The Ethics of Digital Touch

Nicholas A. Barrow, Patrick N. Haggard

***Abstract*—This paper aims to outline the foundations for an ethics of digital touch. Digital touch refers to hardware and software technologies, often collectively referred to as 'haptics', that provide somatic sensations including touch and kinaesthesis, either as a stand-alone interface to users, or as part of a wider immersive experience. Digital touch has particular promise in application areas such as communication, affective computing, medicine, and education. However, as with all emerging technologies, potential value needs to be considered against potential risk. We therefore identify some areas where digital touch raises ethical concerns, and we identify *why* these concerns arise, based on the distinctive physiological and functional properties of the human somatosensory system. Most scientific research in digital touch has focused on user interaction with external objects (active touch). However, the most pressing ethical concerns with digital touch technologies arise when users are being passively touched. Our analysis identifies several important questions about control, transparency, and epistemic procedure in digital touch scenarios. First, human somatosensation is "always on", and many digital touch technologies take advantage of this (e.g., alerting systems). As a result, digital touch technologies can undermine individuals' sensory autonomy (i.e., the right to choose what sensations one experiences). Second, users may reasonably want to know who or what is touching them, and for what purpose. Consent for digital touch will therefore need to be carefully and transparently transacted. Third, because touch gives us a special, direct experience of interacting with our physical environment, digital touch technologies that manipulate this interaction could potentially provide a major epistemic challenge, by changing a user's basic understanding of reality and their relation to it. Informed by this discussion, we conclude by suggesting a basis for an ethical design framework for digital touch systems.**

***Index Terms*— haptics, ethics, haptic technology, digital touch, touch**

I. INTRODUCTION

Recent progress in haptic technology (HT) has given rise to the concept of digital touch. We broadly define digital touch as a process in which a technological device produces specific sensations and experiences in the user by stimulating the body's surface. We note that 'digital' is a term of convenience in this context since the stimulation itself is typically delivered by analogue hardware. Haptic devices can produce vibrotactile, kinaesthetic, or thermal sensations [1] and can be classified into two types: those that require direct contact (such as joysticks [2]) and those that do not, utilising methods such as ultrasound to stimulate tactile sensation in midair [3].

Examples of existing digital touch technology range from 'wearables' such as full-body exoskeletons and remote-operation bracelets/smartwatches, to games console controllers, helicopter joysticks, and steering wheels (see Fig 1). Early haptics interfaces often used touch to convey relevant information to guide behaviour without distracting a user's overt attention. For example, digital touch systems are commonly used in the automotive industry to warn drivers of lane departure, allowing them to remain visually focused on the road [4]. In contrast, more recent systems have much wider ambitions to reproduce, enhance, and extend through technology the capacities and functions of touch (and other bodily sensations) in everyday life. For instance, modern haptics aims to reproduce the distinctive qualitative feeling of touching specific objects, such as their texture [5], warmth [6], or stiffness [7]. One particular interest (and challenge) is the technological simulation and mediation of human-human tactile interaction and communication. The possibilities of digital handshakes [8] and programmable long-distance kissing [9], for example, have attracted both engineering and media interest but such devices remain technically difficult and subjectively unconvincing.

Most modern haptic systems aim at 'realism'. Yet, what this means is rarely discussed. The working assumption is that haptic technologies aim to accurately reproduce everyday experiences of tactile sensation. For example, in many virtual reality (VR) and haptic display set-ups the user can both see and feel an object they can haptically interact with [10]. The tactile experience should then be 'as if' the user is really touching the object that they see and indistinguishable from the tactile experiences that would be provided by the real-world objects being simulated. The assumed goal, therefore, is that users should have an experience 'as if' they were actually handling a piece of fabric, or squeezing a loved one's hand, or holding a ball, and so on. This goal still eludes haptic designers.

The stimulator hardware used to generate these experiences are generally very physically different from the actual piece of fabric, the actual hand of a loved one, or the actual squishiness of ball that is being simulated. Yet, the user's sensory experiences should be qualitatively equivalent. The digital control of the stimulator/actuator therefore needs to be controlled in a way that reproduces the physical events of typical interactions [11]. Further, the user will often *know* that the digital touch sensations they are experiencing are artificial. How will their knowledge of what *really* underlies their experiences change the experiences themselves? How will the knowledge that they are actually touching a stimulator that is simulating a soft fabric change their experience of the softness of the fabric? If the experience is unaffected by knowing that it is not what it seems to be, it is said to be "cognitively impenetrable" [12]. Interestingly, many tactile illusions seem to be cognitive impenetrable, in the sense that the tactile experience can be convincing, or at least engaging, even when the user knows it is an illusion [13].

A. Realism and immersion: ethical implications of belief engineering

We distinguish three levels of realism for digital touch and assume that they are hierarchically connected with each other. A perfect haptic system will facilitate all three levels (see Table I).

First, all digital touch involves real physical stimulation of touch receptors. Haptic technologies produce physical events – such as vibrations or ultrasound waves – that induce sensations by activating somatosensory receptors in the skin or underlying body tissue. Importantly, on this definition, auditory and visual stimuli that produce illusions of touch without any physical effect on skin receptors do not qualify as digital touch – they are faked touch, even if the fake is convincing.

Second, we can ask whether a tactile stimulation produces experiences that are sufficiently familiar to their 'real' counterparts to provide the user with a functional, convincing experience of touch. Digital touch systems aim to create experiences that are qualitatively similar to established tactile experiences. For instance, a haptic display setup that aims to replicate the sensation of handling a cotton fabric, should reproduce the particular softness, temperature and weight of cotton. The more the user perceives these particular qualities, the more 'realistic' the tactile experience will be, and the more convincing the interaction will be.

Third, a digital touch system might reproduce some of the qualitative experiences of touching a particular object, but it may still fail to convince a user they are *really* feeling *that particular object*. The "as if" element of the tactile illusion can be only partial. For example, a haptic system might give us a feeling somewhat like touching a loved one's hand. However, the experience might not be convincing, either for tactile reasons (e.g., the digital touch stimulation of their skin might be convincingly soft, but the surface temperature might be implausible) or for other reasons (when we look, we might not see a hand like our loved one's). Interestingly, partial tactile illusions can be interesting and absorbing, even when other sources of belief tell us the tactile experience is not veridical – this fact motivates attempts to improve digital touch technologies. In the context of VR technologies, "immersion" and "presence" are used in evaluating perceived realism, and digital touch potentially facilitates immersion [14].

A perfect haptic display would produce like-real tactile experiences that successfully induce beliefs about the objects, surfaces, and people that one touches. This seems to be the aim of many current haptic technologies. In one sense, these beliefs will be false: we won't *really* be feeling *that* particular object. In fact, in some cases, the beliefs we derive from HT might not pick out anything in the *real* world at all. For example, robotic manipulation systems are often used to give the feeling of one's hand banging into a wall, or following along the side of a wall, yet in reality there is no wall. The challenge for haptic designers and digital touch interface engineers is to provide experiences that induce these (false) beliefs as effectively as possible. But in another sense, such experiences can provide true beliefs about *what* it is like to feel something. For example, a haptic system could provide an accurate experience of what it is like to touch my loved one's hand, even though any associated belief that the hand is really there would be false. From this experience, we would derive a true qualitative belief about what it is like to touch my loved one's hand. Sometimes the experience will induce false beliefs in *both* senses. For instance, a HT could simulate an experience of (a) touching an object that isn't actually present and (b) the object has atypical phenomenal qualities that the real object wouldn't have. To give an example, a HT could potentially simulate the size, texture, and compliance of an ice cube, but couple these with the sensation of heat rather than cold.

Recasting digital touch engineering as the inculcation of false beliefs shows why an ethics of digital touch is important: false belief scenarios are sometimes justifiable or valuable, but sometimes raise ethical concerns.

In summary, HT's potentially raise two distinct sets of ethical concerns, corresponding to two distinct cognitive levels of somatosensory processing. First, HT's can potentially cause sensory *experiences* that are aversive or unpleasant, notably pain. This potentially raises ethical concerns related to affect alone, whether or not these experiences align with expectation or challenge our beliefs about reality. What we have called 'perfect' haptic systems, those that *can* update our beliefs, raise *additional* ethical concerns related to belief, deceit, and manipulation.

B. Ethical particularities of touch

To our knowledge, no previous work has focussed specifically on developing an applied ethics for digital touch. This may be surprising given the burgeoning research and commercialization efforts in this area. Therefore, we outline here an intellectual and scientific foundation for the ethics of both natural and digital touch. Our remarks are directed at those working in research, development, implementation, and use of digital touch technologies, but also in technology ethics generally. Importantly, an ethics of digital touch must be informed by scientific knowledge regarding how the sense of touch works, and by clear theories of touch processing. For example, ethically relevant attributes and content should be considered separately for the stimulus level, experiential level, and belief level described above. Similarly, we argue that the ethical concerns for digital touch differ strongly between cases where the system focusses on the user's active, exploratory tactile interactions with the external environment, and situations where the user is passively touched.

We proceed as follows. First, we outline some specific and relevant facts about the biological functions of touch and somatosensory systems in humans, focussing on their specific role in affect. Second, we show how these specific considerations raise a unique set of ethical problems. Third, we consider how digital touch systems change our everyday understanding of touch and consider the ethical issues these changes will bring. Finally, we illustrate how these insights can be applied within a practical ethical design framework.

