Making sense of haptics

van Beek, F.E.

2016

document version
Publisher's PDF, also known as Version of record

citation for published version (APA)
van Beek, F. E. (2016). *Making sense of haptics: fundamentals of haptic perception and their implications for haptic device design*. [PhD-Thesis - Research and graduation internal, Vrije Universiteit Amsterdam].

Chapter 9

General discussion

The aim of this thesis was to investigate parameters in haptic perception that are important for designing haptic devices and haptic feedback. The data gathered in this investigation cover fundamentals of static haptic perception, fundamentals of dynamic haptic perception and studies into the design of intuitive haptic guidance. In this discussion chapter, these data will be discussed in the light of the aim of the thesis. First, the fundamental insights which this thesis provides for haptic perception will be discussed. Next, the reproducibility of the results is discussed. Finally, recommendations on how these results could be applied to haptic device and haptic guidance design will be given.

9.1 Fundamental insights

As stated in the Introduction (Chapter 1), biases and discrimination thresholds are the main parameters that are used in psychophysical research to describe human performance (Jones & Tan, 2013). Knowing these parameters can help to determine how to design your haptic device in a way that it fits the human needs and is intuitive to use. Using this fundamental approach can be seen as working from general knowledge (knowledge on human perception) to a more specific application (the design of a particular device), so this approach could be called a deductive one (Kimmig, 2013). In most technology-oriented research, the problem is usually approached from the other direction: devices are designed and tuned based on previous designs and tunings, and then they are tested in a human factors experiment. If the results are satisfactory, the device is deemed good enough. If not, the tuning or design is changed and the human factors experiment is repeated. When a lot of these experiments have been performed, general design rules might be inferred from the experiments, so this would be a more inductive way of building theory (Heit, 2000). Obviously, there are advantages to the inductive approach: designers perform tests on the parameter that they are actually interested in, which is the performance of the operator when using that particular device. However, the pitfall of this approach is that it is difficult to understand why some devices and tuning parameters perform well and others do not. When no fundamental knowledge on the problem is gained, the best-case-scenario is that general guidelines can be formulated. So, the approach in this thesis was a deductive one: by investigating human perception, information about the precision and accuracy of perception was obtained. This information can be used to understand how to design haptic devices in general, since the information does not apply to one specific device.

9.1.1 Part I: Static perception

The first part of the thesis describes a situation in which haptic information is acquired in a static fashion, so in which stimuli are applied to the stationary human hand. The obvious parameter to look at in this context is force perception: how do humans perceive that someone is pulling or pushing their hand? In literature, a lot of attention has been devoted to investigating *discrimination thresholds* of force perception (see Jones (1986) for a review). However, *biases* in force perception in directions other than that of gravity have hardly been studied, so that is the topic of the first part of this thesis.

In Chapters 2 and 4, direction-dependent biases in force magnitude perception were found, which were very consistent across participants. Forces exerted perpendicular to the line between shoulder and hand were perceived as being 50% larger than forces exerted along this line. These biases were not directly related to the measured arm dynamics parameters, so the arm dynamics parameters cannot be the sole contributors to the perceptual effect, but they can still play a role in causing it.

In Chapters 2 and 3, direction-dependent biases were also found in force direction perception. The biases ranged between -30° to 60° and varied greatly between participants and force directions, resulting in participant-dependent error patterns. No groups of similar patterns could be discerned within the total group of participants, and the patterns could not be explained using general subject parameters, like arm length (Chapter 3). Nonetheless, the patterns were found to be consistent within participants, even when they were measured a month later (Chapter 3).

9.1.2 Part II: Dynamic perception

The second part investigates what happens when humans start moving their hands. Firstly, they need to sense the hand movements, so the perception of movement is the first topic of interest. Secondly, they make hand movements in order to interact with something, so the perception of objects during movement is the second topic of interest.

