

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Uncovering the Role of Intention in Active and Passive Perception

Permalink

<https://escholarship.org/uc/item/6wj762mq>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Sangati, Ekaterina

Lobo, Lorena

Estelle, Stephen

et al.

Publication Date

2023

Peer reviewed

Uncovering the Role of Intention in Active and Passive Perception

Ekaterina Sangati¹, Lorena Lobo², Stephen Estelle¹, Federico Sangati¹, Samira Tavassoli¹ and Tom Froese¹

¹Okinawa Institute of Science and Technology, Onna, Okinawa 904-0495 Japan

²Department of Psychology and Health, Universidad a Distancia de Madrid, Madrid, Spain
kat.sangati@gmail.com

Abstract

According to the embodied approach to cognition, perception and action are tightly intertwined, as perception is for action and is guided by action. To better understand what this view implies behaviorally, we studied how active movement and intentionality during perceptual exploration affect perceptual accuracy. Participants explored two-dimensional objects using a sensory substitution device, then reported their object size estimates. We manipulated 1) their control over exploratory movements as being either Active (control present) or Passive (control absent) and 2) their knowledge of the task goals, being either Specific (task-focused) or Generic. We found no difference between the Active and Passive conditions but significantly higher perceptual accuracy in Specific Intention trials compared to Generic Intention ones. These results clarify the nature of active perception and contribute to the growing body of evidence that higher level cognitive goals shape how we dynamically sample even low level sensory information from the world.

Keywords: active perception; sensory substitution; goal-directed perception; intentional action

Introduction

Research in embodied cognitive science has emphasized that perception is an active process integrated with an agent's goals and actions (Gibson, 1979; Held & Hein, 1963; Parr & Friston, 2017), a part of the sensorimotor loop (O'Regan & Noë, 2001) rather than a passive first step in adaptive behavior. The specific implications of this idea, however, have yet to be unequivocally demonstrated.

The most frequently explored aspect of active perception is the involvement of motor activity, or movement, in various perceptual tasks. Studies of touch perception, for instance, compared the accuracy of shape recognition between conditions in which participants could move their hand over the stimuli (*active condition*) to those in which their hand was moved over the stimulus (*passive condition*; Heller et al. 1991) and to those in which they held their hand in place and the stimulus was moved over it (*static condition*; Heller 1986). The results from such tactile perception studies so far are contradictory. While there are many that find superiority of active condition in terms of higher perceptual accuracy (Austin & Sleight, 1952; Heller, 1989; Heller et al., 1990; Loomis, 1985), there are also studies in which passive condition seems superior (Heller & Boyd, 1984; Heller et al., 1991; Magee & Kennedy, 1980) or in which there is no difference between the two (Lederman, 1981; Loo et al., 1983).

In addition to merely demonstrating action effects, there is also a question of when action is more crucial for perceptual function—for instance, to what extent does action matter in perceptual learning versus when well established sensorimotor skills are being exercised? This has traditionally been explored with perceptual interfaces called sensory substitution devices (SSDs). In such devices, one type of sensory information (e.g., spatial proximity or sound intensity) is delivered through another sense (e.g., vibrations on the skin). Some studies showed that active control over SSDs is essential for users to learn how to perceive with them (Bach-y-Rita et al., 1969) but contradictory results also exist (Díaz et al., 2012; Froese & Ortiz-Garin, 2020).

The inconsistency of findings could be explained through various means. First, we could do a bottom-up analysis of the conditions in which active or passive perception is better based on particular tasks or stimuli used in them (Symmons, 2004). However, this would not explain why such tasks or stimuli should favor a certain mode of exploration. Second, we could examine whether the sensorimotor loop in both conditions is actually sufficiently matched since if not, then any difference in perception might be attributable to other factors. For instance, static conditions may break the sensorimotor loop rather than merely removing motor activation. Finally, there is also a possibility that perception is never truly “passive” even if the movement itself is not driven by an agent, because the relevant notion of activity is a higher-level one. On such a view, perception is active when it involves exploration in the service of some goal or intention, when it is aimed at detecting task-relevant features of the environment, and when the specifics of that exploration are tuned to the goal and information state of the agent (Bajcsy et al., 2018; Bermejo et al., 2020; Mossio & Taraborelli, 2008). In this broader sense, purposiveness, relevance and attunement could be realized covertly by, e.g., attention modulation, even in situations in which the agent cannot control the motor apparatus or senses it entirely statically (Myin, 2016; Prescott et al., 2011).

The effects of intention on perceptual exploration have been investigated in a number of studies (Lederman & Klatzky, 1987). For example, saccade movement patterns have been shown to change depending on the goal of the observer (Yarbus, 1967). Moreover, Arzamarski et al. (2010) have demonstrated that a change of intention leads to a change in the exploration pattern and a shift of attention in dy-

dynamic touch (touch combined with active manipulation). Intention to perceive particular object features has been shown to affect even such implicit and “passive” movement features as postural sway (Palatinus et al., 2014). Given this evidence, it is plausible that static or passive conditions are detrimental to perceptual performance not simply due to the lack of motor activity involved but rather because this lack is often tied to exploration patterns not sufficiently attuned to the perceptual goals of the particular agent in particular circumstances. Conversely, an adequate exploration pattern can be good enough even in seemingly passive conditions, especially if the agent can attend to relevant features of the task at hand.

Our study strives to make two contributions to this discussion. First, we developed and implemented a novel experimental interface for comparing active and passive spatial perception in a way that makes the process of sensation in both modalities as similar as possible, to result in comparable sensorimotor loops. Second, we jointly examined both aspects of activity (i.e., movement and intentionality) by manipulating 1) the presence or absence of volitional control subjects had over their own movements and 2) the presence or absence of specific intentions when subjects perceived a certain feature of the perceptual stimuli. We reasoned that exploration under a specific intention would lead to a movement pattern better suited to accomplishing the assigned task, thereby leading to better perceptual performance. We also predicted that this effect would hold more strongly for passive participants.

