wherein the device moves the finger, in addition to the already present condition wherein the participant taps normally. Previous work has also shown that it is possible to simulate anticipatory feedback via adaptation (Sugano et al., 2010). The researchers' participants adapted to feedback lags of up to 150 ms for both visual and auditory stimuli, such that any smaller lag was perceived as feedback occurring before their own finger tap. A combination of these researchers' adaptation procedure and a tap predictor could allow us to deliver stimuli that are perceived to occur even earlier, which in particular would be useful for expanding the temporal range in temporal order judgment experiments involving self-generated motor actions.

We observed a difference in the precision of our tap predictor between **Chapters 3** and **4**, which may shed some light on ways to improve the predictor itself. In **Chapter 4**, participants were instructed to tap rhythmically for roughly ten minutes at a time, resulting in fast, consistent finger taps. In **Chapter 3**, participants performed one tap at a time and had to pause to consider the order of events that had just occurred. Their taps tended to be slower and more hesitant in this experiment. It may be possible to get better performance out of our tap predictor by having participants perform multiple rhythmic taps before being presented with an external stimulus, whose timing they would then have to compare with the most recent tap. Further research in this direction would generally benefit from developing a more robust tap prediction method capable of predicting further in advance, or otherwise circumventing this with the inclusion of an adaptation phase.

In **Chapter 4**, we arbitrarily chose perturbation periods of 20 and 40 taps for Experiments 1 and 2, respectively. It would have been better to keep this value consistent across the two experiments, as well as possibly testing with multiple periods in order to properly measure lag effects. With this said, we do not expect that our choice of perturbation period had an effect on our results.

This thesis has demonstrated that vision, audition, and haptics interact intricately to inform us accurately about the state of the world around us. We have demonstrated that spatial perception of the world immediately around us develops and functions in the absence of vision, even in ecologically unlikely scenarios. Furthermore, we have confirmed that temporal perception, and the precision thereof, function best in the presence of motor actions coupled with auditory stimuli. We found that people can most accurately tap along with a guide rhythm when the guide and feedback share a sensory modality, and that perturbed feedback can alter people's tap times if the feedback and guide differ in modality. Finally, we show that interception, which is dependent on both spatial and temporal perception, can be reliably performed with only tactile input, and that we execute this action similarly to a visual interception task.

7

Summary

Our perception of the world around us depends on input from many sense organs working together. In particular, our senses of vision, hearing, and touch are crucial for sensing the present state of our environment. Each sense is particularly suited for certain kinds of information, which we compile in order to perform simple and complex motor tasks. For example, normally sighted individuals rely in large part on the visual system to determine the locations of, and interact with objects in our environment, but we are also able to reach for and manipulate objects in the dark by relying on haptic inputs. Our sense of timing of brief events, which for example is necessary for actions such as singing, dancing, and playing musical instruments, is largely informed by auditory and haptic input, but in cases such as orchestral performances the players must synchronize with the visual signals of the conductor. Interception, a common action in sports, is commonly thought of as a visually guided task, though many animals use olfaction, echolocation, or electroreception to hunt. It is therefore not immediately clear if, or why, interception should be only possible in the visual domain for humans. These questions led us to wonder how each of our senses informs our ability to perform certain tasks, and how generalizable these tasks are to other, less common sensory modalities.

We conducted four studies to understand the influence of our senses on space and time, as well as how we combine spatial and temporal information in an interception task. The study conducted in **Chapter 2** concerns our perception and ability to act in the space within reach of our own bodies. We hypothesized that blind individuals would perform differently in a line-drawing task from blindfolded sighted individuals, due to a lack of development of visual experience. We set up a booth at the ZieZo Beurs in Utrecht and collected data from 132 participants, all of whom were attendees of the convention. These participants, whose visual capacity ranged from being fully blind to being normally sighted, were instructed to perform a series of tasks, wherein they matched the location of a pen in their hand to that of their other hand under the table, and then drew lines of a specific length in a specific direction. We were surprised to find no differences in performance, with regard to positioning of the pen, or the length or angle of the drawn lines, between our groups. We conclude that the development of visual experience has no effect on people's sense of the space around their own bodies.

Chapter 3 explores the precision with which we can determine the timing of brief events, across differing sensory modalities. It is well known that we have better temporal precision for auditory and haptic stimuli than for visual ones, but the extent to which motor actions play a role in our perception is not yet fully understood. We developed a novel method to track the position of participants' index finger and predict when they would perform a finger tap, before said tap actually occurred. Using this method, we set up a temporal order judgment task, wherein participants had to determine which of two brief stimuli occurred first. The experiment consisted of three conditions, each of which explored a different stimulus pair: a flash-tap condition, a beep-tap condition, and a beep-flash condition. We found that participants were most precise in the beep-tap condition, followed by the flash-tap condition. Participants' precision in the beep-flash condition was much poorer, to such an extent that our experiment could not accurately capture it. Our findings contribute to growing evidence that self-generated motor actions benefit precise temporal perception.

In **Chapter 4** we discuss a rhythmic tapping task, meant to explore the effect of perturbed feedback on participants' ability to synchronize with a simple guide rhythm across different sensory modalities. Most rhythmic tapping studies involve an auditory guide rhythm, and provide participants no feedback about their own performance (except for the haptic information given by tapping). Our study considered the effect of both visual and auditory guide stimuli. Each time the participant tapped their finger, they also caused a visual or auditory stimulus to occur. We used the same tap predictor as in the previous chapter to modify the exact timing of this feedback stimulus, such that it could occur too early or too late, relative to the finger tap. We found that in conditions where the modality of the guide and feedback differed, participants showed a greater bias and worse precision in trying to follow the guide than in conditions where the two modalities were the same. Particularly in the visual-auditory condition, the timing of participants' taps relative to the guide was dragged along with the offset of the feedback. Taken together, this indicates that participants rely more on the feedback in the cross-modal conditions than in unimodal conditions, and that they synchronize best with the guide when external stimuli are limited to a single sensory modality.

