

Tactual Articulatory Feedback on Gestural Input

Bert Bongers ^{1,*} and G. C. van der Veer ²

¹ Faculty of Design, Architecture and Building, School of Design, University of Technology Sydney (UTS), Ultimo, NSW 2007, Australia

² Faculty of Sciences, Department of Computer Science, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, The Netherlands; g.c.vander.veer@vu.nl

* Correspondence: bert.bongers@uts.edu.au

Abstract: The role of the sense of touch in Human–Computer Interaction as a channel for feedback in manipulative processes is investigated through the research presented here. The paper discusses how information and feedback as generated by the computer can be presented haptically, and focusses on the feedback that supports the articulation of human gesture. A range of experiments are described that investigate the use of (redundant) tactual articulatory feedback. The results presented show that a significant improvement of effectiveness only occurs when the task is sufficiently difficult, while in some other cases the added feedback can actually lower the effectiveness. However, this work is not just about effectiveness and efficiency, it also explores how multimodal feedback can enhance the interaction and make it more pleasurable—indeed, the qualitative data from this research show a perceived positive effect for added tactual feedback in the overall experience. The discussion includes suggestions for further research, particularly investigating the effect in free moving gestures, multiple points of contact, and the use of more sophisticated actuators.

Keywords: multimodal interaction; tactual articulatory feedback; haptics



Citation: Bongers, A. J., and G. C. van der Veer. 2023. Tactual Articulatory Feedback on Gestural Input. *Arts* 12: 146. <https://doi.org/10.3390/arts12040146>

Academic Editors: Justin Paterson and Marcelo M. Wanderley

Received: 16 February 2023

Revised: 11 April 2023

Accepted: 18 April 2023

Published: 10 July 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Background

Many everyday interactions with physical objects that we manipulate are guided by a variety of (redundant) feedback modalities; for instance, through touch, movement, balance, sound, smell, or vision. However, interactions with digital technologies (computers, smartphones, tablets) are often restricted to the visual modality, and touch screens and trackpads offer very little clues in other modalities about what is being ‘touched’ or manipulated on the screen or on the keyboard. The human sense of touch is barely actively addressed by these systems. All one can feel are things like the texture of the mouse pad, the resistance of keys, or the surface of the glass, which is only limited as a source of information, as it is not related to the virtual world of data inside the computer.

The computer has a wide variety of information to (re)present, such as feedback about processes, information about objects, representation of higher-level meaning, and indications of relevant aspects of the system state. Computers, smartphones, and tablets are presenting this information mainly visually, and in some cases supported with sounds.

This paper describes a research project that investigated if there is an advantage in using tactual feedback as secondary notation or added feedback, when manipulating the virtual objects in the computer (such as widgets, icons, and menus). We were looking particularly at feedback that supports the *articulation* of a gesture, improving the fine control of a gestural expression as it progresses (in Section 1.1 below, different types of feedback are described). A standard mouse input device was used, with an added vibrotactile feedback element (custom-made, based on a small loudspeaker) placed under the tip of the index finger of the preferred hand, as shown in Figure 1.



Figure 1. Mouse with the vibrotactile element.

The core of this research took place around 2005, the experiments were carried out in an academic setting with sufficient methodological rigour as part of the first author's PhD research at the Vrije Universiteit Amsterdam (Bongers 2006). For various reasons, the work was never published at the time. We believe the findings still have merit; even though some of the technologies applied in this research have changed, none of it affects the findings and the main conclusions around the actual interactive haptic experiences and measurements. The introductory paragraphs in Section 1 have all been thoroughly rewritten or added later to reflect the current state of the art and technological advances. Although the research was carried out with a computer mouse, the notion of the importance of tactual feedback in the articulation of gestures comes from earlier work carried out in musical gestures and new electronic musical instruments. This is discussed in Section 1.1 below, which furthermore discusses several different forms of feedback and their role in interaction. To analyse interactions in sufficient detail, Levels of Interaction are discussed in Sections 1.2 and 1.3, placing the level of physical interaction and articulation in a wider interactional context. A concise summary of the various sub-modalities of tactual perception is given in Section 1.4. Section 1.5 gives an overview of haptic interaction devices and related research in this field, and Section 1.6 describes previous work on vibrotactile feedback.

Section 2 describes the experimental setup, method, and results. In Section 3, the results are discussed and concluding remarks and future research directions are presented in Section 4.

1.1. Feedback and Feed-Forward, Active and Passive

The feedback that is presented in order to guide the users' actions, to support them in articulating their intentions, is called *articulatory feedback* (Hix and Hartson 1993, p. 40) or synchronous feedback (Erickson 1995). Articulatory feedback is particularly relevant in gestural interfaces, and particularly traditional musical instruments show the importance of this, often using tactual modalities (Chafe 1993; Chafe and O'Modhrain 1996). A recent book, *Musical Haptics*, presents a range of concepts and research from relevant practitioners in this field (Papetti and Saitis 2018). An earlier exploration by the first author in the early 1990s involved a custom-made tactile actuator (based on a relay electromagnetic actuator) built in a ring worn by the player, presenting articulatory feedback on free-moving gestures of the player, in the context of new electronic musical instruments (Bongers 1998, 2007).

When the system is generating information that is actively drawing the user to something or somewhere, it is generally referred to as *feed-forward*. The strongest examples of this can be found in haptic systems, that actively 'pull' the user towards some location. When a target or goal is reached, *confirmation feedback* can be generated and this active haptic event can be very convincing and appropriate.

Feedback can come from elements specially designed for that purpose, such as message boxes and widgets that allow manipulation, or from the content itself (*content feedback*, sometimes called function feedback (Wensveen et al. 2004)).

Most feedback is actively generated by the system, but some feedback can come from passive elements of the system. An example of such passive feedback, sometimes called inherent feedback (Schomaker et al. 1995; Wensveen et al. 2004), is the mouse click felt when the button is pressed—it is useful information but not generated by the system. The system therefore cannot be held responsible for misinterpretations—the click will still be felt even if there is nothing to click on the screen, the machine has crashed, or is not even on. In fact, this is a bit of real-world information, blending in an often useful way with the virtual world information. Mixing real and virtual phenomena is a good thing as long as one is aware of it, and it is designed as a whole. In the example of the mouse click, it can be said that it is estimated information—usually true, but not always.

Feedback information can be displayed in various ways, addressing various human senses. In each case, particularly as part of parallel and multitasking situations, the most appropriate modality has to be chosen to convey the meaning and optimally support the interaction.

1.2. Levels of Interaction

In order to understand the interaction between human and technology, it is useful to discern various *levels of interaction*. An action is usually initiated in order to achieve some higher order goal or intention, which has to be prepared and verbalised, and finally presented and articulated through physical actions and utterances. The presentation and feedback by the computer passes through several stages as well before it can be displayed, possibly in various modalities including the haptic, in order to be perceived by the user. The actual interaction takes place at the physical level. A dialogue is often described in three levels: semantic, syntactic, and lexical (Dix et al. 1993, chp. 8), but for more specific cases, more levels can be described. Jakob Nielsen's virtual protocol model (Nielsen 1986) is an example of this, specifying a task and a goal level above the semantic level, and an alphabetical and physical level below the lexical level. This can also be applied on direct manipulation interface paradigms. The Layered Protocol is an example of a more complex model (Taylor 1988), particularly to describe the dialogue using the speech modality, but also applied on general user interface issues (Eggen et al. 1996). When more sensory modalities are included in the interaction, models often have to be refined. Applying the Layered Protocol on the interaction which includes active haptic feedback, introduces the idea of (higher level) E-Feedback which has to do with *expectations* the user has of the system, and the I-Feedback which communicates the lower-level *interpretations* of the user's actions by the system (Engel et al. 1994).

A framework for analysing and categorising interaction by levels of interaction, modalities, and modes has been developed by the authors and presented elsewhere (Bongers and van der Veer 2007).

The articulatory feedback, which is studied in the research described in this paper, takes place at the physical level and potentially extends into the semantic levels.

1.3. Feedback and Articulation

In the final phase of the process of making an utterance or a gesture, we rely on feedback in order to shape our actions. Most actions are taking place in a continuous loop of acting, receiving and processing feedback, and adjusting. For instance, when speaking, we rely on the auditory feedback of our utterances and continuously adjust our articulation. When this is absent or suppressed it becomes harder to articulate. (One can try this when speaking while wearing active noise cancelling headphones.) When manipulating objects or tools, we rely on the information conveyed to our senses of vision, touch, sound, and smells, including our self-perception, in order to articulate. We have, of course, many more than the proverbial 'five senses', particularly in the somatosensory channels, bodily senses which include pain, temperature, balance (vestibular senses), vibration, pressure, hairs, muscles, tendons, and joints; in total, it could be up to 32 senses as a recent popular book argues (Young 2021).

The mainstream computer interface paradigm relies almost entirely on the visio-motor loop, sometimes with an added layer of auditory feedback such as sounds built into the operating system (an early example of this was the Sonic Finder as developed by Bill Gaver in the 1980s (Gaver 1989), and his other projects around that time).