II. Key Features of Touch

The sense of touch is thought to play both a functional/perceptual and affective role. Touch gives us structural information about the objects we are touching, such as their size, shape, and compliance. Touch also conveys affect, though experiences of warmth, pain, and social contact. Four distinctive properties of functional touch, which differ importantly from other sensory modalities (vision, audition, and smell), are particularly relevant for the ethics of digital touch. To be specific, touch is physically direct; touch is epistemically private; touch is self-related; and touch can be emotionally salient. All these features raise special ethical issues for touch in general, and digital touch in particular.

A. Directness

Touch provides us with special information about the external world. Classically, it complements other exteroceptive senses such as vision. For example, we may use touch to confirm that what we are seeing is *actually* there: that there *is* an external world, rather than a chimera, a hologram, or a hallucination. Touch also confirms the location of our body relative to other objects within that external world: our tactile experience is restricted to those objects in the world that abut our own location.

In fact, touch has two unique features as a spatial sense. First, it only signals objects that *directly* contact the body, offering no *direct* information about objects at a distance. We can only have tactile perceptions of objects that mechanically interact with our skin in certain ways, either directly or indirectly. Most objects at a distance induce no touch experience at all, even though I may be able to see, smell, or hear them. Information about remote objects can be extracted through tool use, but only when mediated by *direct* contact between the body and the tool. This type of tactile perception has been referred to as distal [15]. In such cases, touch signals provide information directly about the object that contacts the body, such as a walking stick, but also provide mediated information about the remote object touched by the walking stick. But, even in these cases, tactual perception is clearly reliant on *some* object coming into direct contact with the body.

Second, the sense of touch has a low 'bandwidth': people have a limited ability to experience multiple touches simultaneously. Canonical tactile experience typically involves touching a single object [16]–[18]. Taken together, these features mean that our tactile experience in everyday interactions is typically focal. By combining tactile and kinaesthetic inputs, haptic experiences can be considered spatial (one can tell apart the different keys in one's pocket by touch alone). Some argue that touch alone (i.e., skin sensation without proprioceptive body configuration information) lacks a spatial perceptual "field" of the kind that features in visual experience[16]. Others, however, have suggested it does [19], [20].

The directness of touch gives rise to two distinctive properties of everyday manual touch experience. First, tactile experience gives us access to properties of objects unavailable to us from other senses such as their mass, compliance, and temperature. Second, and perhaps for these preceding reasons, touch often seems to be our ultimate arbiter of reality, providing highly authoritative and compelling evidence about physical existence. Indeed, despite the existence of many tactual illusions, appeals to the epistemic authority of touch are notorious in philosophy. Descartes – famous for his sensory scepticism – declared that "of all our senses, touch is the one considered least deceptive and the most secure" [21, p.5]. Johnson, too, in refutation of Berkely's immaterialism, demonstrated the reality of physical objects by kicking a stone. Such thinking is intuitive: if we are unsure whether our visual percept of a tree is genuine or hallucinatory, we might reach out to try and touch it, utilising tactile perception as a test. Indeed, recent evidence shows that people report high confidence with respect to tactile sensations when visual stimuli are ambiguous [23]. The status of touch as our ultimate arbiter of physical reality applies as much to our own bodies as to external bodies. Indeed, the capacity for self-touch plays an important role in generating and maintaining a sense of self that is linked to the body [24]. Further, from what we can and cannot touch, we are able to derive a sense of our bodily boundaries and limitations in relation to the wider world [25].

B. Epistemic privacy

Touch provides us with special information about the external world and many digital touch systems create or enhance this capacity by focussing on exteroception. However, the same technologies can also induce experiences which are linked to the body and the self, rather than to the specific external objects that the body touches. For example, a haptic display could cause a diffuse, caressing feeling of being warm, instead of, or as an alternative to, the impression of touching a warm object.

In general, touch is also widely recognised to be both an exteroceptive and an interoceptive, or at least proprioceptive, sense [26]. Other examples of such 'internal', somatic sensations include itches, pains, and feelings of warmth and cold. This internal, or somatic aspect of tactile sensation is particularly important for the affective qualities of touch and has a particular epistemic and ethical significance. Some somatic sensations seem to possess only the internal aspect (in the sense that they are not linked to any obvious external object of perception). There seems to be nothing more to these sensations than the subjective experience of them. To have an itch is simply to experience that itch. Sometimes there is an external itchy object, like the label in your shirt, but sometimes there is no such object – for example, the itch on my arm caused by a mosquito bite is typically felt long after the mosquito has gone. Examples like this contrast with the deep intuition that exteroceptive perceptions are related to external objects.

Philosophers often debate whether and how we can know the subjective experiences of others [27]– [30]. For somatic sensations dominated by the internal aspect of experience, epistemic privacy becomes a particularly acute problem: how can *I* know what *your* itch feels like? Suppose that we see someone removing their hand from a hot stove and crying out in pain. Although we can *see* and *hear* that they are experiencing pain – and might even empathise with this pain [31] – we do not experience the pain itself. We do not, and cannot, perceive someone else's pain (or itch, ache, patch of warmth) in the same way we can our own. Often, we may understand the external aspect of a person's touch experience – for example, we can tell that the stove is hot because we can see the flame of the gas burner and hear the hot oil sizzling in a pan. We then use our existing knowledge to infer that what they are experiencing is similar to our own previous experience of cooking burns. This allows us to represent and perhaps understand their current experience, but we can have no direct phenomenal access to it. No matter how many years of experience a doctor may have, for example, they will remain reliant on *you* reporting where *you* feel pain, how intense it is, and how long *you* have felt it. The exteroceptive, object-referencing aspects of touch can therefore provide *some* guidance as to what the user of a haptic device may feel, but the internal aspect of somatosensory phenomenology is available only for the perceiver. The person who experiences a somatic sensation has epistemic authority with respect to their experience of that sensation.

C. Self-relatedness

Somatic sensations have a specific feeling of 'mineness'. Each sensation of touch is imbued with the feeling that I am actively touching something, or that something is passively touching me. Because the skin is the body boundary, tactile sensations seem to have a specific role in creating the experience of a bodily self, continuous through time, present in each sensory interaction, and identifiable across different multisensory channels. The implicit presence of a bodily self in each individual sensory experience is often thought to be stronger for somatic senses than for exteroceptive senses [32]. Touching a surface with my hand seems to have a stronger involvement of an "I" or a "me" than, for example, hearing a sound [33].

The term "body ownership" is sometimes used to refer to this aspect of self-consciousness and to discuss its special link to somatic sensation. The nature and origin of bodily self-consciousness is a major topic of current research, with at least three discernible research traditions. First, the distinctive awareness of one's own body may be a by-product of voluntary agency – one can directly control and move one's own body at will, but one cannot control other objects in the same direct way [34]. A second theory links the bodily self to an immediately-surrounding defensive peripersonal space, and thus ultimately to responses to threat [25]. A third view links bodily awareness to visual self-recognition [35]. It remains controversial whether any of these theories sufficiently explain the 'mineness' that accompanies somatic sensations.

D. Affective touch

The affective and emotional significance of touch is evident from everyday life, and forms an essential part of our biological existence through activities such as parental care [36], communicating emotion [37], and social interaction [38]. Current understanding of the emotional element of touch was significantly advanced by the discovery that a specific class of peripheral afferents – C-tactile fibres – preferentially respond to the slow stroking movements typical of affectionate touch [39], [40]. Experientially, the sensations caused by C-tactile afferents are consistently described as pleasant, agreeable, and relaxing [41] but there is also some evidence they may be responsible for evoking painful sensations [40], [42]. Further, these afferents are found primarily in hairy skin rather than glabrous skin [43] and project preferentially to the insular cortex rather than the primary somatosensory cortex [44]. Whereas functional, informative touch provides a channel optimized for perceptual pickup during manual interactions with the external environment, the affective touch system appears optimized as a channel for emotion regulation, interpersonal relations, and social bonding.

A distinctive feature of affective touch in humans (and other social animals) is its strong involvement in behaviour, experience, and, consequently, our cultural and social structures [45]. Interpersonal touch is considered appropriate and positive in some contexts, but highly inappropriate in others [46]. This fact may be related to the importance of monitoring and controlling possible threats, since interpersonal affective touch necessarily involves admitting another agent into one's defensive peripersonal space [25]. Interestingly, even affective touch directed to one's own body, in the absence of another person, shows strong cultural influences and modifications, with many cultural norms regulating behaviours such as thumb-sucking.

III. The Ethics of Touch

Specific discussion of the ethics of touch seems almost absent from the academic literature, although some applied ethics literature exists regarding touch in therapeutic contexts. The preceding review of the psychological and cognitive mechanisms of touch points to a number of ways in which touch sensations and experiences may be ethically significant. In this section, we expand on this to outline the broad principles for an ethics of touch. The ethical concerns surrounding touch are more strongly related to being touched, than to touching. Indeed, haptically exploring my environment and actively touching the objects within it seems to raise no particular ethical problems over and above looking around me. However, if the objects I touch include another person's body, the ethical concerns around being touched clearly apply.