The thesis concludes with a task which requires a sense of both space and time to execute successfully. In **Chapter 5**, we carried out a study on interception in the tactile domain. Interception in humans is usually a visual endeavor, as our visual system is best suited to track fast and distant objects. It has previously been shown that people's ability to hit or catch a moving object can be predicted from how precisely they can determine both the object's position at any given time, as well as its speed. In our study, we sought to determine whether interception could also be executed on the basis of tactile information, and if so, whether people's performance similarly depended on a combination of position and velocity information. In each of our experiments we blindfolded participants and used a robotic device to deliver stationary or moving tactile stimuli to the participants' left arm. The study consisted of three experiments: one to see how precisely participants could locate a stationary target they felt on their arm, one to test their ability to move their hand at the same speed as a moving target, and finally, one wherein they had to hit the target as it moved along their arm. Our results showed that participants could locate the static target on the arm, albeit with relatively poor precision; that they were able to distinguish between the various velocities of the moving object; and that performance in the Interception experiment could be predicted based on performance in the Position and Velocity experiments, similarly to what is found in visual interception experiments.

Our studies have shed new light on how we perceive the state of the world around us. We have shown in **Chapters 2** and **5** that spatial perception and interception, while often informed via the visual system, are at the same time not always dependent on it. We found in **Chapter 3** that people are far more precise at comparing the timing of a self-generated motor action and that of an external stimulus than they are at comparing the timing of two external stimuli. In **Chapter 4** we showed that participants rely more on feedback during cross-sensory rhythmic tapping tasks than in unisensory ones, and that people's ability to maintain rhythm is hindered when they are asked to pay attention to external stimuli of more than one sensory modality.

8

References

- Alais, D., & Burr, D. (2004). The Ventriloquist Effect Results from Near-Optimal Bimodal Integration. *Current Biology*, *14*(3), 257–262.
<https://doi.org/10.1016/j.cub.2004.01.029>
- Angelaki, D. E., Gu, Y., & DeAngelis, G. C. (2009). Multisensory integration: Psychophysics, neurophysiology, and computation. *Current Opinion in Neurobiology*, *19*(4), 452–458. <https://doi.org/10.1016/j.conb.2009.06.008>
- Aschersleben, G., Gehrke, J., & Prinz, W. (2001). Tapping with peripheral nerve block. *Experimental Brain Research*, *136*(3), 331–339.
<https://doi.org/10.1007/s002210000562>
- Bartlett, N. R., & Bartlett, S. C. (1959). Synchronization of a motor response with an anticipated sensory event. *Psychological Review*, *66*(4), 203–218.
<https://doi.org/10.1037/h0046490>
- Blais, M., Albaret, J.-M., & Tallet, J. (2015). Is there a link between sensorimotor coordination and inter-manual coordination? Differential effects of auditory and/or visual rhythmic stimulations. *Experimental Brain Research*, *233*(11), 3261–3269.
<https://doi.org/10.1007/s00221-015-4394-6>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.
<https://doi.org/10.1163/156856897X00357>
- Brandes, J., & Heed, T. (2015). Reach Trajectories Characterize Tactile Localization for Sensorimotor Decision Making. *The Journal of Neuroscience*, *35*(40), 13648–13658.
<https://doi.org/10.1523/JNEUROSCI.1873-14.2015>
- Brenner, E., Rodriguez, I. A., Muñoz, V. E., Schootemeijer, S., Mahieu, Y., Veerkamp, K., Zandbergen, M., van der Zee, T., & Smeets, J. B. J. (2016). How Can People Be so Good at Intercepting Accelerating Objects if They Are so Poor at Visually Judging Acceleration? *I-Perception*, *7*(1), 1–13. <https://doi.org/10.1177/2041669515624317>
- Brenner, E., Smeets, J. B. J., & de Lussanet, M. H. E. (1998). Hitting moving targets: Continuous control of the acceleration of the hand on the basis of the target's velocity. *Experimental Brain Research*, *122*(4), 467–474.
<https://doi.org/10.1007/s002210050535>
- Brenner, Eli, & Smeets, J. B. J. (1997). Fast Responses of the Human Hand to Changes in Target Position. *Journal of Motor Behavior*, *29*(4), 297–310.
<https://doi.org/10.1080/00222899709600017>
- Brenner, Eli, & Smeets, J. B. J. (2010). How well can people judge when something happened? *Vision Research*, *50*(12), 1101–1108.
<https://doi.org/10.1016/j.visres.2010.03.004>
- Brenner, Eli, & Smeets, J. B. J. (2015). How people achieve their amazing temporal precision in interception. *Journal of Vision* (2015), *15*(3):8. <https://doi.org/10.1167/15.3.8>
- Brenner, Eli, & Smeets, J. B. J. (2018). Continuously updating one's predictions underlies successful interception. *Journal of Neurophysiology*, *120*(6), 3257–3274.
<https://doi.org/10.1152/jn.00517.2018>

- Brenner, Eli, van Dam, M., Berkhout, S., & Smeets, J. B. J. (2012). Timing the moment of impact in fast human movements. *Acta Psychologica*, *141*(1), 104–111. <https://doi.org/10.1016/j.actpsy.2012.07.002>
- Brooks, J., Seizova-Cajic, T., & Taylor, J. L. (2019). Biases in tactile localization by pointing: Compression for weak stimuli and centering for distributions of stimuli. *Journal of Neurophysiology*, *121*(3), 764–772. <https://doi.org/10.1152/jn.00189.2018>
- Brown, M. S., & Stratton, G. M. (1925). The spatial threshold of touch in blind and in seeing children. *Journal of Experimental Psychology*, *8*(6), 434–442. <https://doi.org/10.1037/h0063926>
- Burgess, N. (2006). Spatial memory: How egocentric and allocentric combine. *Trends in Cognitive Sciences*, *10*(12), 551–557. <https://doi.org/10.1016/j.tics.2006.10.005>
- Castro-Meneses, L. J., & Sowman, P. F. (2018). Stop signals delay synchrony more for finger tapping than vocalization: A dual modality study of rhythmic synchronization in the stop signal task. *PeerJ*, *6*, e5242. <https://doi.org/10.7717/peerj.5242>
- Cattaneo, Z., Fantino, M., Silvanto, J., Tinti, C., Pascual-Leone, A., & Vecchi, T. (2010). Symmetry perception in the blind. *Acta Psychologica*, *134*(3), 398–402. <https://doi.org/10.1016/j.actpsy.2010.04.002>
- Cholewiak, R. W., & Collins, A. A. (2003). Vibrotactile localization on the arm: Effects of place, space, and age. *Perception & Psychophysics*, *65*(7), 1058–1077. <https://doi.org/10.3758/BF03194834>
- Cody, F. W. J., Garside, R. A. D., Lloyd, D., & Poliakoff, E. (2008). Tactile spatial acuity varies with site and axis in the human upper limb. *Neuroscience Letters*, *433*(2), 103–108. <https://doi.org/10.1016/j.neulet.2007.12.054>
- Cohen, Y. E., & Andersen, R. A. (2002). A common reference frame for movement plans in the posterior parietal cortex. *Nature Reviews Neuroscience*, *3*(7), 553–562. <https://doi.org/10.1038/nrn873>
- Crevecoeur, F. (n.d.). *A sub-movement time scale of human motor adaptation*. 30.
- Cuturi, L. F., Aggius-Vella, E., Campus, C., Parmiggiani, A., & Gori, M. (2016). From science to technology: Orientation and mobility in blind children and adults. *Neuroscience & Biobehavioral Reviews*, *71*, 240–251. <https://doi.org/10.1016/j.neubiorev.2016.08.019>
- Dallmann, C. J., Ernst, M. O., & Moscatelli, A. (2015). The role of vibration in tactile speed perception. *Journal of Neurophysiology*, *114*(6), 3131–3139. <https://doi.org/10.1152/jn.00621.2015>
- de la Malla, C., & López-Moliner, J. (2015). Hitting moving targets with a continuously changing temporal window. *Experimental Brain Research*, *233*(9), 2507–2515. <https://doi.org/10.1007/s00221-015-4321-x>
- de la Malla, C., López-Moliner, J., & Brenner, E. (2012). Seeing the last part of a hitting movement is enough to adapt to a temporal delay. *Journal of Vision* (2012), *12*(10):4. <https://doi.org/10.1167/12.10.4>