However, in a real world ‘direct manipulation’ action, the closest sense involved in that act is usually the sense of touch. Musicians know how important the touch feedback on articulation is, and so does the craftsperson. Computers seem to have been conceived as a tool for mental processes such as mathematics (or administrative applications) rather than as a tool for manual workers, and have inherited a strong tendency of anti-crafts, and therefore anti-touch (McCullough 1996). The developments in Tangible Interaction in the last decades have challenged this notion, with many projects presented at the TEI (Tangible, Embodied, and Embedded Interaction) conferences, and books and journals.

Anticipating a further development and emphasis on gestural control of computers, including the fine movements and issues of dexterity, there will be a need for the feedback to be generated properly in order to be adequate. That is, we need feedback that supports the articulation of the gesture, enabling the user to refine the action on the fly, and becoming more articulate.

1.4. Tactual Perception

The human sense of touch gathers its information and impressions of the environment through various interrelated channels, together called *tactual perception* (Loomis and Leederman 1986) or the *somatosensory* system (Hayward 2018). These channels and their sub-channels can be functionally distinguished; in practice, however, they often interrelate and are therefore harder to separate.

The tactile perception receives its information through the *cutaneous* sense, from the several different mechanoreceptors in the skin (Sherrick and Cholewiak 1986; Linden 2015). *Proprioceptors* (mechanoreceptors in the muscles, tendons, and joints) are the main input for our *kinaesthetic* sense, which is the awareness of movement, position, and orientation of limbs and parts of the human body (Clark and Horch 1986). (Often, the term *proprioception*, meaning self-perception, inside the body, is used as a synonym for kinaesthesia, and indeed it is a bit similar but we prefer to keep this terms for its original intended meaning of self-perception).

A third source of input for tactual perception is ‘efference copy’, reflecting the state of the nerves that control the human motor system. The brain effectively has access to the signals it sends to the muscles, the efferent signals, which form an actual input to perceptual processes—*efference copy* signals. In other words, when we move (as opposed to being moved), we know that we are moving and this is a source of input as well. This is called *active touch*, as opposed to passive touch. The importance of the difference between passive and active movement was identified early in the 20th century by the Gestalt psychologist David Katz, presented in his 1925 book *Der Aufbau der Tastwelt* (much later translated in English as *The World of Touch* by Lester Krueger (Katz 1989)). The importance of this was further established by J. J. Gibson’s work on active touch in the 1960s (Gibson 1962).

Haptic perception uses information from both the tactile and kinaesthetic senses, particularly when actively gathering information about objects outside of the body. Most of the interactions in the real world involve haptic perception, and it is the main modality for applying in HCI. Haptic perception is the combination of touch (cutaneous sensitivity, the receptors in the skin) and kinaesthesia (receptors in the body). These modes can be either passive or active, the efference copy signal is a potential third source for perceptual processes. Passive touch (*tactile* perception) can only exist without self-initiated movement. Kinaesthesia can be passive or active (imposed or obtained). Active touch inherently involves movement of the body part, and therefore kinaesthesia, so the result is the (active) *haptic* mode. The *active haptic* mode is the most common and most meaningful tactual mode, used when exploring the environment and objects, particularly the manipulation of

musical instruments. The field of HCI often refers to all modes of touch and movement (particularly active) as 'haptic' (Ziat 2023).

To summarise, the five sub-modalities of tactual perception as defined by Loomis and Leederman (1986) are:

- (1) tactile perception (always passive), based on the cutaneous sensitivity of the skin,
- (2) passive kinaesthetic perception (imposed movement, picked up with the proprioceptors inside the body),
- (3) passive haptic perception (imposed movement, picked up by the receptors in the skin, and the proprioceptors inside the body),
- (4) active kinaesthetic perception (as a result of self-initiated movement, based on signals from internal receptors, and efference copy), and
- (5) active haptic perception (from self movement, in contact with objects, based on signals from internal receptors, cutaneous receptors, and efference copy).

The feedback discussed in this paper mainly involves the cutaneous sense, which is actively addressed by the system, in addition to the participant's movements. Cutaneous sensitivity is based on four sense systems involved relating to four types of mechanoreceptors in the skin, making up all four combinations of the parameters adaptivity (slow and fast, having to do with frequency) and spatial sensitivity (diffuse and punctuate) (Bolanowski et al. 1988; Linden 2015). The experiments described in this paper are particularly about feedback addressing the fast adapting and diffuse system, based on one of the four mechanoreceptors, the Pacinian corpuscles which are important for perceiving textures but also vibrations. The sensitivity of the Pacinian receptors overlaps with the audible range and therefore play a crucial role in musical applications. The vibration sensitivity range of this system has a U-shaped curve from about 40 Hz up to about 1000 Hz with a peak sensitivity between 200 and 300 Hz. This range of frequencies picked up is in the audible range of the ear, and encompasses the range of the human voice, and most of the musical pitches of traditional instruments, though with less pitch discrimination than the auditory system (20–25% vs. 1%, respectively) (Verrillo 1992).

In the section below, some examples are presented of devices and techniques to address these tactile, kinaesthetic, and haptic modes. In this interactive situation, where both technology and the person can be active, it is sometimes difficult to distinguish the actual channels but it is important to be clear. For instance, a vibrating actuator addressing the cutaneous sense (the Pacinian system) of the fingertip of a person who moves their hand would be a haptic perception (movement and touch), but the display is called vibrotactile (the movement is made by the person). This is the situation in the experiments presented in this paper. If the movement was displayed by the system, for instance through force feedback, the kinaesthetic sense (through the body) would be addressed, and usually the tactile as well (though the skin), so then it would be appropriate to describe it as an active haptic display.

1.5. Tactual Feedback Devices

The tactual senses can be addressed with a variety of technologies and displays. Reflecting the tactual modes and their interrelationships as described above, the tactile sense can be addressed with vibrotactile displays (such as speakers and other electromagnetic transducers, piezoelectric actuators, shape memory alloys (SMA)), and the haptic modalities can be influenced with force feedback and displays (such as motors, solenoids, linear actuators) which usually influence the cutaneous as well as the kinaesthetic senses. These are active tactual displays, where the technology has the ability to generate movement, texture, and shape change. Passive display comes from any shape or material that is manipulated or explored by a person, where the person is active and the technology is passive.

Several devices have been developed to actively address the human sense of touch, and many studies have shown the improvement in speed or accuracy in the interaction.

Already in the early 1990s, existing devices such as the mouse have been retrofitted with solenoids for tactile feedback on the index finger and electromagnets acting as brakes on an iron mousepad for force-feedback (Akamatsu and Sato 1992, 1994; Münch and Dillmann 1997), or solenoids on the side of the mouse for vibrotactile feedback (Göbel et al. 1995). The principle of electromagnetic braking mouse was also applied in an art interface, an interesting application where the emphasis was not on improving efficiency (Kruglanski 2000). The Logitech force-feedback mouse (introduced in 1999 and only on the market for a few years), and its predecessor the Immersion FEELit mouse, have two motors influencing the mouse movement through a pantograph mechanism. These have been used in several studies (Dennerlein et al. 2000), including studies which are a continuation of earlier work with the vibrotactile device as described in this paper but then using force-feedback (Hwang et al. 2003). All show performance improvements by applying various forms of feedback as described above.

Several motor-based force-feedback joysticks have been used for generating virtual textures (Hardwick et al. 1996, 1998), and other experiences (Minsky et al. 1990). These joysticks have become cheaply available for gaming applications, such as the devices from Microsoft and Logitech using the Immersion TouchSense protocol. Low-end devices such as the force-feedback joysticks and also steering wheels and foot pedals are found in gaming applications.

The Phantom is a well-known device; it is a multiple degree-of-freedom mechanical linkage that uses motors to generate force-feedback and feed-forward on the movements. It has been used in many studies that investigate the advantages of tactual feedback (Oakley et al. 2000, 2002). Active haptic feedback has also been applied in automotive situations, traditionally an area that has paid a lot of attention to the ‘feel’ of controls, as the eyes are (expected to be) occupied with watching the road and traffic.

To investigate the notion of E- and I-Feedback as described above, a force-feedback trackball was developed at the Institute for Perception Research in Eindhoven, The Netherlands (a collaboration between the Philips corporation and the Technical University there) using computer-controlled motors to influence the movement of the ball, enabling the system to generate feedback and feed-forward (Engel et al. 1994). This device was further used for studies in tactual perception (Keyson and Houtsma 1995; Keyson 1996), and the development of multimodal interfaces (Keyson and van Stuivenberg 1997; Bongers et al. 1998).