A. Sensory Autonomy

A first and central concern for an ethics of touch is care for and autonomy over one's own body. Autonomy refers to the capacity for self-determination: to be in control over particular facets of our lives and to live according to our own reasons and motives, free from external manipulation [47]. Concerning the body in particular, this control is most acutely reflected by the general thought that we should have control over what (and particularly who) can touch us, how they touch us, where they touch us, and when they touch us. Because the skin is a sensory organ enclosing the entire body, it matters what comes into contact with our skin. Losing the ability to control skin inputs is a threat to our bodily autonomy and thus to the self. This is further reflected by the fact we often accord a special ethical status to the body. Consider, for instance, how punishments that restrict our social and environmental autonomy – such as imprisonment – are generally accepted, while those that infringe upon the body – branding, torture, and so on – are considered unacceptable [48].

What is the special link between somatic sensation and bodily autonomy? Indirectly, somatic sensation clearly informs the relation between the body and the sense of agency and also informs the relation between the body and personal survival. Concerning the former, somatic sensation facilitates the feeling of controlling and influencing the external world. Concerning the latter, somatic sensation functions as an indicator for the preservation of the body and, therefore, the self [25].

We suggest, however, that there is an additional *direct* and ethically relevant relation between somatic sensation and bodily autonomy. This relation exists because of the specific conjunction of two factors: an important but under-appreciated concept of sensory autonomy, and the distinctive neurophysiology of somatic sensation.

Sensory autonomy refers to the ability and right of individuals to control their sensory inputs. Moral agents are assumed, *ceteris paribus*, to seek out, control, and enjoy experiences as they wish. If I want to listen to Beethoven, look at the clouds, or enjoy the taste of an apple, I am free to choose those sensory inputs. Equally, if I do not want to hear rock music, see public nudity on the beach, or smell fried food, then I am free to wear earplugs, close my eyes, or move to a new flat that is not next to a fish and chip shop. Indeed, many traditional rights, such as the right to free movement, can be recast in terms of rights to sensory autonomy – if people have a right to move to the environments that they wish, this is in part related to the sensory experiences produced by such environments. The right to control one's sensory inputs is clearly qualified rather than absolute: my right to listen to Beethoven must be balanced against your right to silence, or to rock music; everyone is expected to listen to the safety briefing on an aeroplane, and so on.

Sensory autonomy is something we value. But sensory autonomy implies some functional method of controlling one's sensory inputs. Human physiology allows us to close our eyes, plug our ears, and hold our noses. Yet this is not the case for the sense of touch, which cannot be turned off. An object in contact with the skin will generate a tactile experience whether the agent wishes to have the experience or not. There is no tactile physiological equivalent to closing the eyes, covering one's ears, or holding one's nose. Instead, the agent's ability to control their tactile experience is physiologically limited and is completely dependent on being able to physically move away from the object of stimulation.

Further, sensory autonomy applies rather differently to different subtypes of somatic sensations. We argue that these differences can be understood with reference to the sensory receptors that give rise to each sensation. With cutaneous mechanosensation (i.e., the sense of touch as conventionally understood) we should be able to choose what we touch and what touches us. Concerning painful mechanical stimuli, for example, such as a bite or a bump, we very strongly expect to have our autonomy respected: other things being equal, nothing and nobody should cause us pain. However, in this case, autonomy is not linked with voluntary choice and preference. Rather, the nociceptive flexion reflex mechanism withdraws the body part that encounters a painful stimulus, operating outside of voluntary control [49]. Sensory autonomy is therefore not only a behavioural matter of preferences, values, and voluntary choices over what sensations an agent actively seeks out or avoids, but a sentient physiological mechanism that activates irrespective of which sensations an agent wishes to have or not have. We accept that we have little or no sensory autonomy over internal skin sensations such as itches and deeper visceral sensations such as aches and cramps – they happen to us, and we cannot choose to have them or not have them. In the specific case of active and passive touch, however, our sensory experiences reflect the interaction of our body with the external environment, so there is a clear convergence of sensory autonomy and body autonomy.

B. Sensory Consent

An important counterpart to sensory autonomy is sensory consent. Consent can have several meanings and is conventionally required or sought in a variety of circumstances from accepting web browser cookies to engaging in physical intimacy. The general principle underlying these situations involves A consenting for another party, B, to act in a particular way Φ, where A has authority to do so. Typically, if B does Φ, it will impact A. It therefore *matters* to A whether B Φs. Consent to interpersonal touch is often carefully transacted and managed, particularly in sensitive contexts such as medical care, personal care, and intimate partner relations. The attention and importance given to consensual interpersonal touch is evidence of its ethical importance – particularly in the context of individual sensory autonomy.

When we consent to B Φing, we are expressing our capacity to either allow or disallow B to Φ. We are allowing something to happen that we also have the authority and right to refuse. This sort of consent is referred to as *permissive* consent as it permits actions that otherwise would be wrong to be right: it relieves others of particular duties [50]. To give an example, I have a right to bodily autonomy – to control what comes into contact with my body. From my right, you derive a duty not to violate this autonomy by touching my body. It would be wrong for you to do this. But I can consent for you to – by exercising my autonomy – thereby relieving you of this particular duty. This sort of transaction happens with *sensory* autonomy too. It follows from *sensory* autonomy that we ought to be able to control our sensory environment. However, there are often circumstances where we waive the right to sensory autonomy: when we enter a cinema, we consent to watch the film, and perhaps also the accompanying trailer and advertisements. In this case, we implicitly consent for the cinema operator to control our visual sensory environment.

A further important concept in this area is *informed* consent. This plays a foundational role in medical ethics by ensuring the autonomy of patients (and their families) are respected [51]. Informed consent requires not only that we approve of a particular act (or medical procedure) but that we sufficiently understand what it will entail. It is, therefore, particularly relevant for an ethics of touch given the epistemic opacity of somatic sensation. Epistemic opacity refers to how the experience of touch is privately and uniquely accessible to the one who is experiencing it. Therefore, it is difficult (and perhaps impossible), that another party could fully inform you about what a particular tactile experience will be like for you.

Finally, and importantly, the status of consent is ongoing and subject to updating. Consent can be withdrawn or affirmed at different stages of a series of events, and this updating can be either explicit, or implicit through contextual inferences [52].

IV. The Ethics of Digital Touch

Based on this analysis of tactile sensation and sensory ethics, we can now turn to the specific ethical issues raised by *digital* touch (see Table II). Many digital touch scenarios envisaged by haptics designers and engineers do not raise any specific ethical concerns. The fact that a device provides mechanical stimulation to the body surface is not in itself ethically problematic. For example, manual interaction with haptically-active objects [2], the use of haptic devices for teleoperation [53], and simple tactile alerting devices [54], represent enhancements of normal sensorimotor activity that are in common use. These scenarios do not raise concern if they are used safely to avoid noxious levels of stimulation. In some cases, however, digital touch does raise specific ethical concerns that need to be explored, considered, and mitigated. This is often because of specific features of touch as a sensory channel. In particular, the directness and privacy of touch, its links to bodily self-awareness, and its unique standing with respect to sensory autonomy.

This section sketches an approach to these issues. We base our approach on the forward-looking principle of the European Commission's Responsible Research and Innovation framework [55]. To our knowledge, none of these issues have yet arisen in practice. But we believe they could arise in the future and, if they did, would cause significant ethical implications. Current systems for digital touch remain relatively unsophisticated, but the range and realism of digital touch sensations are growing rapidly, particularly with novel hardware devices for improved stimulation. For this reason, researchers and designers will benefit from upstream thinking about the ethical hazards of future digital touch systems (see Table II).

A. The 'always-on' nature of tactile sensation and sensory autonomy

The first ethical hazard for digital touch lies in the potential for users' sensory autonomy to be compromised. This concern primarily arises in situations where a user is being passively touched by a digital touch system, rather than where the user is actively using a digital touch interface (for example, during haptic exploration [56]). However, where a user's active engagement with a haptic interface passively stimulates another user – as is the case in virtual environments [57] – autonomy becomes a central concern.

Many wearable digital touch systems are based on tactile stimulation focally or on wide body regions. Mid-air haptics allows similar stimulation to be applied even without contacting a stimulation device – even through clothing [58]. These systems are often designed to provide alerts, or to deliver specific touch sensations for mood induction, and may be useful in that context. But, because of the 'always on' nature of tactile sensation, haptic devices – especially those that do not require direct contact – may limit and compromise a user's ability to control their own tactile sensations and experiences. It follows from the principle of sensory autonomy that a user should be able to make an active and deliberate choice regarding whether they wish to experience a particular sensation – they should consent to have that experience.

With current systems, this remains relatively straightforward. Consent can be established explicitly or more implicitly as appropriate. For example, one might argue that the user consents to a particular tactile sensation at the point where they put on a wearable appliance or approach a stimulating device. In either case, it must be informed, continuous, maintained, and re-established as required. This will be easier with some devices than others.

Establishing consent means ensuring that a user understands as far as possible what a particular sensation will be like. Of course, some digital touch experiences will be entirely new to the user. As a result, fully informed consent may be impossible. It follows from the first-person nature of experience that a person cannot be fully informed until they have had that experience. Facilitators of digital touch should ensure that prior to a user using a device they disclose relevant information about what the experience will be like by, for example, appealing to similar familiar experiences.

However, some devices may be capable of delivering many sensory experiences differing in intensity, quality, and impact. Since on any given occasion, a user may wish to be touched in one way, but not in another way, these devices cannot assume global consent. The design of such systems should therefore include a continual 'checking in' process to ensure the user can and does give informed consent for each particular type of sensation.