- de la Malla, C., Smeets, J. B. J., & Brenner, E. (2017). Potential Systematic Interception Errors are Avoided When Tracking the Target with One's Eyes. *Scientific Reports*, 7(1). <https://doi.org/10.1038/s41598-017-11200-5>
- de la Malla, C., Smeets, J. B. J., & Brenner, E. (2018). Errors in interception can be predicted from errors in perception. *Cortex*, 98, 49–59. <https://doi.org/10.1016/j.cortex.2017.03.006>
- Delignières, D., & Marmelat, V. (2014). Strong anticipation and long-range cross-correlation: Application of detrended cross-correlation analysis to human behavioral data. *Physica A: Statistical Mechanics and Its Applications*, 394, 47–60. <https://doi.org/10.1016/j.physa.2013.09.037>
- Donohue, S. E., Woldorff, M. G., & Mitroff, S. R. (2010). Video game players show more precise multisensory temporal processing abilities. *Attention, Perception, & Psychophysics*, 72(4), 1120–1129. <https://doi.org/10.3758/APP.72.4.1120>
- Doumas, M., Wing, A. M., & Wood, K. (2008). Interval timing and trajectory in unequal amplitude movements. *Experimental Brain Research*, 189(1), 49–60. <https://doi.org/10.1007/s00221-008-1397-6>
- Drewing, K. (2018). The extent of skin bending rather than action possibilities explains why holes feel larger with the tongue than with the finger. *Journal of Experimental Psychology: Human Perception and Performance*, 44(4), 535–550. <https://doi.org/10.1037/xhp0000453>
- Eagleman, D. M. (2008). Human time perception and its illusions. *Current Opinion in Neurobiology*, 18(2), 131–136. <https://doi.org/10.1016/j.conb.2008.06.002>
- Essick, G. K., Franzen, O., & Whitsel, B. L. (1988). Discrimination and Scaling of Velocity of Stimulus Motion across the Skin. *Somatosensory & Motor Research*, 6(1), 21–40. <https://doi.org/10.3109/08990228809144639>
- Fardo, F., Beck, B., Cheng, T., & Haggard, P. (2018). A mechanism for spatial perception on human skin. *Cognition*, 178, 236–243. <https://doi.org/10.1016/j.cognition.2018.05.024>
- Fitch, W. T., & Rosenfeld, A. J. (2007). Perception and Production of Syncopated Rhythms. *Music Perception: An Interdisciplinary Journal*, 25(1), 43–58. <https://doi.org/10.1525/mp.2007.25.1.43>
- Franklin, D. W., & Wolpert, D. M. (2008). Specificity of Reflex Adaptation for Task-Relevant Variability. *Journal of Neuroscience*, 28(52), 14165–14175. <https://doi.org/10.1523/JNEUROSCI.4406-08.2008>
- Frissen, I., Ziat, M., Campion, G., Hayward, V., & Guastavino, C. (2012). The effects of voluntary movements on auditory–haptic and haptic–haptic temporal order judgments. *Acta Psychologica*, 141(2), 140–148. <https://doi.org/10.1016/j.actpsy.2012.07.010>
- Fujii, S., Hirashima, M., Kudo, K., Ohtsuki, T., Nakamura, Y., & Oda, S. (2011). Synchronization Error of Drum Kit Playing with a Metronome at Different Tempi by Professional Drummers. *Music Perception: An Interdisciplinary Journal*, 28(5), 491–503. <https://doi.org/10.1525/mp.2011.28.5.491>