I realised at the time that this idea of active force feedback is even important at the level of an individual switch, and proposed a dynamic or ‘active button’, the feedback click would be generated by the system and could therefore give meaningful (contextual) feedback. With this, the click would only be made tangible if there was an active response. For instance, with a light switch, the haptic feedback would not be generated if the light was broken, or not needed to turn on a sunny day. In the mouse, trackpad, or touchscreen, the click would only be tangible if there was something to click on. It is a way of giving active feedback on the click action, and can also be used for indicating that a function is disabled. Philips didn’t pursue this further, and I moved on and watched over the years several similar implementations of the idea, particularly combined with touch screens where it is clearly beneficial (Fukumoto and Sugimura 2001; Poupyrev et al. 2002). This active button idea is more or less what Apple has introduced in the recent years with the Taptic Engine in the latest MacBook trackpads (since 2014), using electromagnetic transducer to generate the feel of a click, and in the iPhones (since the 6S in 2015) and iPads. Apple has been working on haptic feedback since the mid-1990s, at the same time as we were at Philips, apparently experimenting with using electromagnetic brakes on a mouse (which presumably could display ridges and edges, but not feed-forward) and force feedback on a trackpad-like device using rods. It is interesting to observe how long it can take for an idea to go from the R&D stage to a commercial product. The Force Click application is really well done, it feels very good, and in fact so good that many people I asked don’t even know it is a ‘virtual’ click (an active electromagnetic device instead of a passive mechanical system). The amount of feedback can be set in the System Preferences,

switchable between “light”, “medium”, and “firm” (these models are also pressure sensitive, hence the name Force Touch). Only when the MacBook is turned off entirely (and the battery removed or fully depleted, in later models) is the click gone and it is revealed as an active feedback on an essentially isometric input (it works at a very low level, even while restarting the computer the click is still there) (it is coupled to the capacitive gesture sensing though—if the system doesn’t sense the presence of the finger, for instance when wearing gloves or a band aid, or using a pencil, the click is not generated). The Steam Controller, a game controller developed by gaming company Valve Corporation and introduced in 2015, has a similar technology inside the (round) trackpad, less used to generate a key click but more for textures and other effects related to the game played. It would be interesting to explore if the Apple Force Click display can also be used to present textures and more contextually relevant variations, through Apple’s CoreHaptics API, so far this is barely applied as such (except for example in the haptic cues when snapping to a grid in Apple’s Keynote or Pages applications).

Preceding Apple’s Taptic Engine is a similar and more sophisticated device, the Haptuator by Hsin-Yun Yao and Vincent Hayward which is able to simulate textures (Yao and Hayward 2010), as applied in research on texture perception (Strohmeier and Hornbæk 2017).

Smartphone and tablet manufacturers have improved the haptic actuators in their devices. Examples are the Apple iPhone (since the 6S in 2015) and the Apple Watch, linked to pressure sensitive input (Apple iPhones have a pressure sensitive home button and touch screen since the model 6S, until the XR model in 2018) which can give a response based on a variable pressure threshold, similar to the Force Touch in the trackpad mentioned above. This Taptic Engine in the iPhone (and trackpads) is a custom-made linear transducer (somewhat like a solenoid), indeed it creates the sensation of a ‘tap’, which is a gentle and useful presentation (Strictly speaking, following the framework presented in Section 1.4, the Taptic Engine is not haptic, it only creates the ‘tap’ addressing the cutaneous sense rather than the kinaesthesia, any actual movement in the interaction is by the person). This combination of pressure sensitivity and ‘taptic’ feedback then led to the omission of mechanical (passive) feedback entirely in the next model (iPhone 7 and up, and in the trackpads), essentially an isometric action with the key click simulated by the taptic engine. While this is designed really well and overall surprisingly convincing, it can feel a bit artificial (which it is, after all), particularly when the timing is not always perfect.

1.6. Vibrotactile Feedback

The most common application of tactual (vibrotactile) feedback in consumer devices is the use of the ‘vibromotor’, a small motor spinning an eccentric weight (Eccentric Rotating Mass, ERM), which creates the vibrotactile cues. Another technology is similar, smaller motors and based on the resonating frequency of the weight, the Linear Resonant Actuator (LRA) (Papetti et al. 2018, p. 259). The range of expression of these technologies is quite limited, it is just one frequency, which can be constant or pulsated. It is kind of the haptic equivalent of the omnipresent buzzer (the cheap piezoelectric beep that make many consumer electronics so annoying) or the blinking red LED, and except for alerting (which it does very well), it is not very appropriate for most of the other applications. Larger versions of this technique are the ‘rumble packs’ in game controllers (often at two different frequencies), and in the late 1990s, several mice have been introduced with larger vibromotors inside. Logitech tried this in several products but the placement of the motor resulted in the vibration being displayed near the palm of the hand, which is a location we usually associate with the party trick of tickling someone’s hand on the inside when shaking hands and an otherwise not very useful sensation. A mouse by AVB Tech from 2001 was a bit more convincing, being able to subtly wiggle the whole device. It was still a bit limited in its application, but it worked on the sound output (just like my palpable pixels work), particularly noticeable in games.

Somewhat similar to a solenoid is the (miniature) relay coil that can generate tactile cues. Such a relay coil has been used in the Tactile Ring developed for gestural articulation in musical contexts (Bongers 1998). The advantage of using a miniature loudspeaker however is that many more pressure levels and frequencies can be generated. This flexibility and potential for rich display of vibrotactile information is the rationale behind the developments described in this paper and earlier research (Bongers 2002). This led to the formulation of our research question of whether vibrotactile feedback could be used to improve articulation, in general applications but with the musical context firmly in mind.

2. Investigating Tactual Articulatory Feedback

As stated before, in the current computer interaction paradigm, articulatory feedback is given visually, while in the real world many movements and manipulative actions are guided by touch, vision, and audition. The rationale behind the work as described in this paper is that it is expected that adding auditory and vibrotactile feedback to the visual articulatory feedback improves the articulation, either in speed or accuracy, and that of these two the vibrotactile feedback will give the greatest benefit as it is the most natural form of feedback in this case.

The work reported here focusses on the physical level of the layered interaction as described before, where the articulation of the gesture takes place. It is at this stage limited to a one degree-of-freedom movement, of a mouse on a flat surface. An essential part of our approach is to leave some freedom for the participants in our experiments to explore.

The research method is described in Section 2.1. Section 2.2 describes the experiment, which is set up in five phases. Phase 1 (Section 2.2.1) is an exploratory phase, where the participants' task is to examine an area on the screen with varying levels of vibrotactile feedback, as if there were invisible palpable pixels. This phase is structured to gain information about the participants' sensitivity to the tactual cues, which is then used in the following phases of the experiment. In the second phase (Section 2.2.2), a menu consisting of a number of items has been enhanced with tactual feedback, and performance of the participants is measured in various conditions. In the third and fourth phase (Sections 2.2.3 and 2.2.4), an on-screen widget, a horizontal slider, has to be manipulated. Feedback conditions are varied in a controlled way, applying different modalities, and performance (reaction times and accuracy) is measured. The last phase (Section 2.2.5) of the experiment consisted of a questionnaire.

Section 2.3 gives the results of the measurements and statistical analyses.

2.1. Method

In the experiments, the participants are given simple tasks, and by measuring the response times and error rates in a controlled situation, differences can be detected. It must be noted that, in the current desktop metaphor paradigm, a translation takes place between hand movement with the mouse (input) and the visual feedback of the pointer (output) on the screen. This translation is convincing due to the way our perceptual systems work. The tactual feedback in the experiments is presented on the hand, where the input takes place, rather than there where the visual feedback is presented. To display the tactual feedback, a custom built vibrotactile element is used. In some of the experiments, auditory articulatory feedback was generated as well, firstly because this often happens in the real world, and secondly to investigate if sound can substitute for tactual feedback. (Mobile phones and touch screens often have such an audible 'key click').

A gesture can be defined as a multiple degree-of-freedom meaningful movement. In order to investigate the effect of the feedback, a restricted gesture was chosen. In its simplest form, the gesture has one degree-of-freedom and a certain development over time. As mentioned earlier, there are in fact several forms of tactual feedback that can occur through the tactual modes as described above, discriminating between cutaneous and proprioceptive, active and passive. When interacting with a computer, cues can be generated by the system (active feedback) while other cues can be the result of real-world elements

(passive feedback). When moving the mouse with the vibrotactile element, all tactual modes are addressed, actively or passively; it is a haptic experience drawing information from the *cutaneous, proprioceptive, and efference copy* (because the participant moves actively), while our system only actively addresses the cutaneous sense. All other feedback is passive, i.e., not generated by the system, but it can play some role as articulatory information.

It was a conscious decision not to incorporate additional feedback on the reaching of the goal. This kind of ‘confirmation feedback’ has often been researched in haptic feedback systems, and has proven to be an effective demonstration of positive effects as discussed in Section 1.1. However, we are interested in situations where the system does not know where the user is aiming for, and are primarily interested in the feedback that supports the articulation of the gesture—not the goal to be reached, as it is not known. For the same reason, we chose not to guide (pull) the users towards the goal, which could be beneficial although following Fitts’s Law (about movement time in relation to distance (‘amplitude’) and target size (Fitts and Posner 1967, p. 112); what happens in such cases is that effectively the target size gets increased (Akamatsu and MacKenzie 1996). Another example is the work of Koert van Mensvoort, who simulated the force-feedback and feedforward visually, which is also very effective (van Mensvoort 2002, 2009).