Current digital touch systems are designed to require implicit consent – for example, a user must (themselves) switch the device on. However, the same technologies could allow one person to administer unwanted touch to another, having initially coerced them into turning the device on. Digital touch systems, therefore, could be used in ways that enable or accentuate asymmetric power relations between individuals. This power aspect of digital touch could be particularly problematic in scenarios involving intimidation and/or gender violence. For example, a malevolent individual could use digital touch to exert remote coercive control over a person's body from a distance. A victim might then be unable to physically get away from their abuser. Digital touch systems will need to make sure they do not become a channel for abusive, unwanted, and unavoidable interpersonal touch that violates the principle of sensory autonomy. Power asymmetries in digital touch operate between individuals engaged in interpersonal interaction, but specific groups of people may be affected in specific ways. For example, digital touch may be experienced differently, and the negative effects of power asymmetry may represent a greater risk for children, the disabled, and LBGT+ communities. Rich traditions in women's studies, queer studies, and other areas, have highlighted the specific issues surrounding body concepts and bodily sensations in such groups and the risks of harm due to power asymmetries in interpersonal interactions [59]. Future digital touch technologies will need to attend to the diversity of body experiences, and the needs and interests of specific groups of people, in order to be as inclusive as possible.

One way to ensure a user's sensory autonomy is respected is by implementing a readily available off-switch which terminates the sensory experience, or the relevant part of it, when activated. The ethical principle of sensory autonomy requires that this off-switch should always remain under the control of the person being touched – the experiencer – not the person doing the touching. Further, even if a digital touch stimulus seems to be innocuous, the epistemic authority of touch means that a user could potentially experience that stimulus as unwanted and unpleasant. Providing the user with the means to control the tactile stimulations they receive is key to respecting their autonomy, even when the stimulation may seem to be non-noxious or has been designed to be pleasant.

This point may seem trivial from a systems and interaction viewpoint, but it has high ethical significance. Many digital touch systems are designed precisely around the always-on nature of touch, aiming to provide alerting functions or continuous background input. Stick-shakers in modern aircraft cockpits, for example, give a tactile alert if an emergency situation arises. Here, there are good safety reasons to exploit the always-on nature of touch and for pilots to waive their normal sensory autonomy.

However, other scenarios that exploit the 'always on' nature of touch might be unwanted and potentially problematic. Vibrotactile alert systems in mobile phones, for example, are important from a user experience and design perspective because they provide useful, silent, and covert notifications. However, they have also been shown to negatively affect attention [60], [61] and have phantom physiological effects [62]. Whilst these notifications can be turned off, their intensity cannot be altered – limiting the user's control over them. Recent studies have considered the use of tactile stimulation to 'nudge' users on shopping apps to influence purchases [63]: such use of HT severely undermines user autonomy and control [64].

It is also easy to imagine future ethical edge cases. Employees in a future packing or assembly line might be required to wear a tactile vest whose stimulators aim to increase their arousal levels, and thus their productivity. Or advertising boards might be able to reach out and 'touch' you as you walk by. In the future, digital touch designers will need to balance the attractions of touch as an always-on channel for human-machine communication with the ethical requirements of individual sensory autonomy.

B. Transparency of interpersonal digital touch

A central underlying concern is the question of who or what is controlling a particular digital touch system: 'who (or what) is touching me?' Digital touch systems offer the possibility for one person to stimulate another. These stimulations might be intimate, such as kisses or caresses [9], or less intimate, such as handshakes [8] and similar gestures. In general, it is important that we know who is touching us, particularly in the former cases. Where haptic interaction is devoid of a wider sensory context – for example, through haptic bracelets [65] – users may not be able to see, hear, or know whom they are interacting with and who is touching them. Users are unlikely to use such a device if they fear being touched by an unknown agent, potentially without their consent. Therefore, the security of such devices will be essential for their acceptability and ethical use.

Within the emerging technology literature, especially the ethics of artificial intelligence, 'transparency' is often invoked in discussions emphasising the need for interpretable autonomous decision procedures – to avoid the so-called 'black box' problem [66]. We argue that transparency for digital touch will be rather different and will depend on understanding the agents and reasons behind tactile stimulation, rather than understanding how the system itself works. Users who are receiving tactile stimulation will reasonably want to know who or what is controlling these sensations and for what purpose. Consent is often considered particular and personal: whilst we may consent for A to act in a particular way, this does not mean we consent for B, C, D etc. to too [51]. System designers, therefore, ought to ensure it is clear to users who (or what) is touching them. This information should be clearly and continually available to the user so that their consent is ongoing. Moreover, the system should ensure that *only* those agents that a user has consented to send them touch experiences are able to do so.

Interpersonal touch involves potential threats and risks and is therefore carefully regulated by explicit regulations (i.e., laws) and cultural norms. These norms incorporate a large range of contextual factors that surround any instance of interpersonal touch, including the identity and capacity of the individuals involved, their consent to being touched, the location of the touch on the body, and social/environmental contextual factors (such as public versus private spaces). Further, the norms regarding permissible and appropriate touch vary between one society and another. Designers and engineers should consider embedding in digital touch systems the same kind of user identification and consent mechanisms that govern how individuals touch each other in everyday human-to-human interaction. This will require a degree of attention to cultural norms surrounding touch that is lacking in current tactile engineering practices.

Finally, digital tactile communication systems will need careful privacy planning to prevent third parties from disingenuously obtaining the right to touch individuals or information about how they have been touched. Privacy and data security will be essential to provide assurance to users about who is involved in any tactile communication and how they are involved [67].

C. Directness and realism

Digital touch breaks and manipulates the traditional link between our sense of touch and our direct access to the physical world. Because we ordinarily consider touch as our ultimate arbiter of reality, digital touch systems open the possibility for powerful forms of deception. For example, an immersive VR system involving digital touch can disassociate appearance from reality by making objects that we ordinarily associate with being cold – such as ice – feel hot. Digital touch systems already manipulate haptic feedback so that users can manually interact with objects that are not physically present but whose physical characteristics are simulated by a haptic robot. The same systems could, in principle, give haptic feedback consistent with an object not being present when, in reality, it is – for example, a full-body haptic suit could reproduce an illusion of floating even as one touches physical objects and remains in contact with the environment.

Although current haptic devices can produce compelling illusions in manual object perception [68], they are far from producing convincing global misrepresentations of (virtual) reality. However, this may come with improved hardware, particularly given the strength of multisensory illusions in which touch inputs are combined with other sensory modalities such as vision and audition. Mid-air haptics, for example, could 'confirm' hallucinatory visual perceptions of an object that does not exist, or provide tactile sensations in the absence of any visible object.

Any device which has the potential to undermine this fundamental epistemic function of fitting the mind to the world carries ethical risk. The relation between sensory experience and physical environment may be so deeply rooted that major transformation of these experiences might bring unacceptable hazards. Several philosophical thought experiments are based on decoupling the mind/brain from its environment such as the brain-in-a-vat scenario [69], [70]. This thought experiment is normally used to analyse the necessary and sufficient conditions for consciousness. For instance, enactivist and embodied views of mind typically stress the richness of the interface between the mind and the body and conclude that a brain in a vat would not have human-like consciousness [71], [72]. We raise the question about whether an ethical element is implicit in the thought experiment: if a physical body is required for consciousness, is it wrong to remove brains from bodies and put them in vats for the specific reason that interference in the interface between mind and world is wrong?

The extent to which these considerations actually become ethically problematic will depend strongly on context and use case. Novels, films, and video games frequently represent alternate realities, yet they (often) do not raise particular ethical concern. Further, digital touch offers numerous potential benefits by manipulating reality. Tactile stimulation in conjunction with VR for exposure therapy in treating PTSD and trauma survivors [73] offers one example. Indeed, an improved capacity to alter reality is considered to be a beneficial innovation for therapy *because* of its immersion.

Some general worries, however, have been raised concerning immersive technologies. Slater and colleagues suggest the following concerns which, when combined with the immersive potential of haptic technologies, might be accentuated [74]:

- Uncertainty regarding whether past and current events are real.
- Carrying out a physical act in VR that does not correspond to reality – for example, attempting to sit on a virtual chair which is not there (in reality) might lead to physical harm.
- Difficult real-world transitions after a particularly intense and emotional virtual experience.

These concerns will presumably apply more strongly to VR systems that involve digital touch than to those that do not. We suggest that the primary reason for this is the special link between our tactile experiences and sense of self, leading to stronger effects of immersion. We expand on this point in the next section.

D. Body-ownership and the self

Ordinarily, the body is itself an object of touch. Touching our bodies, therefore, serves to signal bodily self-awareness. Digital touch systems could potentially interfere with this. For example, a system could be designed so that I reach out with my right hand to touch my left, but do not feel my left hand being touched, or I feel it being touched in a way other than expected. Experimental studies confirm that this produces a potentially challenging experience of body disownership as if my left hand is no longer mine [75], [76]. Moreover, because our tactile perception of the world is necessitated by the capabilities of our bodies, the body's interaction with the external environment becomes constitutive to individual experience and identity, grounding a sense of self [77]. Digital touch can intervene in the processes of building and maintaining a basic bodily self in interaction with the external environment.