- Gentaz, E., & Hatwell, Y. (1996). Role of gravitational cues in the haptic perception of orientation. *Perception & Psychophysics*, *58*(8), 1278–1292. <https://doi.org/10.3758/BF03207559>
- Gentaz, E., & Hatwell, Y. (1998). The haptic oblique effect in the perception of rod orientation by blind adults. *Perception & Psychophysics*, *60*(1), 157–167. <https://doi.org/10.3758/BF03211925>
- Goldreich, D., & Tong, J. (2013). Prediction, Postdiction, and Perceptual Length Contraction: A Bayesian Low-Speed Prior Captures the Cutaneous Rabbit and Related Illusions. *Frontiers in Psychology*, *4*, 221. <https://doi.org/10.3389/fpsyg.2013.00221>
- Gori, M., Sandini, G., Martinoli, C., & Burr, D. (2010). Poor Haptic Orientation Discrimination in Nonsighted Children May Reflect Disruption of Cross-Sensory Calibration. *Current Biology*, *20*(3), 223–225. <https://doi.org/10.1016/j.cub.2009.11.069>
- Gouras, P., & Zrenner, E. (1981). Color Vision: A Review from a Neurophysiological Perspective. In H. Autrum, E. R. Perl, R. F. Schmidt, & D. Ottoson (Eds.), *Progress in Sensory Physiology* (pp. 139–179). Springer. https://doi.org/10.1007/978-3-642-66744-2_4
- Green, B. G. (1982). The perception of distance and location for dual tactile pressures. *Perception & Psychophysics*, *31*(4), 315–323. <https://doi.org/10.3758/BF03202654>
- Grondin, S. (2010). Timing and time perception: A review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, & Psychophysics*, *72*(3), 561–582. <https://doi.org/10.3758/APP.72.3.561>
- Haggard, P., Newman, C., Blundell, J., & Andrew, H. (2000). The perceived position of the hand in space. *Perception & Psychophysics*, *62*(2), 363–377. <https://doi.org/10.3758/BF03205556>
- Harrar, V., & Harris, L. R. (2008). The effect of exposure to asynchronous audio, visual, and tactile stimulus combinations on the perception of simultaneity. *Experimental Brain Research*, *186*(4), 517–524. <https://doi.org/10.1007/s00221-007-1253-0>
- Heller, M. A., & Joyner, T. D. (1993). Mechanisms in the haptic horizontal-vertical illusion: Evidence from sighted and blind subjects. *Perception & Psychophysics*, *53*(4), 422–428. <https://doi.org/10.3758/BF03206785>
- Hermelin, B., & O'Connor, N. (1975). Location and Distance Estimates by Blind and Sighted Children. *Quarterly Journal of Experimental Psychology*, *27*(2), 295–301. <https://doi.org/10.1080/14640747508400488>
- Hollins, M., & Kelley, E. K. (1988). Spatial updating in blind and sighted people. *Perception & Psychophysics*, *43*(4), 380–388. <https://doi.org/10.3758/BF03208809>
- Hove, M. J., Iversen, J. R., Zhang, A., & Repp, B. H. (2013). Synchronization with competing visual and auditory rhythms: Bouncing ball meets metronome. *Psychological Research*, *77*(4), 388–398. <https://doi.org/10.1007/s00426-012-0441-0>
- Hudson, T. E., & Landy, M. S. (2012). Measuring adaptation with a sinusoidal perturbation function. *Journal of Neuroscience Methods*, *208*(1), 48–58. <https://doi.org/10.1016/j.jneumeth.2012.04.001>

- Iachini, T., Ruggiero, G., & Ruotolo, F. (2014). Does blindness affect egocentric and allocentric frames of reference in small and large scale spaces? *Behavioural Brain Research*, 273, 73–81. <https://doi.org/10.1016/j.bbr.2014.07.032>
- Jay, M. F., & Sparks, D. L. (1987). Sensorimotor integration in the primate superior colliculus. II. Coordinates of auditory signals. *Journal of Neurophysiology*, 57(1), 35–55. <https://doi.org/10.1152/jn.1987.57.1.35>
- Johnson, D., Jürgens, R., & Kornhuber, H. H. (1980). Somatosensory-evoked potentials and perception of skin velocity. *Archiv Für Psychiatrie Und Nervenkrankheiten*, 228(2), 95–100. <https://doi.org/10.1007/BF00365597>
- Johnson, M., Madsen, P. T., Zimmer, W. M. X., Aguilar de Soto, N., & Tyack, P. L. (2004). Beaked whales echolocate on prey. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(suppl_6). <https://doi.org/10.1098/rsbl.2004.0208>
- Jones, B. (1972). Development of cutaneous and kinesthetic localization by blind and sighted children. *Developmental Psychology*, 6(2), 349–352. <https://doi.org/10.1037/h0032120>
- Jones, M. B., & Vierck, C. (1973). Length Discrimination on the Skin. *The American Journal of Psychology*, 86(1), 49–60. JSTOR. <https://doi.org/10.2307/1421847>
- Kappers, A. M. L. (2002). Haptic perception of parallelity in the midsagittal plane. *Acta Psychologica*, 109(1), 25–40. [https://doi.org/10.1016/S0001-6918\(01\)00047-6](https://doi.org/10.1016/S0001-6918(01)00047-6)
- Kappers, A. M. L. (2003). Large systematic deviations in a bimanual parallelity task: Further analysis of contributing factors. *Acta Psychologica*, 114(2), 131–145. [https://doi.org/10.1016/S0001-6918\(03\)00063-5](https://doi.org/10.1016/S0001-6918(03)00063-5)
- Kappers, A. M. L. (2004). The contributions of egocentric and allocentric reference frames in haptic spatial tasks. *Acta Psychologica*, 117(3), 333–340. <https://doi.org/10.1016/j.actpsy.2004.08.002>
- Kappers, A. M. L., & Koenderink, J. J. (1999). Haptic Perception of Spatial Relations. *Perception*, 28(6), 781–795. <https://doi.org/10.1068/p2930>
- Keetels, M., & Vroomen, J. (2012a). Exposure to delayed visual feedback of the hand changes motor-sensory synchrony perception. *Experimental Brain Research*, 219(4), 431–440. <https://doi.org/10.1007/s00221-012-3081-0>
- Keetels, M., & Vroomen, J. (2012b). Exposure to delayed visual feedback of the hand changes motor-sensory synchrony perception. *Experimental Brain Research*, 219(4), 431–440. <https://doi.org/10.1007/s00221-012-3081-0>
- Kooij, K. van der, Brenner, E., Beers, R. J. van, & Smeets, J. B. J. (2015). Visuomotor Adaptation: How Forgetting Keeps Us Conservative. *PLOS ONE*, 10(2), e0117901. <https://doi.org/10.1371/journal.pone.0117901>
- Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10(7), 319–326. <https://doi.org/10.1016/j.tics.2006.05.003>