2.2. Experimental Set-Up

The graphical, object-based programming language MaxMSP was used. MaxMSP is a programme originally developed for musical purposes and has therefore a suitable real time precision (the internal scheduler works at 1 ms) and many built-in features for the generation of sound, images, video clips, and interface widgets. It has been used and proved useful and valid as a tool in several psychometric and HCI-related experiments (Vertegaal 1998, p. 41; Bongers 2002; de Jong and Van der Veer 2022). Experiments can be set up and modified very quickly.

The custom-developed MaxMSP software (‘patches’) for this experiment ran on an Apple PowerBook G3 with OS 8.6 (but also runs under OSX on newer Macs), used in dual-screen mode; the experimenter used the internal screen for monitoring and controlling the experiment, and the participant used an external screen (Philips 19” TFT). The participant carried out the tasks using a standard Logitech mouse (First Pilot Wheel Mouse M-BE58, connected through USB), with the vibrotactile element positioned under the index finger of the participant where the feedback would be expected. For the right-handed participants, this was on the left mouse button; for the left-handed participants, the element could be attached to the right mouse button. The vibrotactile element is a small loudspeaker (\varnothing 20 mm, 0.1 W, 8 Ω) with a mylar cone, partially covered and protected by a nylon ring, as shown in Figure 2. The ring allows the user to avoid pressing on the loudspeaker cone which would make the stimulus disappear.



Figure 2. The vibrotactile element in position, on the mouse used in the experiments reported in this paper (**left** and **middle**), and a later version on a newer mouse and computer (**right**).

The vibrotactile element is covered by the user’s finger so that frequencies in the auditory domain are further dampened. The signals sent to the vibrotactile element are low frequency sounds, making use of the sound generating capabilities of MaxMSP.

Generally, the tactual stimuli, addressing the Pacinian system which is sensitive between 40 and 1000 Hz, are chosen with a low frequency to avoid interference with the audio range. The wavelengths of the (triangular) tactual pulses were 40 and 12 milliseconds resulting in a base frequency of 25 and 83.3 Hz, respectively. The repeat frequency of the pulses would be depending on the speed of the movement made with the mouse, and in practice is limited by the tracking speed. To investigate whether the tactual display really would be inaudible, measurements were performed with a Brüel and Kjær 2203 Sound Level Meter. The base frequency of the tactual display was set to 83.3 Hz and the amplitude levels at 100% and at 60% as used in the experiment. At a distance of 50 cm and an angle of 45 degrees, the level was 36 dB(A) (decibel adjusted for the human ear) when the device was not covered, and 35 dB(A) when the device was covered by the human finger as in the case of the experiments. At an amplitude of 60%, these values were 35 dB(A) and 33 dB(A), respectively. Therefore, in the most realistic situation, i.e., not the highest amplitude and covered by the finger, the level is 33 dB(A), which is close to the general background noise level in the test environment (30 dB(A)) and furthermore covered by the sound of the mouse moving on the table-top surface (38 dB(A)).

Auditory feedback was presented through a small (mono) loudspeaker next to the participant's screen (at significantly higher sound pressure levels than the tactile element).

In total, 35 subjects participated in the trials and pilot studies, mainly first and second year students of Computer Science of the VU Amsterdam university. They carried out a combination of experiments and trials, as described below, using their preferred hand to move the mouse and explore or carry out tasks, and were given a form with open questions for qualitative feedback. In all phases of the experiment, the participants could work at their own pace, self-timing when to click the "next" button to generate the next cue (and the software would log their responses).

Not all subjects participated in all experiments. A total of 4 out of 35 were pilot studies, the 31 main ones all participated in the most important phases of the experiment (2, 3, and 5) and 11 of them participated in the 4th phase as well (which was an extension).

Special attention was given to the participants' posture and movements, as the experiment involved lots of repetitive movements which could contribute to the development of RSI (repetitive strain injury, or WRULD, work related upper limb disorders, in this case particularly carpal tunnel syndrome as can involve the wrist movements—the mouse ratio was set on purpose to encourage larger movements with the whole arm, instead of just wrist rotations). The overall goal of this research is to come up with paradigms and interaction styles that are more varied, precisely to avoid such complaints. At the beginning of each session therefore, participants were instructed and advised on strategies to avoid such problems, and monitored throughout the experiment.

2.2.1. Phase 1: Threshold Levels

The first phase of the experiment was designed to investigate the threshold levels of tactile sensitivity of the participants, in relationship to the vibrotactile element under the circumstances of the test set-up. The program generated in random order a virtual texture, varying the parameters base frequency of the stimulus (25 or 83.3 Hertz), amplitude (30%, 60%, and 100%) and spatial distribution (1, 2, or 4 pixels between stimuli) resulting in 18 different combinations. A case was included with an amplitude of 0, e.g., no stimulus, for control purposes. In total, therefore, there were 24 combinations. The software and experimental set-up was developed by the first author at Cambridge University in 2000 and is described in more detail elsewhere (Bongers 2002).

The participant had to actively explore an area on the screen, a white square of 400×400 pixels, and report if anything could be felt or not. This could be done by selecting the appropriate button on the screen, upon which the next texture was set. Their responses, together with a code corresponding with the specific combination of values of the parameters of the stimuli, were logged by the programme into a file for later analysis. In this phase of the experiment, 26 subjects participated.

This phase also helped to make the participants more familiar with the use of tactual feedback.

2.2.2. Phase 2: Menu Selection

In this part of the experiment, participants had to select an item from a list of numbers (a familiar pop-up menu), matching the random number displayed next to it. There were two conditions, one visual only and the other supported by tactual feedback where every transition between menu items generated a tangible ‘bump’ in addition to the normal visual feedback. The ‘bump’ was a pulse of 12 ms. (83.3 Hz, triangular waveform), generated at each menu item.

The menu contained 20 items, a list of numbers from 100 to 119 as shown in Figure 3. All 20 values (the cues) would occur twice, from a randomly generated table in a fixed order so that the data across conditions could be easily compared—all distances travelled with the mouse were the same in each condition.

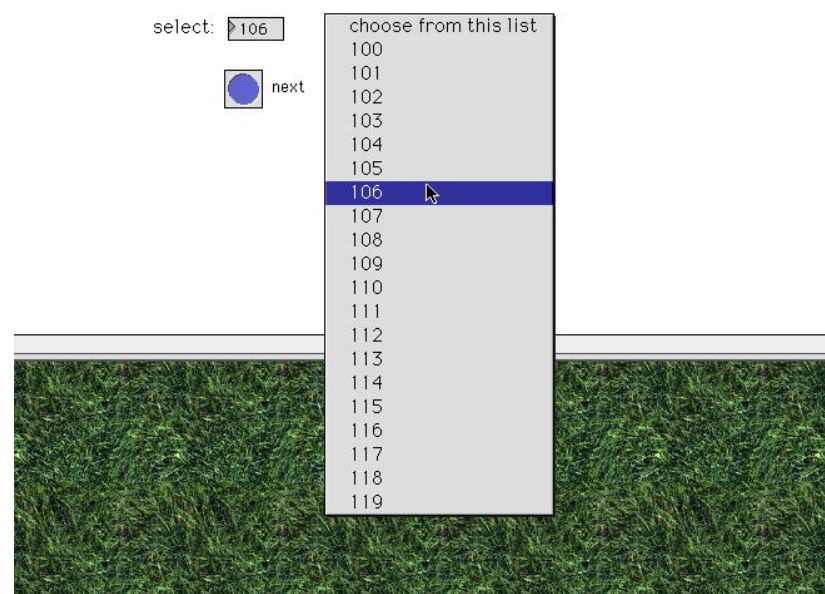


Figure 3. The list of menu items for selection.

Response times and error rates were logged into a file for further analysis. An error would occur if the participant selected the wrong number, the experiment carried on. In total, 30 subjects participated in this phase of the experiment, balanced in order (15–15) for both conditions.

2.2.3. Phase 3: Slider Manipulation

In this experiment, a visual cue was generated by the system, which moved a horizontal slider on the screen to a certain position (the system slider). The participants were instructed to place their slider to the same position as the system slider, as quickly and as accurately as they could. The participant’s slider was shown just below the system slider, each had 120 steps and was about 15 cm wide on the screen.

Articulatory feedback was given in four combinations of three modalities: visual only (V), visual + tactual (VT), visual + auditory (VA), and visual + auditory + tactual (VAT). The sliders were colour coded to represent the combinations of feedback modalities: green (V), brown (VT), violet (VA), and purple (VAT).

The tactual feedback consisted of a pulse for every step of the slider, a triangular wave form with a base frequency of 83.3 Hz. The auditory feedback was generated with the same frequency, making a clicking sound due to the overtones of the triangular wave shape.

These combinations could be presented in 24 different orders, but the most important effect in the context of user interface applications is from Visual Only to Visual and Tactual

and vice versa. It was decided to choose the two that would reveal enough about potential order effects: V-VT-VA-VAT and VAT-VA-VT-V.

Figure 4 shows the experimenter's screen, with several buttons to control the set-up of the experiment, sliders and numbers that monitor system and participant actions, and numeric data indicating the number of trials, and the participant's screen with the cue slider and response slider. All the data was automatically logged into files. The participant's screen had only a system slider (giving the cues), their response slider, and some buttons for advancing the experiment.