These interventions to the minimal bodily self could of course be valuable, in prosthetics applications perhaps, but they could also carry potential ethical risks. For example, future digital touch systems could manipulate touch sensations to produce dramatic distortions of bodily self-awareness. I might reach up to run my hands through my hair or put a piece of food in my mouth, but the tactile feedback from the system could mean that I do not feel my head or my mouth at all. My tactile experience engineered by the digital touch system could be consistent with having no hair, no head, and no mouth. These experiences could be challenging, upsetting, and difficult to process [78]. Future digital haptic devices might potentially reproduce the pathological experiences of depersonalisation (being disconnected from one's body [79]) or Cotard's syndrome (the delusion that a particular part of the body, or the whole body, is failing, dead, or lost [80]).

These concerns will be particularly acute when combined with VR and experienced in virtual worlds. Within these worlds, users are often represented by avatars. Avatars are our virtual representations. They capture our expressions and movements performed in the real world and translate them into virtual environments [81]. The role of avatars in immersion, realism, and social interaction has received substantial scholarly attention. Studies have, for example, investigated its effect in video games such as *"World of Warcraft"* [82] and 3D virtual worlds such as *"Second Life"* [83]. As representations of ourselves in virtual worlds, avatars may be, to an important degree, extensions of ourselves and our bodies.

As Chalmers argues, the parallelism between the avatar as one's body in the virtual environment and one's own body in real life need not involve any illusion [84]. To some extent, your avatar really is *your* body, just a virtual one. And we might come to identify with our avatars as being a part of who we really are. This might raise ethical concerns. Could a user begin to prefer their virtual self and body over their real self and body? This could lead to real world body dysmorphia and disassociation. Could prolonged use of an avatar-based system desensitise a user to *real* negative experiences? This could lead to a user losing their ability to accurately assess and respond to *real* potential threats. Prolonged experience through a virtual self – that is *not* harmed by negative experience – could alter learned responses to threats important to *real* selves.

In general, digital touch designers and engineers should be cautious about interventions in touch that fundamentally and continuously alter the user's relation to their body and the surrounding environment. In a previous paper on the ethics of VR [74], the authors pointed out that the ethical aspects of long-term immersive use of VR had barely been studied. Similarly, current digital touch systems provide brief, episodic experiences. New ethical implications of digital touch might become apparent if such systems become immersive, pervasive, and persistent.

V. Toward an Ethical Design Framework for Digital Touch Systems

Based on the above considerations, Fig. 2 illustrates an ethical design process for haptic designers and engineers to consider when building digital touch systems. Its purpose is to guide the initial design process to ensure the ethical use of digital touch technology and to flag where mitigation might be required. Fig. 2(a) presents four hierarchically ordered ethical design questions. They ask: whether the system facilitates active or passive touch; *where* on the body it applies stimulation; *how* the system evokes tactile sensations; and whether the user can *control* these sensory inputs. Fig. 2(b) presents these questions as a decision tree. It demonstrates how they can be used by considering a hypothetical haptic alert system in a smartwatch.

A. Does it facilitate active (touching) or passive (being touched) touch?

We have argued that the distinction between active and passive touch is crucial for the ethics of digital touch. Active touch implies a degree of implicit consent and expectation, whereas passive touch refers to tactile sensations that may occur without any decision to directly engage with one's environment. A haptic device providing passive touch may cause tactile sensations without the user first explicitly engaging. Indeed, this is the core function of tactile alerting systems. Passive tactile sensations of this kind could potentially bypass consent and expectation. Passive tactile sensations, therefore, threaten to infringe upon our autonomy as they are often out of our control and, because of this, can be unwanted. Thus, for a tactile alerting device, the user ordinarily consents and expects, at the time of putting on or setting the device, that they may stimulated by it later without any further decision process or consent. Thus, the initial expectations of the user are essential to guarantee their sensory autonomy at the later moment of tactile stimulation.

B. Where is the tactile stimulation administered?

Where on the body a digital touch device stimulates the body is clearly ethically important. Some parts of the body are more sensitive – both physiologically and socially – than others. Smartwatches are generally worn on the wrist. This body location may produce different ethical concerns, and often fewer ethical concerns than other body areas such as the face, or torso. Body location is clearly relevant in relation to risks, possible harms and safety. However, body location is also relevant in relation to user perceptions, user well-being, and the ethical meaning of the interaction between user and device. Further, there are strong cultural norms and expectations regarding different parts of the body, which vary across different societies, and across individuals in any particular society. The ethical issues surrounding wearable tactile devices, in particular, need to take careful account of body location for both physiological and socio-cultural reasons.

C. How is it administered?

We distinguish two ethically relevant considerations regarding *how* a haptic device facilitates tactile stimulation, focussing on the case of interpersonal tactile communication between two individuals. The first consideration is: 'who (or what) is in control of the tactile sensations I am experiencing?' There are different possible answers to this question: the system alone may control the tactile stimulation; a human user may send tactile stimulations to another user without being able to receive a returning sensation back themselves; or there may be bidirectional tactile stimulation communicated between two human users (see Fig. 3). The primary ethical concern related to these different dynamics is the potentiality for malevolent control and coercion. For example, where a user can only receive and not send tactile sensations, this has the potential to reinforce undesirable power relations between people. Interpersonal tactile communication systems that are symmetric and reciprocal raise fewer ethical concerns than systems that are asymmetric and hierarchical.

Second, we consider the *method* of haptic feedback. There are various haptic feedback methods, including kinaesthetic, stretch, MidAir, vibrotactile, and thermal. It is important to disclose which method is being used for informed consent. It is especially important for touchless MidAir technology as this method creates a 'field' of stimulation, which may make users feel trapped and unable to escape.

D. How easily can it be stopped?

The user's ability to stop haptic stimulation is crucial: the off switch is always fundamentally important. For many grounded haptic devices, the user can simply let go of the device, or move away from it, providing a very natural equivalent to the off switch. However, for devices attached to the body, such as tactile vests, exoskeletons, and bracelets, this option is less readily available. The ethical importance of an off switch ensuring sensory autonomy increases for stimulation devices in temporally continuous and spatially broad contact with the body.

E. Ethical Reflection

The paths through the decision tree demonstrate a passive touch stimulation on the wrist. The interaction dynamic is classified as *system→user* because the primary function is to communicate notifications to the user (but see [67, p.4]) for nuance). By identifying these properties, we can identify the key ethical considerations that arise. In this case, they include potential intrusiveness, potential lack of sensory autonomy, and potential lack of control over the tactile input. In fact, many existing smartwatches use vibrotactile feedback and users seem to benefit from them without widespread ethical concerns being reported. Nevertheless, because they are wearable and designed to be always on, the wearer is always tactually reachable. This might cause sensory intrusions from a lack of sensory autonomy. However, designers can protect users' sensory autonomy in several ways: by making the tactile-enabled watch easy to take off, by providing controls to cancel or pause tactile stimulation, and by ensuring prior consent to subsequent stimulation. Building such simple design principles into digital touch devices and interactions will play a major role in managing the ethical risks associated with digital touch.

VI. Conclusion

This paper has outlined some principles and problems for the ethics of digital touch. Digital touch technologies will allow enhanced experience and improved functionality in many human interaction applications. However, as with any emerging technology, it is important to consider ethical implications *before* harms come about, not least because those harms can thus be a11voided or mitigated. The Responisble Research and Innovation framework [55] offers one set of tools for including scientists and engineers, as well as users and ethicists, in such upstream ethical reflection.

The challenge for researchers and designers is to incorporate these ethical considerations into system design. We have argued for an important ethical distinction between touching and being touched, and we have highlighted some of the ethical risks associated with being touched by digital touch systems. We then developed a principle of sensory autonomy which will have particular importance for those digital touch systems that are directly body-mounted and always-on. We also noted that many digital touch systems involve interpersonal digital touch, highlighting that transparency over who is touching whom, and power asymmetries between the sender and receiver of digital touch, both raise important ethical questions in digital touch. Finally, we have indicated that touch has a special epistemic status in defining the bodily self and its relation with the world. Digital interventions into this relation are in their infancy, but could potentially become powerful in the future. The ethics of far-reaching epistemic challenges that future digital touch systems will provide have not yet been considered.

Several future directions for ethical research remain. First, our concept of sensory autonomy could be developed, compared across different sensory modalities, and related to existing concepts in ethics and wider social sciences, such as privacy and individual liberty. Second, the overlaps and interactions between sensory technologies, such as HT, and neurotechnologies, such as brain-machine interfaces, should be further explored. Third, our study has provided analyses of potential ethical concerns, but has not developed actionable guidelines or red lines for haptic designers and engineers. The high personal relevance of touch means that further reflection and consensus-building will be required in development of actionable guidelines. For example, are there any tactile experiences that HTs should should not deliver under any circumstances? Pain might seem like one obvious example. On the other hand, pain might be valuable as a trigger for defensive reactions and self-preservation. Future guidelines will have to consider whether haptic systems should provide such capabilities, or should be limited to certain types of tactile experience, and how misuse of such systems can be avoided.

Acknowledgment

This study was supported by funding from the European Union's Horizon 2020 research 562 and innovation programme under grant agreement No 101017746 for both NB and PH. We gratefully acknowledge the valuable assistance of Pierre Dragicevic, Cristina Gonzalez-Liencres, and Orestis Georgiou for their comments on an earlier draft, as well as Jing Xue and Madhan Vasudevan for their helpful comments on a previous iteration of Table II.