In Chapter 5, discrimination thresholds for movement distance were investigated using a task in which two movement distances had to be compared. In contrast to the literature on force perception, the literature on movement distance mainly reports *biases* (e.g. Bergmann Tiest et al., 2011; Fainetеau et al., 2003; Hermelin & O'Connor, 1975; Wong, 1977) and has not devoted much attention to *discrimination thresholds*. So, in this study, discrimination thresholds were tested in conditions that were known to affect biases in movement distance. In most conditions, we found no effect of that condition on the discrimination threshold. Generally speaking, for movement distances of 25 and 35 cm, a Weber fraction of about 11% was found along all cardinal axes. For passive movements, the threshold was a bit higher, and adding cutaneous information did not improve the precision. Movement parameters, which were recorded during the task, showed that participants generally adopted the strategy to keep the speed profiles of two movements the same and to compare the resulting movement times in order to infer the movement distances.

In Chapter 6, the perception of object hardness was investigated for different types of movements. Object hardness is an important object property in teleoperation tasks, since a proper perception of it is essential to assess how much force is needed to manipulate the object, while not breaking it. A large effect of damping on perceived hardness was found. This effect was task-dependent: for an in-contact task (in which the movement started at the object's surface), adding global damping increased perceived hardness, while for a contact-transition task (in which there was a free-air phase before making contact), adding global damping decreased perceived hardness. The latter effect was much larger than the former. The movement parameters revealed that this task-dependency was not caused by a change of movement strategy, since for both tasks in all experiments, participants used the same parameters to base their perceptual decision on. So, an actual difference in task dynamics probably caused the task-dependency of the effect. In both Chapters 5 and 6, analyzing the movement parameters helped to understand the strategies responsible for the perceptual outcomes. In perceptual research, movement parameters are not always investigated, not even in dynamic tasks. The results in this thesis suggest that it would be interesting to include an analysis of movement parameters in future research on dynamic tasks.

9.1.3 Part III: Applications

The results in the last part of this thesis, consisting of Chapters 7 and 8, could be applied more directly to haptic guidance. However, the chapters also provide insight in multisensory integration, since they involve the perception of multiple sources of information.

Chapter 7 revolves around the integration between position and force information in a task of finding the center of a weak force field. Of course, this easily relates to force fields used in haptic guidance, but it also provides fundamental information about the integration of information. It has been shown that humans can integrate information from multiple information sources (e.g. Van Beers et al., 1999). The data in this chapter show that in this task, force and position are not integrated into a percept of stiffness, but humans rather simply use the position(s) where the force reaches the detection threshold level to estimate the center of the force field. This demonstrates that information does not have to be integrated, even if multiple sources of information are available. The lack of integration could be caused by the serial nature of the task: obtaining enough information to make a stiffness estimate would require storing information from serial exploration, which is a costly process (Dopjans et al., 2012; Loomis et al., 1991), while the observed strategy is of a far less serial nature and might therefore be preferred.

Chapter 8 is most interesting from an application point of view, since it mainly shows that correcting for perceptual biases in haptic guidance can increase user performance. In addition, it shows that visuo-haptic biases in 3D are mostly oriented along the depth direction. Most participants overshot the target, when pointing to a visual target without seeing their hand. This is in accordance with literature on similar tasks (Adamovich et al., 1998; Van Beers et al., 1998; Sousa et al., 2010).

In almost all the chapters, significant perceptual biases were reported. It is thus safe to say that human perception is often not veridical, which means that the perceptual world often does not scale linearly with the physical measures that we use to describe the world with. So, why are we still able to perform actions at all? Why don't we bump in to objects all the time, as, for instance, the perceived length of our arm movements depends on our movement direction? It could be that we have learned strategies to cope with these biases, since we have encountered them throughout our lives. Interestingly, humans tend to believe that others are prone to be biased, but usually they have the feeling that their own observations are bias-free, a phenomenon called 'the bias blind spot' (Pronin, Lin, & Ross, 2002; Pronin, 2007). Of course, humans might be able to correct for their biases subconsciously, even when they are not aware of their own biases. However, it would not make much sense that human beings, who have evolved over the course of millions of years, would end up with being stuck with magically imposed biases for which they have to correct in all their actions. It is far more probable that these biases actually reveal functional systems in sensorimotor control. In other words: perceptual biases are probably often related to the action connected to the percept. In Chapter 4, we could not find a direct link between force magnitude perception and arm biomechanics, but links between perception and action have been described throughout the perceptual literature (e.g. Ahmed, Wolpert, & Flanagan, 2008; McCloskey et al., 1974). So, in future research, it would be insightful to not only measure perceptual biases, but to also speculate — and if possible perform experiments — on the link between perceptual biases and the action connected to the perceived property.