- Krause, V., Pollok, B., & Schnitzler, A. (2010). Perception in action: The impact of sensory information on sensorimotor synchronization in musicians and non-musicians. *Acta Psychologica*, *133*(1), 28–37. <https://doi.org/10.1016/j.actpsy.2009.08.003>
- Kuling, I. A., Brenner, E., & Smeets, J. B. J. (2016). Errors in visuo-haptic and haptic-haptic location matching are stable over long periods of time. *Acta Psychologica*, *166*, 31–36. <https://doi.org/10.1016/j.actpsy.2016.03.011>
- Kuling, I. A., van der Graaff, M. C. W., Brenner, E., & Smeets, J. B. J. (2017). Matching locations is not just matching sensory representations. *Experimental Brain Research*, *235*(2), 533–545. <https://doi.org/10.1007/s00221-016-4815-1>
- Kuling, I. A., van der Graaff, M. C. W., Brenner, E., & Smeets, J. B. J. (2014). Proprioceptive Biases in Different Experimental Designs. In M. Auvray & C. Duriez (Eds.), *Haptics: Neuroscience, Devices, Modeling, and Applications* (pp. 18–24). Springer Berlin Heidelberg.
- La Scaleia, B., Zago, M., Moscatelli, A., Lacquaniti, F., & Viviani, P. (2014). Implied Dynamics Biases the Visual Perception of Velocity. *PLoS ONE*, *9*(3), e93020. <https://doi.org/10.1371/journal.pone.0093020>
- Lackner, J. R., & DiZio, P. (2005). Vestibular, Proprioceptive, and Haptic Contributions to Spatial Orientation. *Annual Review of Psychology*, *56*(1), 115–147. <https://doi.org/10.1146/annurev.psych.55.090902.142023>
- Lange, J., Kapala, K., Krause, H., Baumgarten, T. J., & Schnitzler, A. (2018). Rapid temporal recalibration to visuo–tactile stimuli. *Experimental Brain Research*, *236*(2), 347–354. <https://doi.org/10.1007/s00221-017-5132-z>
- Lederman, S. J., Klatzky, R. L., & Barber, P. O. (1985). Spatial and movement-based heuristics for encoding pattern information through touch. *Journal of Experimental Psychology: General*, *114*(1), 33–49. <https://doi.org/10.1037/0096-3445.114.1.33>
- Lessard, N., Paré, M., Lepore, F., & Lassonde, M. (1998). Early-blind human subjects localize sound sources better than sighted subjects. *Nature*, *395*(6699), 278–280. <https://doi.org/10.1038/26228>
- Ley, I., Haggard, P., & Yarrow, K. (2009). Optimal integration of auditory and vibrotactile information for judgments of temporal order. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(4), 1005–1019. <https://doi.org/10.1037/a0015021>
- Macauda, G., Lenggenhager, B., Meier, R., Essick, G., & Brugger, P. (2018). Tactile motion lacks momentum. *Psychological Research*, *82*(5), 889–895. <https://doi.org/10.1007/s00426-017-0879-1>
- Maij, F., Brenner, E., & Smeets, J. B. J. (2009). Temporal Information Can Influence Spatial Localization. *Journal of Neurophysiology*, *102*(1), 490–495. <https://doi.org/10.1152/jn.91253.2008>
- Manning, F. C., & Schutz, M. (2015). Movement Enhances Perceived Timing in the Absence of Auditory Feedback. *Timing & Time Perception*, *3*(1–2), 3–12. <https://doi.org/10.1163/22134468-03002037>

- Mates, Jiří, Müller, U., Radil, T., & Pöppel, E. (1994). Temporal Integration in Sensorimotor Synchronization. *Journal of Cognitive Neuroscience*, *6*(4), 332–340. <https://doi.org/10.1162/jocn.1994.6.4.332>
- Mates, Jiří, & Aschersleben, G. (2000). Sensorimotor synchronization: The impact of temporally displaced auditory feedback. *Acta Psychologica*, *104*(1), 29–44. [https://doi.org/10.1016/S0001-6918\(99\)00052-9](https://doi.org/10.1016/S0001-6918(99)00052-9)
- Matthews, N., Welch, L., Achtman, R., Fenton, R., & FitzGerald, B. (2016). Simultaneity and Temporal Order Judgments Exhibit Distinct Reaction Times and Training Effects. *PLOS ONE*, *11*(1), e0145926. <https://doi.org/10.1371/journal.pone.0145926>
- Mayville, J. M., Bressler, S. L., Fuchs, A., & Kelso, J. A. S. (1999). Spatiotemporal reorganization of electrical activity in the human brain associated with a timing transition in rhythmic auditory-motor coordination. *Experimental Brain Research*, *127*, 11.
- Merz, S., Deller, J., Meyerhoff, H. S., Spence, C., & Frings, C. (2019). The contradictory influence of velocity: Representational momentum in the tactile modality. *Journal of Neurophysiology*, *121*(6), 2358–2363. <https://doi.org/10.1152/jn.00128.2019>
- Moscattelli, A., Naceri, A., & Ernst, M. O. (2014). Path integration in tactile perception of shapes. *Behavioural Brain Research*, *274*, 355–364. <https://doi.org/10.1016/j.bbr.2014.08.025>
- Müller, K., Aschersleben, G., Schmitz, F., Schnitzler, A., Freund, H.-J., & Prinz, W. (2008). Inter- versus intramodal integration in sensorimotor synchronization: A combined behavioral and magnetoencephalographic study. *Experimental Brain Research*, *185*(2), 309–318. <https://doi.org/10.1007/s00221-007-1155-1>
- Murai, Y., & Yotsumoto, Y. (2018). Optimal multisensory integration leads to optimal time estimation. *Scientific Reports*, *8*(1), 1–11. <https://doi.org/10.1038/s41598-018-31468-5>
- Nelson, J. S., Baud-Bovy, G., Smeets, J. B. J., & Brenner, E. (2019). Accuracy of Intercepting Moving Tactile Targets. *Perception*, *48*(8), 685–701. <https://doi.org/10.1177/0301006619861876>
- Nelson, Jacob S., Kuling, I. A., Gori, M., Postma, A., Brenner, E., & Smeets, J. B. J. (2018). Spatial Representation of the Workspace in Blind, Low Vision, and Sighted Human Participants. *I-Perception*, *9*(3), 204166951878187. <https://doi.org/10.1177/2041669518781877>
- Nelson, M. E., & MacIver, M. A. (1999). Prey capture in the weakly electric fish *Apteronotus albifrons*: Sensory acquisition strategies and electrosensory consequences. *Journal of Experimental Biology*, *202*(10), 1195.
- Nguyen, E. H. L., Taylor, J. L., Brooks, J., & Seizova-Cajic, T. (2015). Velocity of motion across the skin influences perception of tactile location. *Journal of Neurophysiology*, *115*(2), 674–684. <https://doi.org/10.1152/jn.00707.2015>
- Oostwoud Wijdenes, L., Brenner, E., & Smeets, J. B. J. (2011). Fast and fine-tuned corrections when the target of a hand movement is displaced. *Experimental Brain Research*, *214*(3), 453–462. <https://doi.org/10.1007/s00221-011-2843-4>