References

[1] R. Raisamo, K. Salminen, J. Rantala, A. Farooq, and M. Ziat, "Interpersonal Haptic Communication: Review and Directions for the Future," *International journal of human-computer studies*, vol. 166, p. 102881, 2022, doi: 10.1016/j.ijhcs.2022.102881.

[2] A. M. Okamura, C. Richard, and M. R. Cutkosky, "Feeling is believing: Using a force-feedback joystick to teach dynamic systems," *Journal of Engineering Education*, vol. 91, no. 3, pp. 345–349, 2002.

[3] I. Rakkolainen, E. Freeman, A. Sand, R. Raisamo, and S. Brewster, "A survey of mid-air ultrasound haptics and its applications," *IEEE Transactions on Haptics*, vol. 14, no. 1, pp. 2–19, 2020.

[4] K. Suzuki and H. Jansson, "An analysis of driver's steering behaviour during auditory or haptic warnings for the designing of lane departure warning system," *JSAE review*, vol. 24, no. 1, pp. 65–70, 2003, doi: 10.1016/S0389-4304(02)00247-3.

[5] B. Delhaye, V. Hayward, P. Lefèvre, and J.-L. Thonnard, "Texture-induced vibrations in the forearm during tactile exploration," *Frontiers in Behavioral Neuroscience*, vol. 6, Jul. 2012, doi: 10.3389/fnbeh.2012.00037.

[6] L. Peiris, Y.-L. Feng, L. Chan, and K. Minamizawa, "Thermalbracelet: Exploring thermal haptic feedback around the wrist," presented at the Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019, pp. 1–11.

[7] Z. F. Quek, S. B. Schorr, I. Nisky, A. M. Okamura, and W. R. Provancher, "Augmentation of stiffness perception with a 1-degree-of-freedom skin stretch device," *IEEE Transactions on Human-Machine Systems*, vol. 44, no. 6, pp. 731–742, 2014.

[8] N. Bailenson, N. Yee, S. Brave, D. Merget, and D. Koslow, "Virtual interpersonal touch: expressing and recognizing emotions through haptic devices," *Human–Computer Interaction*, vol. 22, no. 3, pp. 325–353, 2007.

[9] E. Y. Zhang, S. Nishiguchi, A. D. Cheok, and Y. Morisawa, "Kissenger–development of a real-time internet kiss communication interface for mobile phones," presented at the Love and Sex with Robots: Second International Conference, LSR 2016, London, UK, December 19-20, 2016, Revised Selected Papers 2, Springer, 2017, pp. 115–127.

[10] M. O. Ernst and M. S. Banks, "Humans integrate visual and haptic information in a statistically optimal fashion," *Nature*, vol. 415, no. 6870, pp. 429–433, Jan. 2002, doi: 10.1038/415429a.

[11] A. M. L. Kappers and W. M. Bergmann Tiest, "Haptic perception," *WIREs Cognitive Science*, vol. 4, no. 4, pp. 357–374, 2013, doi: 10.1002/wcs.1238.

[12] Z. Pylyshyn, "Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception," *Behavioral and Brain*

Sciences, vol. 22, no. 3, pp. 341–365, 1999, doi: 10.1017/S0140525X99002022.
[13] V. Hayward, “A brief taxonomy of tactile illusions and demonstrations that can be done in a hardware store,” *Brain Research Bulletin*, vol. 75, no. 6, pp. 742–752, Apr. 2008, doi: 10.1016/j.brainresbull.2008.01.008.
[14] D. S. Pamungkas and K. Ward, “Electro-tactile feedback system to enhance virtual reality experience,” 2016.
[15] M. Fulkerson, “Touch Without Touching,” *Philosophers’ imprint*, vol. 12, p. 1, 2012.
[16] M. Martin, “Sight and touch,” in *The Contents of Experience*, tim Crane, Ed., Cambridge University Press, 1992, pp. 196–215. doi: 10.1017/CBO9780511554582.010.
[17] E. Schwitzgebel, “The Unreliability of Naive Introspection,” *The Philosophical review*, vol. 117, no. 2, pp. 245–273, 2008, doi: 10.1215/00318108-2007-037.
[18] A. Gallace, H. Z. Tan, and C. Spence, “The failure to detect tactile change: A tactile analogue of visual change blindness,” *Psychonomic bulletin & review*, vol. 13, pp. 300–303, 2006.
[19] P. Haggard, T. Cheng, B. Beck, and F. Fardo, “Spatial perception and the sense of touch,” *The subject’s matter: Self-consciousness and the body*, pp. 97–114, 2017.
[20] C. Serrahima, “Sensory fields: the visual and the bodily,” *Philosophical studies*, vol. 180, no. 2, pp. 679–700, 2023, doi: 10.1007/s11098-022-01838-x.
[21] R. Descartes, *Descartes: The World and Other Writings*. in Cambridge Texts in the History of Philosophy. Cambridge: Cambridge University Press, 1998. doi: 10.1017/CBO9780511605727.
[22] J. Boswell, *The Life of Samuel Johnson... With Copious Notes and Biographical Illustrations.* Jones & Company, 1829.
[23] M. T. Fairhurst, E. Travers, V. Hayward, and O. Deroy, “Confidence is higher in touch than in vision in cases of perceptual ambiguity,” *Scientific Reports*, vol. 8, no. 1, p. 15604, Oct. 2018, doi: 10.1038/s41598-018-34052-z.
[24] C. Damiano, C. Antonio, B. Gabriella, G. Hiroaki, and H. Patrick, “Active self-touch restores bodily self-awareness following disruption by ‘rubber hand illusion,’” 2023, *bioRxiv*: doi: 10.1101/2023.03.06.531301.
[25] F. de Vignemont, *Mind the body : an exploration of bodily self-awareness.* Oxford: Oxford University Press, 2017.
[26] D. Katz, “Der aufbau der tastwelt,” 1925.
[27] A. Avramides, “Other Minds,” The Stanford Encyclopedia of Philosophy. [Online]. Available: https://plato.stanford.edu/archives/win2020/entries/other-minds/
[28] D. Chalmers, “Facing up to the problem of consciousness,” *Journal of Consciousness Studies*, vol. 2, no. 3, pp. 200–19, 1995.
[29] N. Block, “On a confusion about a function of consciousness,” *Brain and Behavioral Sciences*, vol. 18, no. 2, pp. 227--247, 1995.
[30] T. Nagel, “What is it like to be a bat?,” *The philosophical review*, vol. 83, no. 4, pp. 435–450, 1974.
[31] F. de Vignemont and P. Jacob, “What Is It like to Feel Another’s Pain?,” *Philosophy of science*, vol. 79, no. 2, pp. 295–316, 2012, doi: 10.1086/664742.
[32] S. Gallagher, “Bodily self-awareness and object perception,” *Theoria et historia scientiarum*, vol. 7, no. 1, pp. 53–68, 2003.
[33] M. Ratcliffe, “Touch and the sense of reality,” *The hand: An organ of the mind*, pp. 131–157, 2013.
[34] M. Tsakiris, M. R. Longo, and P. Haggard, “Having a body versus moving your body: neural signatures of agency and body-ownership,” *Neuropsychologia*, vol. 48, no. 9, pp. 2740–2749, 2010.
[35] G. Gallup, “Self recognition in primates: A comparative approach to the bidirectional properties of consciousness,” *American psychologist*, vol. 32, no. 5, p. 329, 1977.
[36] L. Duhn, “The importance of touch in the development of attachment,” *Advances in Neonatal Care*, vol. 10, no. 6, pp. 294–300, 2010.
[37] M. J. Hertenstein, R. Holmes, M. McCullough, and D. Keltner, “The communication of emotion via touch,” *Emotion*, vol. 9, no. 4, p. 566, 2009.
[38] A. Gallace and C. Spence, “The science of interpersonal touch: an overview,” *Neuroscience & Biobehavioral Reviews*, vol. 34, no. 2, pp. 246–259, 2010.
[39] Y. Zotterman, “Touch, pain and tickling: an electro-physiological investigation on cutaneous sensory nerves,” *The Journal of physiology*, vol. 95, no. 1, pp. 1–28, 1939, doi: 10.1113/jphysiol.1939.sp003707.
[40] F. McGlone, J. Wessberg, and H. Olausson, “Discriminative and Affective Touch: Sensing and Feeling,” *Neuron (Cambridge, Mass.)*, vol. 82, no. 4, pp. 737–755, 2014, doi: 10.1016/j.neuron.2014.05.001.
[41] I. Morrison *et al.*, “Reduced C-afferent fibre density affects perceived pleasantness and empathy for touch,” *Brain (London, England : 1878)*, vol. 134, no. Pt 4, pp. 1116–1126, 2011, doi: 10.1093/brain/awr011.
[42] J. Liljencrantz *et al.*, “Altered C-tactile processing in human dynamic tactile allodynia,” *Pain (Amsterdam)*, vol. 154, no. 2, pp. 227–234, 2013, doi: 10.1016/j.pain.2012.10.024.
[43] P. Bessou, P. R. Burgess, E. R. Perl, and C. B. Taylor, “Dynamic properties of mechanoreceptors with unmyelinated (C) fibers,” *Journal of neurophysiology*, vol. 34, no. 1, pp. 116–131, 1971, doi: 10.1152/jn.1971.34.1.116.
[44] A. J. Todd, “Neuronal circuitry for pain processing in the dorsal horn,” *Nature reviews. Neuroscience*, vol. 11, no. 12, pp. 823–836, 2010, doi: 10.1038/nrn2947.
[45] L. Crucianelli and M. L. Filippetti, “Developmental perspectives on interpersonal affective touch,” *Topoi*, vol. 39, no. 3, pp. 575–586, 2020.
[46] K. Floyd, *Communicating affection: Interpersonal behavior and social context*. Cambridge University Press, 2006.
[47] J. Christman, “Autonomy in Moral and Political Philosophy,” The Stanford Encyclopedia of Philosophy. [Online]. Available: https://plato.stanford.edu/archives/fall2020/entries/autonomy-moral
[48] T. Douglas, “From bodily rights to personal rights,” in *The Cambridge Handbook of New Human Rights*, A. Andreas von, D. Kerstin von der, and S. Mart, Eds. Cambridge: Cambridge University Press, 2020 pp. 378-384.
[49] V. Skljarevski and N. Ramadan, “The nociceptive flexion reflex in humans–review article,” *Pain*, vol. 96, no. 1–2, pp. 3–8, 2002.
[50] T. Dougherty, *The Scope of Consent*. Oxford: Oxford University Press, Incorporated, 2021. doi: 10.1093/oso/9780192894793.001.0001.
[51] N. Eyal, “Informed Consent,” The Stanford Encyclopedia of Philosophy. [Online]. Available: https://plato.stanford.edu/archives/spr2019/entries/informed-consent
[52] M. J. Anderson, “Negotiating sex,” *Southern California law review*, vol. 78, no. 6, p. 1401, 2005.
[53] M. Shahbazi, S. F. Atashzar, and R. V. Patel, “A systematic review of multilateral teleoperation systems,” *IEEE transactions on haptics*, vol. 11, no. 3, pp. 338–356, 2018.
[54] F. Meng and C. Spence, “Tactile warning signals for in-vehicle systems,” *Accident analysis & prevention*, vol. 75, pp. 333–346, 2015.
[55] E. Commission, *Responsible research and innovation (RRI), science and technology: report*. Publications Office, 2013. doi: 10.2777/45726.
[56] R. L. Klatzky and S. J. Lederman, “Stages of manual exploration in haptic object identification,” *Perception & psychophysics*, vol. 52, no. 6, pp. 661–670, 1992.
[57] P. Sykownik and M. Masuch, “The experience of social touch in multi-user virtual reality,” presented at the Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology, 2020, pp. 1–11.
[58] S. Suzuki, R. Takahashi, M. Nakajima, K. Hasegawa, Y. Makino, and H. Shinoda, “Midair haptic display to human upper body,” presented at the 2018 57th Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE), IEEE, 2018, pp. 848–853.
[59] J. Norsigian, *Our bodies, ourselves*. Simon and Schuster, 2011.
[60] C. Stothart, A. Mitchum, and C. Yehnert, “The attentional cost of receiving a cell phone notification,” *Journal of experimental psychology: human perception and performance*, vol. 41, no. 4, p. 893, 2015.
[61] A. Mehrotra, V. Pejovic, J. Vermeulen, R. Hendley, and M. Musolesi, “My phone and me: understanding people’s receptivity to mobile notifications,” presented at the Proceedings of the 2016 CHI conference on human factors in computing systems, 2016, pp. 1021–1032.
[62] R. Rosenberger, “An experiential account of phantom vibration syndrome,” *Computers in Human Behavior*, vol. 52, pp. 124–131, Nov. 2015, doi: 10.1016/j.chb.2015.04.065.
[63] M. Racat and D. Plotkina, “Sensory-enabling technology in m-commerce: the effect of haptic stimulation on consumer purchasing behavior,” *International Journal of Electronic Commerce*, pp. 1–31, 2023, doi: 10.1080/10864415.2023.2226900.
[64] A. T. Schmidt and B. Engelen, “The ethics of nudging: An overview,” *Philosophy Compass*, vol. 15, no. 4, p. e12658, 2020, doi: 10.1111/phc3.12658.