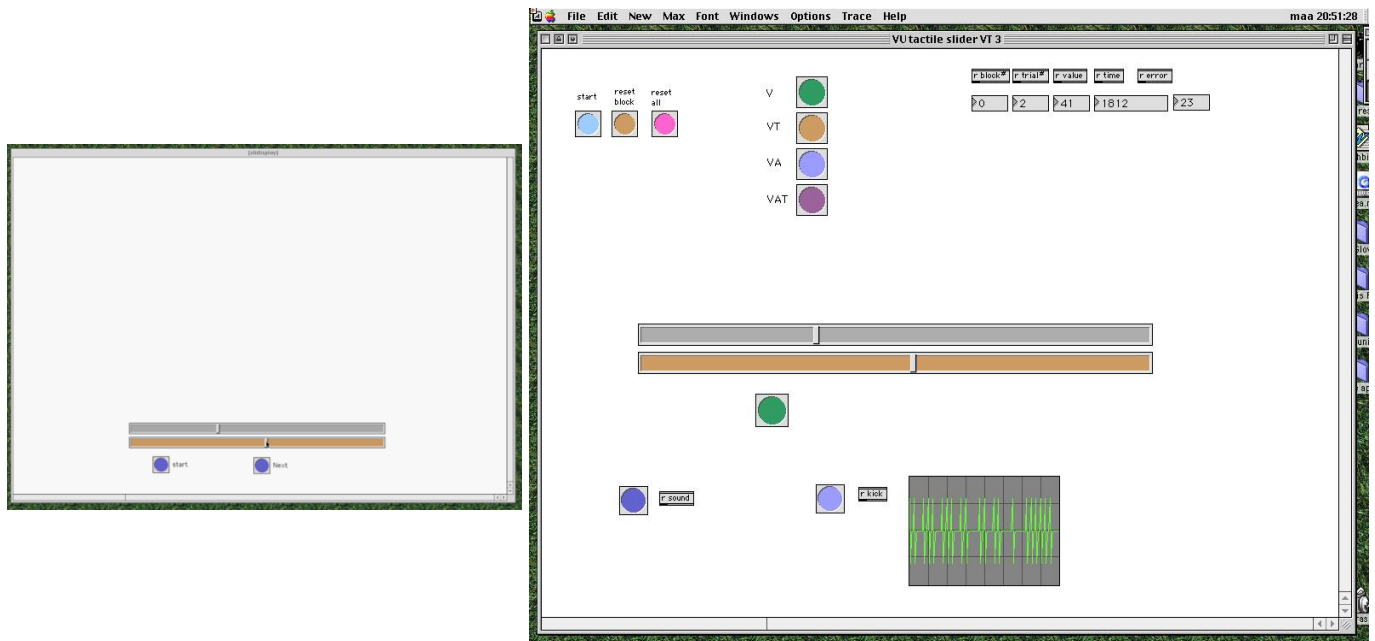


Figure 4. The experimenter's screen (right) of phase 3, the Slider experiment, in the VT condition, showing the controls and monitoring of the experiment and the participant's actions, and the participant's screen (left) (not to scale; the participant's screen was larger than the experimenter's screen).

All 40 cues were presented from a randomly generated, fixed order table of 20 values, every value occurring twice. Values near or at the extreme ends of the slider were avoided, as it was noted during pre-pilot studies that participants developed a strategy of moving to the particular end very fast knowing the system would ignore overshoot. The sliders were 600 pixels wide, mouse ratio was set to slow in order to encourage the participants to really move. Through the use of the fixed values in the table across conditions, it was ensured that in all conditions the same distance would be travelled in the same order, thereby avoiding having to bother with Fitts's law or the mouse-pointer ratio (which is proportional to the movement speed, a feature which cannot be turned off in the Apple Macintosh operating systems).

The cues were presented in blocks for each condition, if the conditions would have been mixed participants would have developed a visual only strategy in order to be independent on the secondary (and tertiary) feedback. Their slider was also colour-coded, each condition had its own colour. A total of 31 subjects participated in this phase, and the orders were balanced as well as possible: in 15 times the V-VAT order and 16 times the VAT-V order.

2.2.4. Phase 4: Step Counting

In Phase 4, the participants had to count the steps cued by the system (Figure 5). It was a variation on the Phase 3 experiment, but made more difficult and challenging for the participants to see if a clearer advantage of the added feedback could be shown. In this experiment, the participants were prompted with a certain number of steps to be taken with

a horizontal slider (the slider was similar to the one in the previous experiment). The range of the cue points on the slider was limited to maximum of 20 steps (from the centre) rather than the full 120 steps in this case, to keep it manageable. The number of steps to be taken with the slider were cued by a message, not by a system slider as in Phase 3, as shown in Figure 5.

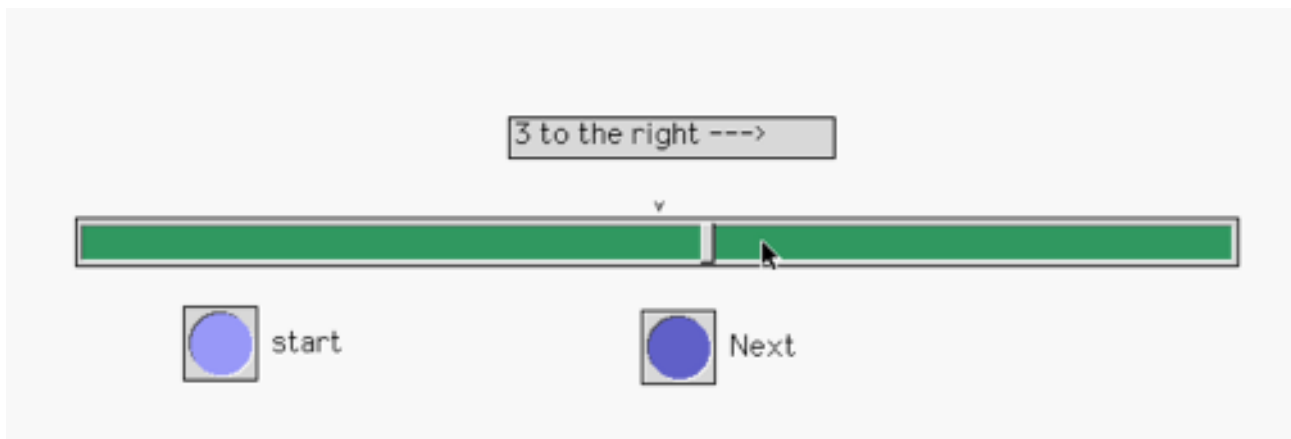


Figure 5. The participant's screen (section), with the cue message and the response slider (V condition). The middle of the slider is marked (and the position the slider resets to after the participants selects the "Next" button to trigger the cue).

The conditions were visual modality only (V) and visual combined with touch modalities (VT). No confirmative feedback was given, the participant would press a button when he or she thought that the goal was reached and then the next cue was generated.

In this phase of the experiment, 11 subjects participated, balanced between the orders in 5 V-VT and 6 VT-V order.

2.2.5. Phase 5: Questionnaire

The last part of the session consisted of a form with questions to obtain qualitative feedback on the chosen modalities, and some personal data such as gender, handedness, and experience with computers and the mouse. This data was acquired for 31 of the participants.

2.3. Results and Analysis

The data from the files compiled by the Max patches logging was assembled in a spreadsheet. The data was then analysed in SPSS.

All phases of the experiment showed a learning effect. This was expected, so the trials were balanced in order (as well as possible) to compensate for this effect.

The errors logged were distinguished in two types: misalignments and mistakes. A mistake is a mishit, for instance when the participant 'drops' the slider too early. A mistake can also appear in the measured response times, some values appeared that were below the physically possible minimum response time, as a result of the participant moving the slider by just clicking briefly in the desired position so that the slider jumps to that position (they were instructed not to use this 'cheating' but occasionally it occurred). Some values were unrealistically high, probably due to some distraction or other error. All these cases contained no real information, and were therefore omitted from the data as mistakes. The misalignment errors were to be analysed as they might reveal information about the level of performance.

In the sections below, the results of each phase are summarised, and the significant results of the analysis are reported.

2.3.1. Phase 1: Threshold Levels

Of the 19 presented textures (averaged over two runs), 12 were recognised correctly (including four of the six non-textures) by all the participants, and a further 10 with more than 90% accuracy (including two non-textures). There were two lower scores (88% and 75%) which corresponded with the two most difficult to perceive conditions, the combination of the lowest amplitude (30%), low frequency (25 Hz), and the lowest spatial distributions of two respectively four pixels.

2.3.2. Phase 2: Menu Selection

Table 1 below shows the total mean response times. All trials were balanced to compensate for learning effects, and the number of trials was such that individual differences were spread out, so that all trials and participants (N in the table) could be grouped. The response times Visual + Tactual condition were slightly higher than for the Visual Only condition, statistical analysis showed that this was not significant. The error rates were not statistically analysed, in both conditions, they were too low to draw any conclusions from.

Table 1. Mean response times and standard deviations for the Menu Selection.