9.2 Reproducibility

An important question in all fields of research is the reproducibility of results. Recently, psychological research received a particularly hard blow when a large consortium tried to reproduce 100 psychological studies, which had all been published in high-quality journals (Open Science Collaboration, 2015). Whereas in the original publications 97 of the 100 papers reported a significant effect, only 36 of the 100 replications showed a significant effect. Fortunately, hardly any of the results in this thesis rely on a single set of data, as will be shown in this section.

In the two chapters on force magnitude perception, Chapters 2 and 4, very similar biases were found, even though the former chapter describes biases measured in 2D using a custom-made setup, while the latter chapter describes biases measured in 3D using the

HapticMaster (compare Figure 2.4 and 4.4). The data on arm biomechanics in Chapter 4 have not been reproduced in this thesis, but are very comparable to values reported in literature (Krutky et al., 2009; Pierre & Kirsch, 2002). For the two chapters on perception of force direction, Chapters 2 and 3, it is a bit harder to draw conclusions, since there were larger differences between participants and different participants took part in the different experiments. However, what is apparent is that most participants showed considerable perceptual biases in at least some directions, and all data sets show large variations between participants (compare Figure 2.6a, 3.2, and 3.3). The conclusions in Chapter 3 on the origin and consistency of the error patterns are the only conclusions in this part of the thesis that rely on a single set of data. The absence of a correlation between the participant's general characteristics and the error patterns in force direction perception could be due to a group size ($n = 25$) that might be too small for this type of correlational research. However, none of the p-values were close to significant (all $p \geq 0.15$), so it is not very likely that a larger group size would have yielded significant results. The other experiment in this chapter, showing the consistency of the patterns within participants, did not yield reproducible patterns for all participants. However, for most participants the patterns were reproducible, therefore the conclusion was drawn that generally speaking, the patterns are reproducible. A replication of this experiment using a larger group of participants could be useful in order to draw firmer conclusions.

The two chapters on perception of dynamic parameters also show consistent results. The data on discrimination of movement distance, presented in Chapter 5, have not been reproduced in this thesis, but all the different research questions yielded biases in the same magnitude range (see Figure 5.4). For research questions in which a main effect of condition was found, this effect was highly significant (all $p \leq 0.003$). The data on the influence of damping on hardness perception, Chapter 6, were reproduced in the same experiment, since Experiments 1 and 2 were performed with different groups of participants and yielded comparable results (compare Figure 6.5 and 6.6). Experiment 3 had to be repeated with a different group of participants because of errors in the measuring procedure (only the data of the last group are shown in Chapter 6), but the results of the replications were very similar.

Finally, the two more applied chapters, Chapters 7 and 8, were both follow-up studies on themes that some of the authors had been working on previously. The current results are all in line with the results from the previous studies. The pattern in the biases found in the chapter on the integration of force and position information are very comparable to the biases reported in Baud-Bovy (2014). The results of the chapter on adjusting haptic guidance to individual biases are also very comparable to the results in Kuling, van Beek, et al. (submitted), even though the biases and effect sizes in the previous study in 2D were much smaller than those in the current study in 3D.

Summarizing, most of the data in this thesis have been reproduced in similar experiments in one of the chapters or are comparable to data in similar experiments in literature, which ensures that the chance that the reported effects are reproducible is considerable. It is hard to make claims on the generalizability of the results, since it is known in perceptual research that small adjustments to research designs can have massive effects on the size of perceptual biases (see for instance the review by McFarland and Soechting (2007) on all the factors influencing the size of the radial-tangental illusion). Some of the experiments in this thesis have been reproduced using different devices, such as the force magnitude perception experiment, which has first been performed on a custom-made device (Chapter 2) and then reproduced using a HapticMaster (Chapter 4), while yielding the same results. The experiment on correcting for visuo-haptic biases (Chapter 8) was performed on a HapticMaster, while the first experiment in that research line was performed on a PHANToM Premium (Kuling, van Beek, et al., submitted). However, these are all still ‘devices’, so whether these results can be translated to a natural situation in which you perform actions with your own hands, without any restrictions on, for instance, hand movement, remains to be investigated.