- Pasqualotto, A., Spiller, M. J., Jansari, A. S., & Proulx, M. J. (2013). Visual experience facilitates allocentric spatial representation. *Behavioural Brain Research*, *236*, 175–179. <https://doi.org/10.1016/j.bbr.2012.08.042>
- Pelli, D. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
- Pfordresher, P. Q., & Palmer, C. (2006). Effects of hearing the past, present, or future during music performance. *Perception & Psychophysics*, *68*(3), 362–376. <https://doi.org/10.3758/BF03193683>
- Postma, A., Zuidhoek, S., Noordzij, M. L., & Kappers, A. M. L. (2007). Differences between Early-Blind, Late-Blind, and Blindfolded-Sighted People in Haptic Spatial-Configuration Learning and Resulting Memory Traces. *Perception*, *36*(8), 1253–1265. <https://doi.org/10.1068/p5441>
- Postma, A., Zuidhoek, S., Noordzij, M. L., & Kappers, A. M. L. (2008). Haptic orientation perception benefits from visual experience: Evidence from early-blind, late-blind, and sighted people. *Perception & Psychophysics*, *70*(7), 1197–1206. <https://doi.org/10.3758/PP.70.7.1197>
- Pouget, A., Deneve, S., & Duhamel, J.-R. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nature Reviews Neuroscience*, *3*(9), 741–747. <https://doi.org/10.1038/nrn914>
- Poulton, E. C. (1979). Models for biases in judging sensory magnitude. *Psychological Bulletin*, *86*(4), 777–803. <https://doi.org/10.1037/0033-2909.86.4.777>
- Pruszynski, J. A., Johansson, R. S., & Flanagan, J. R. (2016). A Rapid Tactile-Motor Reflex Automatically Guides Reaching toward Handheld Objects. *Current Biology*, *26*(6), 788–792. <https://doi.org/10.1016/j.cub.2016.01.027>
- Repp, B. H. (2000). Compensation for subliminal timing perturbations in perceptual-motor synchronization. *Psychological Research Psychologische Forschung*, *63*(2), 106–128. <https://doi.org/10.1007/PL00008170>
- Repp, B. H. (2002). Phase correction in sensorimotor synchronization: Nonlinearities in voluntary and involuntary responses to perturbations. *Human Movement Science*, *21*(1), 1–37. [https://doi.org/10.1016/S0167-9457\(02\)00076-3](https://doi.org/10.1016/S0167-9457(02)00076-3)
- Repp, B. H. (2005). Sensorimotor synchronization: A review of the tapping literature. *Psychonomic Bulletin & Review*, *12*(6), 969–992. <https://doi.org/10.3758/BF03206433>
- Repp, B. H. (2008). Multiple temporal references in sensorimotor synchronization with metrical auditory sequences. *Psychological Research*, *72*(1), 79–98. <https://doi.org/10.1007/s00426-006-0067-1>
- Repp, B. H. (2010a). Sensorimotor synchronization and perception of timing: Effects of music training and task experience. *Human Movement Science*, *29*(2), 200–213. <https://doi.org/10.1016/j.humov.2009.08.002>
- Repp, B. H. (2010b). Sensorimotor synchronization and perception of timing: Effects of music training and task experience. *Human Movement Science*, *29*(2), 200–213. <https://doi.org/10.1016/j.humov.2009.08.002>

- Repp, B. H. (2011). Tapping in Synchrony With a Perturbed Metronome: The Phase Correction Response to Small and Large Phase Shifts as a Function of Tempo. *Journal of Motor Behavior*, 43(3), 213–227. <https://doi.org/10.1080/00222895.2011.561377>
- Repp, B. H., & Doggett, R. (2007). Tapping to a Very Slow Beat: A Comparison of Musicians and Nonmusicians. *Music Perception: An Interdisciplinary Journal*, 24(4), 367–376. <https://doi.org/10.1525/mp.2007.24.4.367>
- Repp, B. H., & Su, Y.-H. (2013). Sensorimotor synchronization: A review of recent research (2006–2012). *Psychonomic Bulletin & Review*, 20(3), 403–452. <https://doi.org/10.3758/s13423-012-0371-2>
- Röder, B., Kusmirek, A., Spence, C., & Schicke, T. (2007). Developmental vision determines the reference frame for the multisensory control of action. *Proceedings of the National Academy of Sciences*, 104(11), 4753–4758. <https://doi.org/10.1073/pnas.0607158104>
- Röder, B., Rösler, F., & Spence, C. (2004). Early Vision Impairs Tactile Perception in the Blind. *Current Biology*, 14(2), 121–124. <https://doi.org/10.1016/j.cub.2003.12.054>
- Röder, B., Teder-Sälejärvi, W., Sterr, A., Rösler, F., Hillyard, S. A., & Neville, H. J. (1999). Improved auditory spatial tuning in blind humans. *Nature*, 400, 162–166. <https://doi.org/10.1038/22106>
- Rohde, M., & Ernst, M. O. (2013). To Lead and To Lag – Forward and Backward Recalibration of Perceived Visuo-Motor Simultaneity. *Frontiers in Psychology*, 3, 599. <https://doi.org/10.3389/fpsyg.2012.00599>
- Rossi, F., Montanaro, E., & de’Sperati, C. (2018). Speed Biases With Real-Life Video Clips. *Frontiers in Integrative Neuroscience*, 12, 11. <https://doi.org/10.3389/fnint.2018.00011>
- Sadibolova, R., Tamè, L., & Longo, M. R. (20180830). More than skin-deep: Integration of skin-based and musculoskeletal reference frames in localization of touch. *Journal of Experimental Psychology: Human Perception and Performance*, 44(11), 1672. <https://doi.org/10.1037/xhp0000562>
- Schot, W. D., Brenner, E., & Smeets, J. B. J. (2010). Robust movement segmentation by combining multiple sources of information. *Journal of Neuroscience Methods*, 187(2), 147–155. <https://doi.org/10.1016/j.jneumeth.2010.01.004>
- Schütt, H. H., Harmeling, S., Macke, J. H., & Wichmann, F. A. (2016). Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research*, 122, 105–123. <https://doi.org/10.1016/j.visres.2016.02.002>
- Seashore, C. E., & Ling, T. L. (1918). The comparative sensitiveness of blind and seeing persons. *Psychological Monographs*, 25(2), 148–158. <https://doi.org/10.1037/h0093120>
- Senot, P., Zago, M., Le Seac’h, A., Zaoui, M., Berthoz, A., Lacquaniti, F., & McIntyre, J. (2012). When Up Is Down in Og: How Gravity Sensing Affects the Timing of Interceptive Actions. *Journal of Neuroscience*, 32(6), 1969–1973. <https://doi.org/10.1523/JNEUROSCI.3886-11.2012>