Condition	N	Mean Response Time (ms)	Standard Deviation (ms)
V	1120	1806	624
VT	1120	1842	531

2.3.3. Phase 3: Slider Manipulation

The table below (Table 2) gives an overview of the means and the standard deviations of all response times (for all distances) per condition.

Table 2. Mean response times and standard deviations for the Slider.

Condition	N	Mean Response Time (ms)	Standard Deviation (ms)
V	1239	1436	530
VT	1240	1512	572
VA	1240	1468	595
VAT	1240	1547	593

The response times were normally distributed and not symmetrical in the extremes, so two-tailed *t*-tests were carried out in order to investigate whether these differences are significant. Of the three possible degrees-of-freedom, the interaction between V and VT as well as VA and VAT was analysed, in order to test the hypothesis of the influence of T (tactual feedback added) on the performance. The results are shown in Table 3.

Table 3. The results of the *t*-tests.

Condition	<i>t</i>	Significance
V-VT	−3.435	0.001
VA-VAT	−3.307	0.001

This table shows that response times for the Visual only condition were significantly faster than the Visual + Tactual condition, and that the Visual + Auditory condition was significantly faster than the similar condition with tactual feedback added.

The differences in response times between conditions were further analysed for each of the 20 distances, but no effects were found. The differences in error rates were not significant.

2.3.4. Phase 4: Step Counting

It was observed that this phase of the experiment was the most challenging for the participants, as it indeed was designed to be. It relied more on the participants' cognitive

skills than on the sensori-motor loop alone. The total processing time and the standard deviation are reported the Table 4.

Table 4. Mean response times and standard deviations for the Counter.

Condition	N	Mean Processing Time (ms)	Standard Deviation (ms)
V	200	3569	2281
VT	200	3078	1868

In order to find out if the difference is significant, an analysis of variance (ANOVA) was performed, to discriminate between the learning effect and a possible order effect in the data. The ANOVA systematically controlled for the following variables:

- between subject variable: sequence of conditions (V-VT vs. VT-V);
- within subject variables: 20 different distances.

The results of this analysis are summarised in the table below (Table 5). From this analysis follows that the mean response times for the Visual + Tactile condition are significantly faster than for the Visual Only condition.

Table 5. The results of the comparison for the two conditions of the Counter phase.

Condition	df	Error df	F	Significance
V/VT	1	8	12.432	0.008

The error rates were lower in the VT condition, but not significantly.

2.3.5. Phase 5: Questionnaire

The average age of the 31 participants who filled in the questionnaire was 20 years, the majority was male (26 out of 31) and right-handed (27 out of 31). The left-handed people used the right hand to operate the mouse, which was their own preference (the set-up was designed to be modified for left hand use). They answered a question about their computer and mouse experience with a Likert scale from 0 (no experience) to 5 (lot of experience). The average from this question was 4.4 on this scale, meaning their computer and mousing skills were high.

The qualitative information obtained by the open questions was of course quite varied, the common statements are however categorised and presented in the Table 6.

Table 6. Categorised answers on open questions, out of 31 participants.

	Tactual	Auditory
effective	22	17
effective but potentially annoying	7	3
annoying/irritating	1	5
not effective	1	6

There was a question about whether it was thought that the added feedback was useful. For added sound, 16 (out of 31) participants answered “yes”, 4 answered “no”, and 11 thought it would depend on the context (issues were mentioned such as privacy, blind people, precision). For the added touch feedback, 27 participants thought it was useful, and 4 thought it would depend on the context. In Table 7, the results are summarised in percentages.

Table 7. Categorised answers on question about usefulness.

Added Feedback	Auditory	Tactual
Useful	52%	87%
Not useful	13%	0%
Depending on context	35%	13%

When asked which combination they preferred in general, 3 (out of 31) answered Visual Only, 13 answered Visual+Tactual, 3 answered Visual+Auditory, 7 Visual+Auditory+Tactual, and 4 thought it would depend on the context. One participant did not answer. The results are shown in percentages in Table 8.

Table 8. Preferred combination of articulatory feedback.

Preferred Combination	V	VT	VA	VAT	Depending on Context
	10%	42%	10%	23%	13%

3. Discussion

The results from the Threshold test (Phase 1) show that our participants had no difficulty recognising the stimuli generated, apart from the cases of the really low levels. For the experiments (Phases 2–4), combinations of the stimuli were chosen that were far above these thresholds that were found in Phase 1.

The results from the menu selection (Phase 2) show no significant difference. This task was easy and familiar for the participants, and standard menus are designed with large selection areas allowing for easy selection without needing the help of and added tactual articulatory feedback.

The results from the slider experiment (Phase 3) show that people tend to slow down when the tactual feedback is added. This can be due to the novelty effect of it, as people are not used having their sense of touch actively addressed by their computer. This slowing down effect may be explained by the factor of perceived physical resistance, as has been found in other research (Oakley et al. 2001), perhaps people are used to trying to be more precise and accurate in these circumstances.

When the task reaches a sufficient level of difficulty, the advantage of the added tactual feedback can be shown. This occurred in the step counter set-up in Phase 4 of the experiment, where task completion times were significantly shorter with the added feedback.

The questionnaires show that people generally appreciate the added feedback, and favour the tactual over the auditory. Sound feedback is sometimes perceived as irritating, depending on the context. Note, however, that computer-generated tactual feedback can be unpleasant as well. Not much research has been carried out investigating this rather qualitative aspect and its relationship to the choice of actuator (e.g., a motor or a speaker).

The participants who carried out the experiments were involved in a lecture series on Information Representation at the CS department of the university, and they were therefore very involved in the experiments. It can be argued that this would bias their responses, particularly in the qualitative parts of the session, but it must be stressed that the work here was carried out in the context of human–computer interaction research and not as pure psychometric experiments. Phase 5 of the experiment was therefore more an expert-appraisal. Moreover, the core idea behind the research into tactual feedback as described in this paper was not part of the lectures, so the participants were not biased towards the proposed type of interaction and enhanced articulatory feedback.

The response times between participants varied largely, which is quite fascinating and a potential subject for further investigation. Clearly, people all have their own individual way of moving (what we call the ‘movement fingerprint’) and these results show that even in the simplest gesture, this idiosyncrasy can manifest itself.

The observation that in some cases some participants actually seem to slow down when given the added tactual feedback has a lot to do with the tasks, which were primarily visual. The tactual feedback can be perceived by the user as resistance. This is interesting, as we often experience in real world interactions (particularly with musical instruments) that indeed effort is an important factor in information gathering of our environment, and for articulation. Dancers can control their precise movements while relying more on vision and internal kinaesthetic feedback, but only after many years of training. It is still expected that the greatest benefit of adding tactual feedback to a gesture will be found in a free moving gesture, without any passive feedback as is the case when moving the mouse. This has to be further investigated. Some preliminary experiments we carried out both with lateral movements as well as with rotational movement in free space show promise. The potential benefit of this has already been shown elsewhere, for instance at Sony Labs in the early 2000s (Poupyrev et al. 2002) where a tilting movement of a handheld device did benefit from the added or secondary feedback (here somewhat confusingly labelled as ‘ambient touch’).

4. Conclusions and Further Research

In the research described in this paper, we have investigated the influence of added tactual feedback in the interaction with the computer. It has been found that in simple tasks, the added feedback may actually decrease performance, while at more difficult tasks an increase in effectiveness can occur.

The vibrotactile element, based on a miniature loudspeaker and low-frequency ‘sounds’, proved to be a cost-effective and flexible solution allowing a wide variety of tactile experiences to be generated using sound synthesis techniques.

In the experiments, the added feedback was redundant. In all these cases, the various feedback modalities were used to convey the same information, following the recommendations of the Proximity Compatibility Principle (PCP) (Wickens 1992, pp. 98–101), (Wickens and Carswell 1995). In a multitasking situation, the feedback would perhaps be divided differently according to PCP, but in our experiment, they are tightly coupled. We are not trying to replace visual feedback by tactual feedback, we are adding this feedback to make the interaction more natural, more enjoyable and pleasurable, and possibly more effective. It is expected that in some cases this may lead to a performance benefit. The widget manipulations we chose to investigate are standard user interactions, the participants were very familiar with it, and they could devote all their attention to fulfilling the tasks. Only when the task is made less mundane, such as in the last phase where the participants had to count steps of the slider widget, adding tactile feedback helps to improve the interaction. This may seem irrelevant in the current computer interaction paradigm (including tablets and smartphones), which is entirely based on the visio-motor loop, but our research has a longer-term goal of developing new interaction paradigms based on natural and multimodal interaction. Multitasking is an element of such paradigms, and a future experiment can be developed where the participants have to divide their attention between the task (i.e., placing the slider under various feedback conditions) and a distracting task (see below).

The set-up as described in this paper mainly addresses the Pacinian system, the one out of the four tactual systems of cutaneous sensitivity that has the rapidly adapting and diffuse sensitivity (see Section 1.4), and that can pick up frequencies from 40 Hz to 1000 Hz (Verrillo 1992). Other feedback can be applied as well, conveying different kinds of information.