[65] M. F. Rotella, K. Guerin, X. He, and A. M. Okamura, “Hapi bands: a haptic augmented posture interface,” presented at the 2012 IEEE Haptics Symposium (HAPTICS), IEEE, 2012, pp. 163–170.
[66] W. J. von Eschenbach, “Transparency and the black box problem: Why we do not trust AI,” *Philosophy & Technology*, vol. 34, no. 4, pp. 1607–1622, 2021.
[67] M. Ley and N. Rambukkana, “Touching at a Distance: Digital Intimacies, Haptic Platforms, and the Ethics of Consent,” *Science and engineering ethics*, vol. 27, no. 5, pp. 63–63, 2021, doi: 10.1007/s11948-021-00338-1.
[68] H. Dostmohamed and V. Hayward, “Trajectory of contact region on the fingerpad gives the illusion of haptic shape,” *Experimental Brain Research*, vol. 164, no. 3, pp. 387–394, Jul. 2005, doi: 10.1007/s00221-005-2262-5.
[69] H. Putnam, *Reason, Truth, and History*. Cambridge: Cambridge University Press, 1981. doi: 10.1017/CBO9780511625398.003.
[70] M. McKinsey, “Skepticism and Content Externalism,” The Stanford Encyclopedia of Philosophy. [Online]. Available: https://plato.stanford.edu/archives/sum2018/entries/skepticism-content-externalism
[71] J. K. O’regan and A. Noë, “A sensorimotor account of vision and visual consciousness,” *Behavioral and brain sciences*, vol. 24, no. 5, pp. 939–973, 2001.
[72] E. Thompson and D. Cosmelli, “Brain in a vat or body in a world? Brainbound versus enactive views of experience,” *Philosophical topics*, pp. 163–180, 2011.
[73] A. S. Rizzo *et al.*, “Development and early evaluation of the Virtual Iraq/Afghanistan exposure therapy system for combat-related PTSD,” *Annals of the New York Academy of Sciences*, vol. 1208, no. 1, pp. 114–125, 2010, doi: 10.1111/j.1749-6632.2010.05755.x.
[74] M. Slater *et al.*, “The Ethics of Realism in Virtual and Augmented Reality,” *Frontiers in virtual reality*, vol. 1, 2020, doi: 10.3389/frvir.2020.00001.
[75] H. H. Ehrsson, N. P. Holmes, and R. E. Passingham, “Touching a rubber hand: feeling of body ownership is associated with activity in multisensory brain areas,” *Journal of neuroscience*, vol. 25, no. 45, pp. 10564–10573, 2005.
[76] A. Cataldo, M. Di Luca, O. Deroy, and V. Hayward, “Touching with the eyes: Oculomotor self-touch induces illusory body ownership,” *Iscience*, vol. 26, no. 3, 2023.
[77] J. Degenaar and J. K. O’Regan, “Sensorimotor Theory and Enactivism,” *Topoi*, vol. 36, no. 3, pp. 393–407, 2017, doi: 10.1007/s11245-015-9338-z.
[78] R. Newport and H. R. Gilpin, “Multisensory disintegration and the disappearing hand trick,” *Current biology*, vol. 21, no. 19, pp. R804–R805, 2011.
[79] E. C. M. Hunter, M. L. Phillips, T. Chalder, M. Sierra, and A. S. David, “Depersonalisation disorder: a cognitive–behavioural conceptualisation,” *Behaviour Research and Therapy*, vol. 41, no. 12, pp. 1451–1467, Dec. 2003, doi: 10.1016/S0005-7967(03)00066-4.
[80] H. Debruyne, M. Portzky, F. Van den Eynde, and K. Audenaert, “Cotard’s syndrome: a review,” *Current psychiatry reports*, vol. 11, no. 3, pp. 197–202, 2009.
[81] G. Freeman and D. Maloney, “Body, Avatar, and Me: The Presentation and Perception of Self in Social Virtual Reality,” *Proceedings of the ACM on human-computer interaction*, vol. 4, no. CSCW3, pp. 1–27, 2021, doi: 10.1145/3432938.
[82] K. Bessière, A. F. Seay, and S. Kiesler, “The ideal elf: Identity exploration in World of Warcraft,” *Cyberpsychology & behavior*, vol. 10, no. 4, pp. 530–535, 2007.
[83] R. M. Martey and M. Consalvo, “Performing the looking-glass self: Avatar appearance and group identity in Second Life,” *Popular Communication*, vol. 9, no. 3, pp. 165–180, 2011.
[84] D. Chalmers, *Reality+: virtual worlds and the problems of philosophy*. London: Penguin Books, 2023.