- Sensation. 2020. In *The American Heritage Dictionary of the English Language*. Retrieved January 13, 2020, from <https://www.ahdictionary.com/word/search.html?q=sensation>
- Shaffer, D. M., Dolgov, I., Mcmanama, E., Swank, C., Maynor, A. B., Kelly, K., & Neuhoff, J. G. (2013). Blind(fold)ed by science: A constant target-heading angle is used in visual and nonvisual pursuit. *Psychonomic Bulletin & Review*, *20*(5), 923–934. <https://doi.org/10.3758/s13423-013-0412-5>
- Smeets, J. B. J., & Brenner, E. (1995). Prediction of a moving target's position in fast goal-directed action. *Biological Cybernetics*, *73*(6), 519–528. <https://doi.org/10.1007/BF00199544>
- Smeets, J. B. J., Frens, M. A., & Brenner, E. (2002). Throwing darts: Timing is not the limiting factor. *Experimental Brain Research*, *144*(2), 268–274. <https://doi.org/10.1007/s00221-002-1072-2>
- Smeets, J. B. J., van den Dobbelen, J. J., de Grave, D. D. J., van Beers, R. J., & Brenner, E. (2006). Sensory integration does not lead to sensory calibration. *Proceedings of the National Academy of Sciences*, *103*(49), 18781–18786. <https://doi.org/10.1073/pnas.0607687103>
- Smeets, J. B. J., Wijdenes, L. O., & Brenner, E. (2016). Movement Adjustments Have Short Latencies because There Is No Need to Detect Anything. *Motor Control*, *20*(2), 137–148. <https://doi.org/10.1123/mc.2014-0064>
- Sowiński, J., & Dalla Bella, S. (2013). Poor synchronization to the beat may result from deficient auditory-motor mapping. *Neuropsychologia*, *51*(10), 1952–1963. <https://doi.org/10.1016/j.neuropsychologia.2013.06.027>
- Spence, C., Baddeley, R., Zampini, M., James, R., & Shore, D. I. (2003). Multisensory temporal order judgments: When two locations are better than one. *Perception & Psychophysics*, *65*(2), 318–328. <https://doi.org/10.3758/BF03194803>
- Spence, C., Shore, D. I., & Klein, R. M. (2001). Multisensory prior entry. *Journal of Experimental Psychology: General*, *130*(4), 799–832. <https://doi.org/10.1037/0096-3445.130.4.799>
- Sugano, Y., Keetels, M., & Vroomen, J. (2010). Adaptation to motor-visual and motor-auditory temporal lags transfer across modalities. *Experimental Brain Research*, *201*(3), 393–399. <https://doi.org/10.1007/s00221-009-2047-3>
- Sugano, Y., Keetels, M., & Vroomen, J. (2016). Auditory dominance in motor-sensory temporal recalibration. *Experimental Brain Research*, *234*(5), 1249–1262. <https://doi.org/10.1007/s00221-015-4497-0>
- Surlykke, A., & Kalko, E. K. V. (2008). Echolocating Bats Cry Out Loud to Detect Their Prey. *PLoS ONE*, *3*(4), e2036. <https://doi.org/10.1371/journal.pone.0002036>
- Thaut, M. H., Miller, R. A., & Schauer, L. M. (1998). Multiple synchronization strategies in rhythmic sensorimotor tasks: Phase vs period correction. *Biological Cybernetics*, *79*(3), 241–250. <https://doi.org/10.1007/s004220050474>

- Tomassini, A., Gori, M., Burr, D., Sandini, G., & Morrone, M. C. (2012). Active movement restores veridical event-timing after tactile adaptation. *Journal of Neurophysiology*, *108*(8), 2092–2100. <https://doi.org/10.1152/jn.00238.2012>
- Tong, J., Ngo, V., & Goldreich, D. (2016). Tactile length contraction as Bayesian inference. *Journal of Neurophysiology*, *116*(2), 369–379. <https://doi.org/10.1152/jn.00029.2016>
- Torre, K., Varlet, M., & Marmelat, V. (2013). Predicting the biological variability of environmental rhythms: Weak or strong anticipation for sensorimotor synchronization? *Brain and Cognition*, *83*(3), 342–350. <https://doi.org/10.1016/j.bandc.2013.10.002>
- Tricas, T. C. (1982). Bioelectric-Mediated Predation by Swell Sharks, *Cephaloscyllium ventriosum*. *Copeia*, *1982*(4), 948. <https://doi.org/10.2307/1444104>
- Trojan, J., Heil, M., Maihöfner, C., Hölzl, R., Kleinböhl, D., Flor, H., & Benrath, J. (2014). Spatiotemporal integration of tactile patterns along and across fingers. *Neuropsychologia*, *53*, 12–24. <https://doi.org/10.1016/j.neuropsychologia.2013.10.019>
- Vaina, L. M. (1989). Selective impairment of visual motion interpretation following lesions of the right occipito-parietal area in humans. *Biological Cybernetics*, *61*(5), 347–359. <https://doi.org/10.1007/BF00200800>
- van Beers, R. J., Sittig, A. C., & Gon, J. J. D. van der. (1999). Integration of Proprioceptive and Visual Position-Information: An Experimentally Supported Model. *Journal of Neurophysiology*, *81*(3), 1355–1364. <https://doi.org/10.1152/jn.1999.81.3.1355>
- van der Burg, E., Alais, D., & Cass, J. (2013). Rapid Recalibration to Audiovisual Asynchrony. *Journal of Neuroscience*, *33*(37), 14633–14637. <https://doi.org/10.1523/JNEUROSCI.1182-13.2013>
- Van der Burg, E., Orchard-Mills, E., & Alais, D. (2015). Rapid temporal recalibration is unique to audiovisual stimuli. *Experimental Brain Research*, *233*(1), 53–59. <https://doi.org/10.1007/s00221-014-4085-8>
- Van Mier, H. I. (2013). Effects of visual information regarding allocentric processing in haptic parallelity matching. *Acta Psychologica*, *144*(2), 352–360. <https://doi.org/10.1016/j.actpsy.2013.07.003>
- Vercillo, T., Carrasco, C., & Jiang, F. (2017). Age-Related Changes in Sensorimotor Temporal Binding. *Frontiers in Human Neuroscience*, *11*, 500. <https://doi.org/10.3389/fnhum.2017.00500>
- Vercillo, T., Milne, J. L., Gori, M., & Goodale, M. A. (2015). Enhanced auditory spatial localization in blind echolocators. *Neuropsychologia*, *67*, 35–40. <https://doi.org/10.1016/j.neuropsychologia.2014.12.001>
- Virsu, V., Oksanen-Hennah, H., Vedenpää, A., Jaatinen, P., & Lahti-Nuuttila, P. (2008). Simultaneity learning in vision, audition, tactile sense and their cross-modal combinations. *Experimental Brain Research*, *186*(4), 525–537. <https://doi.org/10.1007/s00221-007-1254-z>