In the current situation, only pure articulatory feedback was considered, other extensions can be added later. This would also include the results from the ongoing research on virtual textures, resulting in a palette of textures, feels, and other tangible information to design with. This is shown for instance in the recent work of Paul Strohmeier (Strohmeier and Hornbæk 2017).

A logical extension of the set-up is to generate feedback upon reaching the goal, a 'confirmative feedback', which has been proven to produce a strong effect. This could greatly improve performance as it is known that the gesture can be described by dichotomous models of human movement in a pointing task, which imply an initial ballistic movement towards a target followed by finer movements controlled by finer feedback (Mithal 1995), or even three stages: initiative, accelerative and terminal guidance phases of cursor positioning (Phillips and Triggs 2001). These stages were observed in some cases in our experiments, and further added feedback could improve the performance (if there is a known target).

Another development that is thought of, is to have multiple points of feedback addressing more fingers or parts of fingers (Max/MSP works with external hardware to address multiple output channels). The advantage of having multiple points of contact in haptic exploring was already demonstrated by J. J. Gibson with a task of recognising cookie cutters in the paper about active touch (1962), and in many further studies, for instance work at Uppsala University on recognising textures by Gunnar Jansson and Linda Monaci (Jansson and Monaci 2002) (including knäckebröt, Swedish crispbread).

We are also looking for smaller vibrotactile elements, so that more can be put on one finger to facilitate multiple points of contact. This can be used to simulate direction, such as in the effect of stroking velvet. Ideally, the set-up should also incorporate force-feedback, addressing the kinaesthetic sense. Multiple degree-of-freedom input is another expected element of new interaction paradigms.

We are interested to see the effects of added tactual feedback on a free-moving gesture, when the passive feedback (kinaesthetic guidance from the table-top surface in the 2D case) is absent. It is expected that adding active feedback can compensate for this.

In the Interactivation Studio, since 2010, we have applied similar techniques to create articulatory feedback in several projects, using small 'surface transducers', effectively loudspeakers without a cone. One example is a gestural light controller, based on a laser pointer, and a light sensor mounted on the lamp to identify when the lamp was addressed, inside a small handheld unit which contained various sensors and actuators. The research project looked at ambient or peripheral interaction, and participants were tested in a controlled study with varying feedback modalities (visual, auditory, and haptic) to show that they were able to control light levels accurately as a peripheral task when occupied in a focal task simultaneously, significantly better (increased speed and accuracy of the interaction) with the multimodal confirmation and articulatory feedback (Heijboer et al. 2016).

Another example was a musical keyboard retrofitted with these actuators under the keys, to investigate the role of multimodal feedback in a finger setting learning task of playing the keyboard.

This small actuator was until recently sold through SparkFun (COM-10917) with a small amplifier circuit (SparkFun BOB-11044 based on the Texas Instruments TPA2005D1 Class-D amplifier chip), and also available in an improved (more robust) version by Dayton Audio (BCE-1). These are similar to the transducers used in bone conduction headphones, and larger actuators such as the SoundBug that are used to 'shake' a larger surface which can turn a table-top or window into a speaker.

We are also exploring the piezo-electric transducers of Boréas, which have the advantage of being very flat while giving a clear haptic cue using low frequency audio signals, similar to the electromagnetic surface transducers, to simulate textures. Furthermore we are exploring the use of miniature solenoids, for impulse-based feedback.

A new version of the set-up as presented in this paper has been developed, which works in Max8 in OS10.12. This is based on this small surface transducer, which has the advantage that the vibrotactile feedback is much clearer and less dependent on the pressing down force by the finger, as shown in Figure 6.

This is particularly relevant for supporting improved articulation of musical gestures and movements, which some of our current projects are investigating, including the application in the context of audio-editing software.



Figure 6. The current set-up using a small surface transducer for improved tactual feedback.

It is expected that by combining these different actuator technologies, a rich and versatile display can be designed that addresses the various sub-modalities of the tactual perception, leading to a better connection between people and the data and feedback from the computer, in effectivity, but also making the interaction a more pleasurable experience.

Author Contributions: Conceptualization, A.J.B. and G.C.v.d.V.; methodology, A.J.B. and G.C.v.d.V.; software, A.J.B.; validation, A.J.B. and G.C.v.d.V.; formal analysis, A.J.B. and G.C.v.d.V.; investigation, A.J.B.; resources, A.J.B.; data curation, A.J.B.; writing—original draft preparation, A.J.B. and G.C.v.d.V.; writing—review and editing, A.J.B.; visualization, A.J.B.; supervision, G.C.v.d.V. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: The authors wish to thank the student interns Ioana Codoban and Iulia Istrate for their assistance in setting up and carrying out the experiments and their work on the statistical analysis of the data, and Jos Mulder for his help with the sound level measurements.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Akamatsu, Motoyuki, and I. Scott MacKenzie. 1996. Movement characteristics using a mouse with tactile and force feedback. *International Journal of Human-Computer Studies* 45: 483–93. [\[CrossRef\]](#)
- Akamatsu, Motoyuki, and Seguro Sato. 1994. A multi-modal mouse with tactile and force feedback. *International Journal of Human-Computer Studies* 40: 443–53. [\[CrossRef\]](#)
- Akamatsu, Motoyuki, and Sigeru Sato. 1992. Mouse type interface device with tactile and force display. Paper presented at International Conference on Artificial Reality and Telexistence, Tokyo, Japan, July 1–3; pp. 177–82.
- Bolanowski, Stanley J., Jr., George A. Gescheider, Ronald T. Verrillo, and Christin M. Checkowsky. 1988. Four Channels Mediate the Mechanical Aspects of Touch. *Journal of the Acoustic Society of America* 84: 1680–94. [\[CrossRef\]](#)
- Bongers, Albertus J. 1998. Tactual Display of Sound Properties in Electronic Musical Instruments. *Displays Journal* 18: 129–33. [\[CrossRef\]](#)
- Bongers, Albertus J. 2002. Palpable Pixels, A Method for the Development of Virtual Textures. In *Touch, Blindness and Neuroscience*. Edited by Soledad Ballesteros and Morton A. Heller. Madrid: UNED Press.
- Bongers, Albertus J. 2006. Interactivation—Towards an Ecology of People, Our Technological Environment, and the Arts. Ph.D. thesis, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands, July.
- Bongers, Albertus J. 2007. Electronic Musical Instruments: Experiences of a New Luthier. *Leonardo Music Journal* 17: 9–16. [\[CrossRef\]](#)
- Bongers, Albertus J., and Gerrit C. van der Veer. 2007. Towards a Multimodal Interaction Space, categorisation and applications. Special issue on Movement-Based Interaction. *Journal of Personal and Ubiquitous Computing* 11: 609–19. [\[CrossRef\]](#)
- Bongers, Albertus J., J. Hubertus Eggen, David V. Keyson, and Steffen C. Pauws. 1998. Multimodal Interaction Styles. *HCI Letters Journal* 1: 3–5.