TABLES AND FIGURES

Table I

Level of Perceptual Hierarchy	Level of Realisation	Content		Ethical Concern	
		Touching (active)	Being touched (passive)	Touching (active)	Being touched (passive)
Stimulus	Physical Event	Stimulus parameters (e.g., intensity, vibration frequency) Spatial Location relative to the body		• Tissue Damage • Device safety	• Tissue Damage • Device Safety • Effects of surprise
Sensory Experience	Sensation	"It feels smooth (like my cotton shirt)"	"I feel tapping on my back"	• Noxious experiences • Epistemic gap	• Consent • Power asymmetry • Sensory autonomy • Ambiguity about the 'sender' of touch sensations
Belief	Representation	"I am touching my cotton shirt"	"I am being tapped on the back"	• Deception (touch ≠ reality) • Epistemic confusion (what am I touching?)	• Deception (who is touching me?) • Interpersonal confusion

Table II

Application	Active or Passive Touch	Example Device	Body Part	Flow of Interaction	Ethical Considerations
Alerting	Passive	Smart Watches (Wearable)	Wrist	System→User (see Fig.3)	• Sensory intrusion • Always reachable
Object Exploration	Active	Teleoperator (Static)	Hand, Fingers	System→User (see Fig.3)	• Consent ought to be explicitly established for *digital* touch i.e., teleoperation for surgical procedures. • Distraction potential (Safety)
		Automotive Dashboard Interaction (MidAir)			
Communication	Both	Hey Bracelet (Wearable)	Wrist	System Mediated User↔User (see Fig.3)	• Consent needs to be maintained and reaffirmed between parties. • Necessity of easily accessible 'Off Switch' to ensure sensory autonomy
VR Integration	Both	Exoskeletons (Wearable)	Whole body	User↔User (Interaction in Virtual World – see Fig.3) System→User (Interaction with VR environment – see Fig.3)	• Epistemic Deception • Interaction with other sentient agents within virtual environments • Blurred distinction between real and virtual
		Haptic Joysticks (Held)	Hand		

Fig. 1. *Illustrative examples of haptic devices. From left to right: A bHaptics.inc vibrotactile vest; vibrotactile feedback from a video game controller; paired Hey Bracelets for long-distance vibrotactile communication; switching haptic alerts on/off in the settings of an Apple Watch. Images used in this figure are licensed under a Creative Commons Attribution 4.0 International License.*

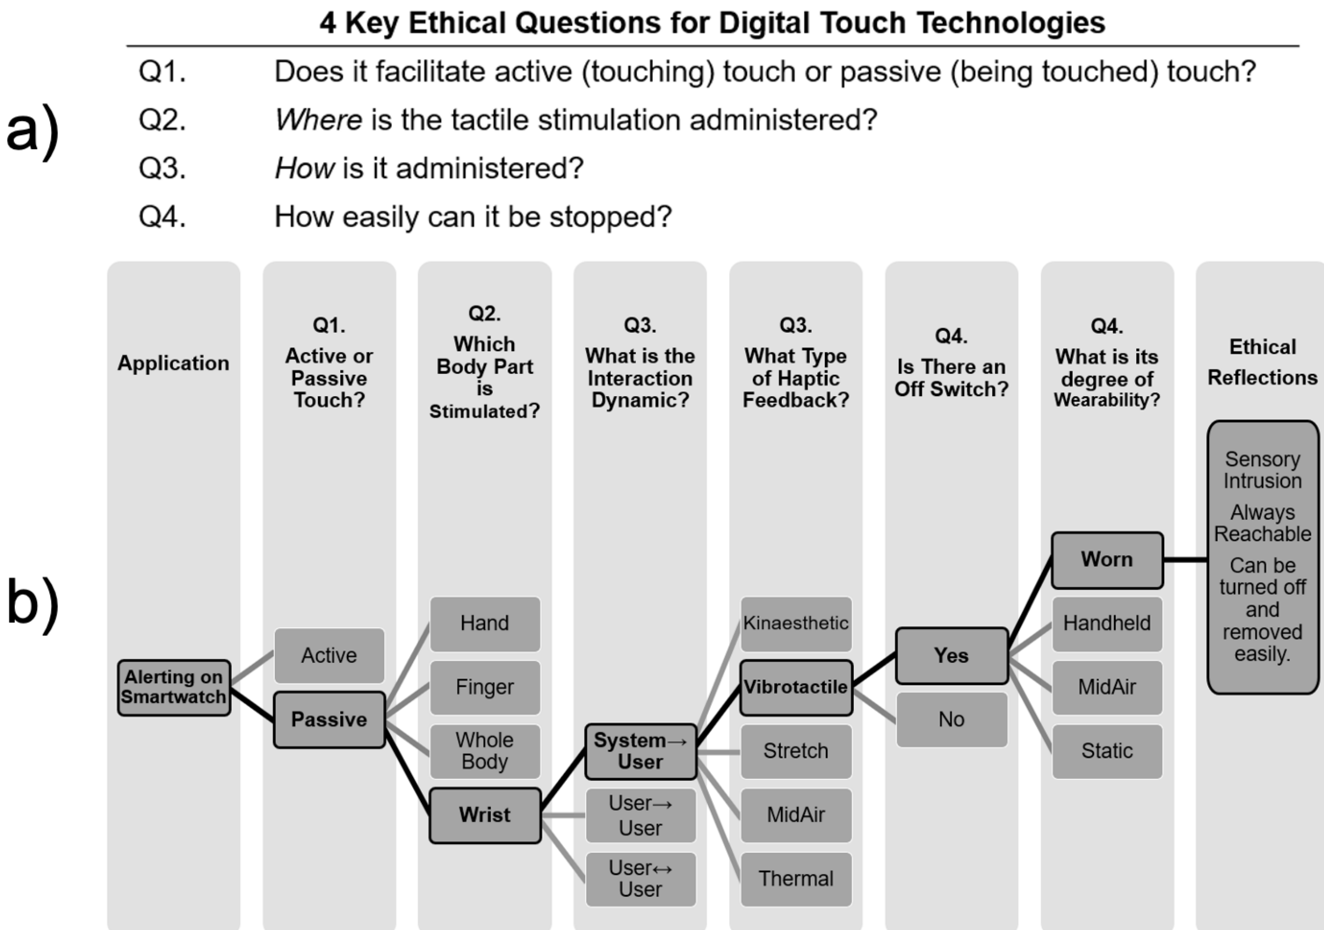

Fig.2. (a) *Four key questions to establish ethical concerns with digital touch systems.* ***(b)*** *The four ethical questions in Fig.2(a) used to construct a decision tree that identifies relevant ethical considerations for a hypothetical tactile alerting device. Black outline and black lines indicate the path taken.*

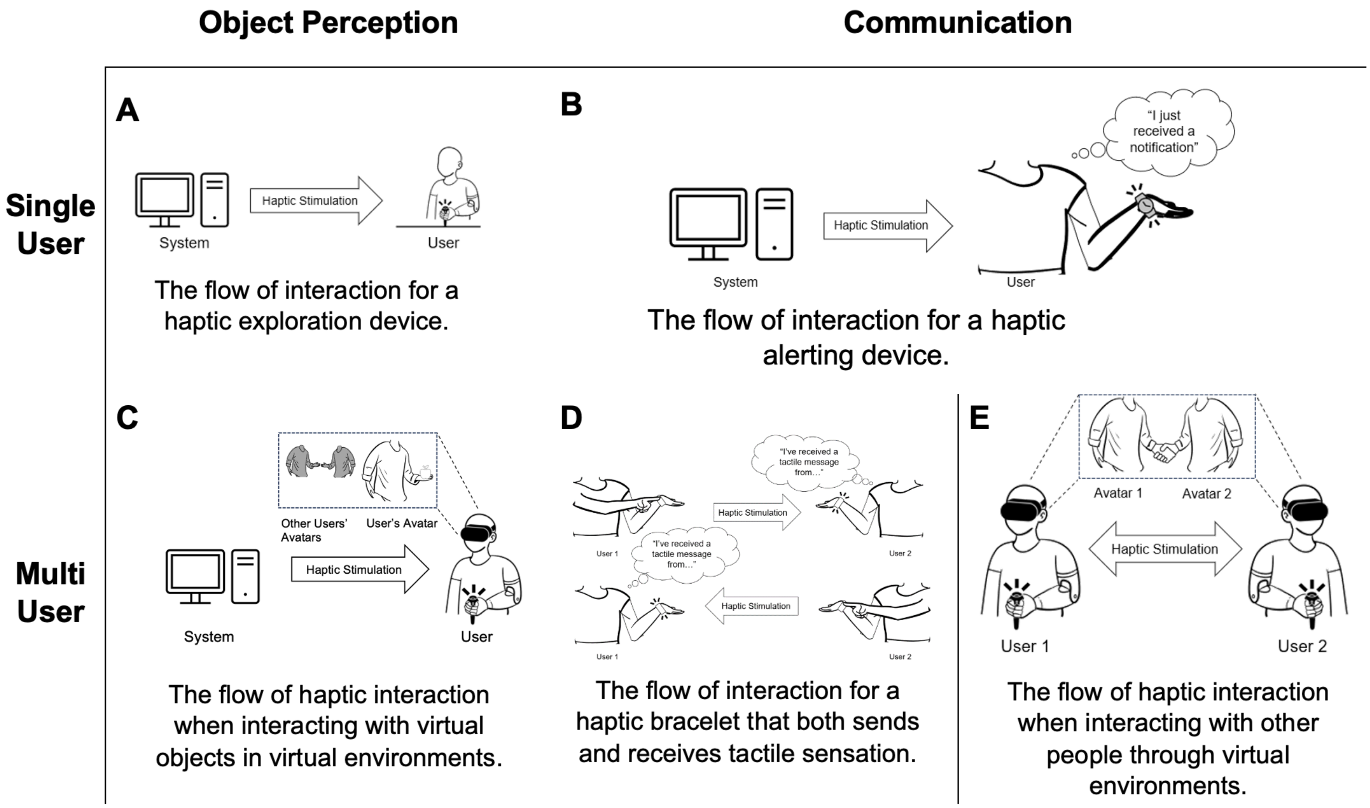

Fig.3 *A schema illustrating flows of haptic interaction and different interaction dynamics for example applications.*