- von Hofsten, C., & Rösblad, B. (1988). The integration of sensory information in the development of precise manual pointing. *Neuropsychologia*, *26*(6), 805–821. [https://doi.org/10.1016/0028-3932\(88\)90051-6](https://doi.org/10.1016/0028-3932(88)90051-6)
- Voss, P., Lassonde, M., Gougoux, F., Fortin, M., Guillemot, J.-P., & Lepore, F. (2004). Early- and Late-Onset Blind Individuals Show Supra-Normal Auditory Abilities in Far-Space. *Current Biology*, *14*(19), 1734–1738. <https://doi.org/10.1016/j.cub.2004.09.051>
- Welch, R. B., & Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin*, *88*(3), 638–667. <https://doi.org/10.1037/0033-2909.88.3.638>
- Westheimer, G. (1999). Discrimination of short time intervals by the human observer. *Experimental Brain Research*, *129*(1), 121–126. <https://doi.org/10.1007/s002210050942>
- Wiener, M., Zhou, W., Bader, F., & Joiner, W. M. (2019). Movement Improves the Quality of Temporal Perception and Decision-Making. *Eneuro*, *6*(4), ENEURO.0042-19.2019. <https://doi.org/10.1523/ENEURO.0042-19.2019>
- Wing, A. M., Dumas, M., & Welchman, A. E. (2010). Combining multisensory temporal information for movement synchronisation. *Experimental Brain Research*, *200*(3–4), 277–282. <https://doi.org/10.1007/s00221-009-2134-5>
- Yang, J., Ouyang, F., Holm, L., Huang, Y., Gan, L., Zhou, L., Chao, H., Wang, M., He, M., Zhang, S., Yang, B., Pan, J., & Wu, X. (2018). Tapping ahead of time: Its association with timing variability. *Psychological Research*. <https://doi.org/10.1007/s00426-018-1043-2>
- Zampini, M., Shore, D. I., & Spence, C. (2003). Audiovisual temporal order judgments. *Experimental Brain Research*, *152*(2), 198–210. <https://doi.org/10.1007/s00221-003-1536-z>
- Zendel, B. R., Ross, B., & Fujioka, T. (2011). The Effects of Stimulus Rate and Tapping Rate on Tapping Performance. *Music Perception*, *29*(1), 65–78. <https://doi.org/10.1525/mp.2011.29.1.65>

9

Acknowledgements

The past few years have been full of ups and downs, and I owe a great deal of gratitude to everyone who has celebrated the highs with me and helped me through the lows.

Eli and Jeroen, thank you both for inviting me to work with you and for all your help along the way. I really appreciated that your doors were always open to me and that you were always eager to think along and provide feedback.

Albert, Fabio, Gabriel, Irene, Luigi, and Monica, thank you all for your input and support during our collaborations. You broadened my horizons and encouraged thoughtful discussions during our time working together.

Emily, Frans-Jozef, Katinka, Kiki, Myrthe, Nina, Rob, and everyone else in TC3, thank you for the tea meetings, the coffee breaks, the help with my experiments, and of course the gezelligheid. Ali, Belle, and Kelsey, thank you for making my final year at the VU so much better. To everyone else in the FGB, I didn't spend nearly enough time with you, but when I did it was always a pleasure.

To all the late-lunch buitenlanders, Basti, Cris, Nico, Yajie, and everyone else, thank you for making me laugh and keeping me sane. My time at the VU would not be the same without you.

The PACE network made my time as a PhD candidate fun and unforgettable. Sarah and Elise, thank you both for your hard work making the ITN everything that it was. To all of the ESRs, thank you for all the fun times we had across Europe. I always looked forward to seeing each of you and wished the meetings could have lasted longer. Adam and Alix, you made the Netherlands feel more like home.

Oran, Susannah, and Thomas, thank you all for the fun (and not so fun) times with bouldering, for our awesome road trip right before the world went into lockdown, for encouraging and supporting my photography hobby, and for our teledrinking sessions when we couldn't see each other in person. You have really helped keep me grounded these past few years.

Alex, Anna, Ben, Chad, Crystal, Erin, Janin, Jess, Katharina, Katie, Lia, Lisa, Marielle, Oscar, Per, Rachel, Simon B, Simon H, and Solomon, I am extremely lucky to have friends like you in my life. I am who I am today thanks to all of you. Our friendships have stood the tests of distance and time, and frankly you each deserve (at least) a full page here.

Voor mijn schoonfamilie, in het bijzonder Jan, Thea, Ceciël, Kees, Nayri, Samuel, Lukas en Jeannette: ik zal nooit vergeten hoe jullie mij hebben opgenomen als lid van de familie vanaf het moment dat we elkaar ontmoetten. Jullie hebben altijd je best gedaan om me welkom te laten voelen bij jullie thuis, tijdens jullie avonturen en tijdens jullie gesprekken. Bedankt dat jullie er altijd voor me zijn.

To my mom, dad, and Glenna, thank you so much for everything you have done and sacrificed for me over the years. I know that raising and growing up with me was often difficult, and I wouldn't be here today without your limitless patience and forgiveness. I love you all and am more grateful to each of you than you will ever know.

And lastly, Carien. Words cannot begin to express how glad I am to have you in my life. You made Boulder even more fun than it already was, and you helped me survive Bonn. You made sure that I always felt like I belonged in the Netherlands. You encourage me to pursue my goals, and you challenge me to be a better person. You bring out the best in me and you put up with the worst in me. Even though we're living in the timeline where there's a deadly pandemic and not the one where we've gotten married and traveled the world, that hasn't stopped us from making the best of things. I can't wait to see where life takes us next, and I am grateful that you have always been there for me. I love you.