- Chafe, Chris. 1993. Tactile Audio Feedback. Paper presented at ICMC International Computer Music Conference, Tokyo, Japan, September 10–15; pp. 76–79.
- Chafe, Chris, and Sile O'Modhrain. 1996. Musical Muscle Memory and the Haptic Display of Performance Nuance. Paper presented at ICMC International Computer Music Conference, Hong Kong, China, August 19–24; pp. 428–31.
- Clark, Francis J., and Kenneth W. Horch. 1986. Kinesthesia. In *Handbook of Perception and Human Performance*. Volume 1—Sensory Processes and Perception. New York: John Wiley and Sons, chp. 13.
- Dennerlein, Jack Tigh, David M. Martin, and Christopher J. Hasser. 2000. Force-feedback improves performance for steering and combined steering-targeting tasks. Paper presented at CHI 2000 Conference on Human Factors in Computing Systems, The Hague, The Netherlands, April 1–6; pp. 423–29.
- Dix, Alan, Janey Finlay, Gregory Abowd, and Russell Beale. 1993. *Human-Computer Interaction*. New York: Prentice Hall.
- Eggen, J. Hubertus, Reinder Haakma, and Joyce H. D. M. Westerink. 1996. Layered Protocols: Hands-on experience. *International Journal of Human-Computer Studies* 44: 45–72. [[CrossRef](#)]
- Engel, Frits L., Peter Goossens, and Reinder Haakma. 1994. Improved efficiency through I- and E-feedback: A trackball with contextual force feedback. *International Journal of Human-Computer Studies* 41: 949–74. [[CrossRef](#)]
- Erickson, Thomas. 1995. Coherence and Portrayal in Human-Computer Interface Design. In *Dialogue and Instruction*. Edited by Robbert-Jan Beun, Michael Baker and Miriam Reiner. Heidelberg: Springer, pp. 302–20.
- Fitts, Paul M., and Michael I. Posner. 1967. *Human Performance*. London: Prentice-Hall.
- Fukumoto, Masaaki, and Toshiaki Sugimura. 2001. ActiveClick: Tactile Feedback for Touch Panels. In *CHI'2001, Extended Abstracts*. New York: ACM, pp. 121–22.
- Gaver, William W. 1989. The SonicFinder, a prototype interface that uses auditory icons. *Human Computer Interaction* 4: 67–94. [[CrossRef](#)]
- Gibson, James J. 1962. Observations on active touch. *Psychological Review* 69: 477–91. [[CrossRef](#)] [[PubMed](#)]
- Göbel, Matthias, Holger Luczack, Johannes Springer, Volkmar Hedicke, and Matthias Rötting. 1995. Tactile Feedback Applied to Computer Mice. *International Journal of Human-Computer Interaction* 7: 1–24. [[CrossRef](#)]
- Hayward, Vincent. 2018. A Brief Overview of the Human Somatosensory System. In *Musical Haptics*. Berlin and Heidelberg: Springer, chp. 3. pp. 29–48.
- Hardwick, Andrew, Jim Rush, Stephen Furner, and John Seton. 1996. Feeling it as well as seeing it—haptic display within gestural HCI for multimedia telematic services. In *Progress in Gestural Interaction*. Berlin and Heidelberg: Springer, pp. 105–16.
- Hardwick, Andrew, Stephen Furner, and John Rush. 1998. Tactile access for blind people to virtual reality on the world wide web. *Displays Journal* 18: 153–61. [[CrossRef](#)]
- Heijboer, Marigo, Elise A. W. H. van den Hoven, Albertus J. Bongers, and Saskia Bakker. 2016. Facilitating Peripheral Interaction—design and evaluation of peripheral interaction for a gesture-based lighting control with multimodal feedback. *Personal and Ubiquitous Computing* 20: 1–22. [[CrossRef](#)]
- Hix, Deborah, and H. Rex Hartson. 1993. *Developing User Interfaces*. New York: John Wiley & Sons, Inc.
- Hwang, Faustina, Simon Keates, Patrick Langdon, and P. John Clarkson. 2003. Multiple haptic targets for motion-impaired users. Paper presented at CHI 2003 Conference on Human Factors in Computing Systems, Fort Lauderdale, FL, USA, April 5–10; pp. 41–48.
- Jansson, Gunnar, and Linda Monaci. 2002. Haptic Identification of Objects with Different Number of Fingers. In *Touch, Blindness and Neuroscience*. Edited by Soledad Ballesteros and Morton A. Heller. Madrid: UNED Press, pp. 209–19.
- de Jong, Staas, and Gerrit C. Van der Veer. 2022. Computational Techniques Enabling the Perception of Virtual Images Exclusive to the Retinal Afterimage. *Journal of Big Data and Cognitive Computation* 6: 97. [[CrossRef](#)]
- Katz, David T. 1989. *The World of Touch*. Translated and Edited version of *Der Aufbau der Tastwelt* [1925]. Edited by L. E. Krueger. Hillsdale: Lawrence Erlbaum Associates.
- Keyson, David V. 1996. Touch in User Interface Navigation. Ph.D. thesis, Eindhoven University of Technology, Eindhoven, The Netherlands.
- Keyson, David V., and Adrianus J. M. Houtsma. 1995. Directional sensitivity to a tactile point stimulus moving across the fingerpad. *Perception & Psychophysics* 57: 738–44.
- Keyson, David V., and Leon van Stuijvenberg. 1997. TacTool v2.0: An object-based multimodal interface design platform. Paper presented at HCI International Conference, San Francisco, CA, USA, August 24–29; pp. 311–14.
- Kruglanski, Orit. 2000. As much as you love me. In *Cyberarts International Prix Arts Electronica*. Edited by Hannes Leopoldseder and Christine Schöpf. Berlin and Heidelberg: Springer, pp. 96–97.
- Linden, David J. 2015. *Touch—The Science of Hand, Heart, and Mind*. London: Penguin Books.
- Loomis, Jack M., and Susan J. Leederman. 1986. Tactual Perception. In *Handbook of Perception and Human Performance*. Volume 2: Cognitive Processes and Performance. New York: John Wiley and Sons, chp. 31.
- McCullough, Malcolm. 1996. *Abstracting Craft, The Practised Digital Hand*. Cambridge: MIT Press.
- van Mensvoort, Koert M. 2002. What you see is what you feel—exploiting the dominance of the visual over the haptic domain to simulate force-feedback with cursor displacements. Paper presented at 4th conference on Designing Interactive Systems, London, UK, June 25–28; pp. 345–48.
- van Mensvoort, Koert M. 2009. What You See Is What You Feel. Ph.D. thesis, Eindhoven University of Technology, Eindhoven, The Netherlands.

- Minsky, Margaret, Ouh-young Ming, Oliver Steele, and Frederick P. Brooks. 1990. Feeling and seeing: Issues in force display. *Computer Graphics* 24: 235–43. [[CrossRef](#)]
- Mithal, Anant Kartik. 1995. Using psychomotor models of movement in the analysis and design of computer pointing devices. Paper presented at CHI, Denver, CO, USA, May 7–11; pp. 65–66.
- Münch, Stefan, and Rüdiger Dillmann. 1997. Haptic output in multimodal user interfaces. Paper presented at International Conference on Intelligent User Interfaces, Orlando, FL, USA, January 6–9; pp. 105–12.
- Nielsen, Jakob. 1986. A virtual protocol model for computer-human interaction. *International Journal of Man-Machine Studies* 24: 301–12. [[CrossRef](#)]
- Oakley, Ian, Alison Adams, Stephen A. Brewster, and Philip D. Gray. 2002. Guidelines for the design of haptic widgets. Paper presented at HCI 2002 Conference, London, UK, September 2–6.
- Oakley, Ian, Marilyn Rose McGee, Stephen A. Brewster, and Philip D. Gray. 2000. Putting the ‘feel’ into ‘look and feel’. Paper presented at CHI 2000 Conference on Human Factors in Computing Systems, The Hague, The Netherlands, April 1–6; pp. 415–22.
- Oakley, Ian, Stephen A. Brewster, and Philip D. Gray. 2001. Solving multi-target haptic problems in menu interaction. Paper presented at CHI 2001, Conference on Human Factors in Computing Systems, Seattle, WA, USA, March 31–April 5; pp. 357–58.
- Papetti, Stefano, and Charalampos Saitis, eds. 2018. *Musical Haptics*. Berlin and Heidelberg: Springer.
- Papetti, Stefano, Martin Fröhlich, Federico Fontana, Sébastien Schiesser, and Federico Avanzini. 2018. Implementation and Characterization of Vibrotactile Interfaces. In *Musical Haptics*. Berlin and Heidelberg: Springer, chp. 13. pp. 257–82.
- Phillips, James G., and Thomas J. Triggs. 2001. Characteristics of Cursor Trajectories Controlled by the Computer Mouse. *Ergonomics* 44: 527–36. [[CrossRef](#)]
- Poupyrev, Ivan, Shigeaki Maruyama, and Jun Rekimoto. 2002. Ambient touch: Designing tactile interfaces for handheld devices. Paper presented at User Interface and Software Technology (UIST) Conference, Paris, France, October 27–30.
- Schomaker, Lambert, Stefan Münch, and Klaus Hartung, eds. 1995. *A Taxonomy of Multimodal Interaction in the Human Information Processing System*. Report of the ESPRIT Project 8579: MIAMI. Nijmegen: Radboud University Nijmegen.
- Sherrick, Carl E., and Roger W. Cholewiak. 1986. Cutaneous Sensitivity. In *Handbook of Perception and Human Performance*. Volume 1—Sensory Processes and Perception. New York: John Wiley and Sons, chp. 12.
- Strohmeier, Paul, and Kasper Hornbæk. 2017. Generating Textures with a Vibrotactile Actuator. Paper presented at CHI 2017, Conference on Human Factors in Computing Systems, Denver, CO, USA, May 6–11.
- Taylor, Martin M. 1988. Layered protocol for computer-human dialogue. I: Principles. *International Journal of Man-Machine Studies* 28: 175–218. [[CrossRef](#)]
- Verrillo, Ronald T. 1992. Vibration sensing in humans. *Music Perception* 9: 281–302. [[CrossRef](#)]
- Vertegaal, Roel. 1998. Look Who’s Talking to Whom. Doctoral thesis, Twente University, Twente, The Netherlands.
- Wensveen, Stephan A. G., Tom Djajadiningrat, and Cees J. Overbeeke. 2004. Interaction Frogger: A Design Framework to Couple Action and Function through Feedback and Feedforward. Paper presented at Designing Interactive Systems DIS’04 Conference, Cambridge, MA, USA, August 1–4.
- Wickens, Christopher D. 1992. *Engineering Psychology and Human Performance*, 2nd ed. New York: Harper Collins.
- Wickens, Christopher D., and C. Melody Carswell. 1995. The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors* 37: 473–94. [[CrossRef](#)]
- Yao, Hsin-Yun, and Vincent Hayward. 2010. Design and analysis of a recoil-type vibrotactile transducer. *Journal of the Acoustic Society of America* 128: 619–27. [[CrossRef](#)] [[PubMed](#)]
- Young, Emma. 2021. *Super Senses—The Science of Your 32 Senses and How to Use Them*. London: John Murray.
- Ziat, Mounia. 2023. Haptics for Human-Computer Interaction: From the Skin to the Brain. *Foundations and Trends in Human-Computer Interaction* 17: 1–194. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.