9.3 Implications for design

The results from the more fundamental parts of this thesis also provide information on human perception that can be useful to consider when designing haptic devices. In the next sections, the possible applications of this fundamental knowledge are explained.

9.3.1 Haptic parameters for haptic device design

Chapters 2 and 3 show that perception of force direction is not veridical, and it is also very different between participants. A general observation that can be drawn from these results is that it is probably not very useful to use force direction as a way to communicate information. Mean individual biases ranged between -30° and 60°, so one physical force direction can result in a perceptual direction with a range of 90°. Note that this does not mean that differences in direction cannot be perceived reliably; the *discrimination threshold* for force direction is 30° (Barbagli et al., 2006; Tan et al., 2006; Ho et al., 2006), so differences in force direction exceeding this value can be perceived reliably. However, this thesis shows that the *accuracy* of the perception of force direction is generally very poor.

It is probably not necessary to build haptic devices that outperform the human capacities, so the results on the precision of distance perception (Chapter 5) show that, to communicate distance information, a precision better than 11% of the movement distance is acceptable. For the smallest movement distance in our study (15 cm), the Weber fraction was higher than for the other two distances, so probably there is an absolute limit in precision for smaller distances, but it is likely that we did not reach that yet at a movement distance of 15 cm. The data also show that adding cutaneous information does not improve precision, so in order to convey precise distance information, it is not necessary to simulate surfaces.

The results on the influence of damping on perceived hardness (Chapter 6) show that adding global damping to a system has a large effect on contact-transition tasks. Since contact-transition situations are common in teleoperation applications (Sarkar & Yun, 1996), designers should be careful about adding damping to their system to increase its stability, if the operator is required to have a proper perception of object hardness. As a start of a guideline, Figure 6.7 could be used, which describes the interplay between stiffness and damping in perceived hardness. Of course, these values have been measured on a specific device (the HapticMaster) in particular tasks, so there is no guarantee that the values are the same for other tasks or devices. Therefore, these specific values can only be seen as a start of a more general guideline. In general, the best solution from a human perspective would be to avoid inserting (a lot of) damping at all, which is possible when using less conventional designs, such as the one proposed in Heck, Saccon, and Nijmeijer (submitted). These authors propose a scheme in which the energy flows of the master and the slave are constantly monitored, and damping is only inserted when the system's stability is in danger.

9.3.2 Haptic parameters for haptic guidance design

How to design the force field in haptic guidance in a way that is intuitive for the human operator? This thesis surely does not provide a full answer to this question, but the presented data do provide ideas on adjustments that could make haptic guidance more intuitive.

Chapters 2 and 4 show that perception of force magnitude is not veridical, but shows a direction-dependent distortion. This means that it might be useful to correct for this distortion when the aim is to provide the same message throughout the force field, independent of the direction of the force. So, it might be useful to design a force field that is perceptually equal in all directions, rather than physically equal. It is not clear how to do this yet, since the data in these chapters have been measured using one arm posture, while in a real teleoperation situation, the operator probably moves his arm. Unfortunately, no direct relationship between arm dynamics and force magnitude perception was found in Chapter 4. There is a strong relation between arm dynamics and posture (Milner, 2002; Tsuji et al., 1995), so if arm dynamics had been directly related to the perceptual distortion, arm posture could have been used to predict the effects on the perceptual distortion.

In future research, it would therefore be useful to measure the distortion in force magnitude perception for various arm postures. This could shed light on relations between the perceptual biases and arm posture, and possibly also on their relation with arm dynamics.

The design of force fields for haptic guidance that are intuitive to use is not very straightforward, as shown in Chapter 7. This chapter shows that, when trying to find the center of a force field, which is also the basic task in a haptic guidance situation, humans do not integrate position and force information into a stiffness percept. Our study showed some consequences of this notion, relating to asymmetric force fields, uni-lateral force fields and visual information. In asymmetric force fields, large biases in finding the center of the force field were found, which were oriented towards the weaker side of the field. A similar result was found for uni-lateral force fields, in which participants were asked to find the edge of the force field. In that situation, participants always perceived the edge of the force field to be somewhat inside the force field. The latter finding raises an interesting point for haptic guidance applications, as operators usually do not overshoot the target very far. Therefore, the chance that their movement reaches the part of the force field with forces above the detection threshold level behind the target is not very large, thus operators are mainly exposed to a uni-lateral situation. In this situation, participants perceive the center of the force field, and thus the target, to be at the position where the force reaches the detection threshold level. So, it could be worthwhile to position the force field in such a way that the force magnitude at the target is not 0, but is at the detection threshold level. This does make the design of the force field more complicated when the force field needs to be multi-dimensional, but it would be worthwhile to test this approach in future research. A final observation from this study is that visual information about hand position, while it was not informative about the position of the force field, still caused an increase in precision. So, depending on which types of information the operator uses to perform a task, it might be worthwhile to provide visual information, even if this seems to be pointless for the direct task objective. All the observations in this chapter could be explained using a mathematical model. When trying to assess the perceptual consequences of haptic guidance design choices, this model could be used to obtain a general idea of the direction of the perceptual effects.

In this thesis, a lot of studies report biases in human perception, of which some are participant-specific, such as the biases in the perception of force direction (Chapter 2 and 3). It has already been shown that correcting the mapping between operator and slave movements, by using parameters that are consistent across participants, increases user performance (Pierce & Kuchenbecker, 2012). In the final chapter, Chapter 8, we aimed to test if it is also useful to correct for participant-specific biases in haptic perception. To investigate this, we wanted to use a well-known paradigm for participant-specific biases, so we chose the paradigm of visuo-haptic biases (Soechting & Flanders, 1989). The task for the participants was to move their unseen hand to a visual target, which usually results in large errors. First, the biases were measured, after which two types of haptic guidance were compared: one that was adjusted to correct for the biases and one that was not. Adjusting for the biases significantly improved user performance. In our study, the feedback was adjusted by correcting the visual information: the visual image was shifted, while the haptic guidance was directed towards the original target position. In the first experiment on this idea, a similar approach was used, but in that case, the position of the haptic guidance was corrected, while the visual image remained at the target location (Kuling, van Beek, et al., submitted). In that approach, participants did not end up at the target location, but at a location shifted by their visuo-haptic bias. However, since the bias was known, their end position was predictable. This was the objective of the study, since in a teleoperation setting, predictable biases on the master side can be corrected for on the slave side, which ultimately ensures that the movement of the slave reaches the target position. Nonetheless, our approach in Chapter 8 seems to be more practical: the only adjustment that is needed is a shift of visual information to increase operator performance. It is noteworthy that the measurement of the visuo-haptic bias was a bit more tricky in our approach: usually the visuo-haptic bias congruent with a visual position is measured, which can easily be done by asking participants to point to a visual target, as was also done in Kuling, van Beek, et al. (submitted). In the current approach, we needed to measure the bias congruent with a haptic location, so an adaptive paradigm was used to find the visual location which matched the haptic representation of the target location. In a practical setting, it is imaginable to measure the operator-specific biases once every month using a task that is comparable to our adaptive paradigm, after which the personalized settings can be used throughout that month. Most of these biases are fairly stable over time (Chapter 3 and Kuling et al. (submitted)), so the benefits should be fairly long-lasting.

In this thesis, it has been shown that human perception is often not veridical, but shows consistent and reproducible biases in both static and dynamic perception. This information can be used in the design of haptic feedback devices and in the design of haptic guidance. Taking human perception into consideration when building devices that humans need to operate is a logical step to enable the construction of intuitive, human-centered systems.