The task we chose was for subjects to use a novel SSD designed by our laboratory, to perceive and estimate object sizes. Sensory substitution was chosen over established sensory modalities to minimize the effect of learned sensorimotor correlations on the differences between conditions of interest. Our hypotheses were as follows:

1. Actively moving (Active Condition) to perceive will lead to more accurate object size estimation than being moved passively (Passive Condition).
2. Having a precise goal during one’s perceptual exploration (Specific Intention) will lead to more accurate object size estimation than exploring without a specific goal (Generic Intention).
3. There will be an interaction effect between the Condition and Intention manipulations.

Methods

Participants

Twenty-seven participants from the University participant pool took part in this study. One participant was excluded from the analysis due to aborted recording session. The remaining participants (15 female, 10 male, 1 non-binary) were all right-handed, had normal or corrected-to-normal vision, and were aged between 24 and 55 years (mean age 32). Each experimental session took a maximum of 90 minutes and participants were paid ¥1500 for their involvement. The

study was approved by the Okinawa Institute of Science and Technology Graduate University Human Subjects Research Review Committee (HSR-2020-019-2) and written informed consent was obtained from each participant.

Participants were semi-randomly assigned to either the Active or Passive Condition group and matched as best as possible on their arm length to ensure anatomical compatibility between the participants in each pair.

Apparatus and Materials

Enactive Torch Sensory substitution device used in this study was Enactive Torch (ET; v.5.4). This tool is composed of a LiDAR distance sensor, a haptic feedback motor and an onboard Arduino controller that translates distance between the Torch and the first-encountered surface into vibration. In this experiment the ET was set to produce a sine wave haptic output with a constant frequency of 155 Hz whenever an object placed within 80 cm distance was encountered. That is, in the context of the present task setup (see Figure 1), the mode of operation was binary: vibration when the ET was in front of the stimulus and no vibration when pointing beyond the stimulus boundaries. Haptic output was applied via an external motor attached to participant’s hand (specifically, their outer thenar webspace) using a soft Velcro strap.




Figure 1: Experimental setup for all trials in the Active and Passive Conditions, with the TA-ET interface on the left.

Robot arm To implement Active and Passive exploration conditions, a robotic arm Torobo Arm (TA; Tokyo Robotics) was used. In the current study, only 2 out of 7 available joints were unlocked, resulting in the movement possibility in only 2 dimensions. TA can operate in two modalities: (1) *external force mode*, in which a human operator controls the movement of the arm while trajectory data (joint positions and angular velocities) are being saved – implementing our Active Condition, and (2) *trajectory mode*, in which the arm can play back a given trajectory from saved data (Passive Condition, from the participant perspective). In this experimental setup, a custom made shelf and handle was attached to the arm. Participants had their arm resting on the shelf, grasped the handle and moved the robot (AC) or were moved by it (PC) while exploring the objects placed in front of them by the experimenter. The ET was inserted into a special slot in the handle and delivered corresponding vibrations.

Stimuli Six rectangles made of white acrylic material were used. The relevant object sizes, that is, the sizes which participants were asked to estimate in every trial were between 10 and 35 cm in increments of 5 cm. The objects were designed such that 1) their two sides had different length and 2) the relevant dimension was shorter or longer in half of the cases. Specifically, the sizes were as follows (the relevant dimension is stated first): 10x15, 15x25, 20x35, 25x10, 30x20, 35x30.

Additional signals Physiological data were collected at 1000 Hz frequency using the Brain Products recording system (BrainAmp, Brain Products GmbH, Gilching, Germany) composed of BrainAmp ExG MR amplifier, BrainVision Recorder software and three types of sensors. First, Brain Products 3D Acceleration Sensor was attached to the Torobo arm's end point to ensure accurate and high-frequency monitoring of the robot movement and enable a precise check of the equivalence of motion trajectories across two conditions. Furthermore, five surface electrodes (Vitrode F disposable electrodes by Nihon Kohden) were applied to participants' skin after it has been rubbed with alcohol. Two electrodes were placed on the biceps brachii muscle following the SENIAM guidelines¹ and a ground electrode on the left clavicle. The EMG activity analysis was done to ensure that Passive Condition participants were indeed passive, i.e., not surreptitiously activating their muscle trying to follow the robot movement. Two further electrodes were used to collect ECG data to explore the possible differences in physiological activity between two conditions – not reported here.

Procedure

After arriving to the lab and giving their informed consent, participants were asked to perform a short Enactive Torch familiarization task that consisted in a custom-built whack-a-mole game. The data from this task was not analyzed. Next, participants were seated at the main experimental interface and fitted with surface electrodes. To become comfortable with a complete Torobo Arm - Enactive Torch (TA-ET) interface, participants performed an additional familiarization task, in which they were asked to explore a shape placed in front of them (with visual feedback) and then identify 2 different shapes (without visual feedback). Both Active and Passive Condition participants performed this part of the experiment actively. Finally, we recorded 2 minutes of baseline activity in a resting state with eyes closed.

In the main task, participants were asked to estimate sizes of rectangles hidden behind the curtain, using the TA-ET interface. There were a total of 3 practice and 72 experimental trials. Trial orders were randomly generated prior to the experiment and were identical for active and passive participant within each matched pair. Experimental interface was programmed using Matlab Psychtoolbox-3 (Kleiner et al., 2007).

On each trial, the experimenter placed the stimuli on the stand in front of the TA-ET interface with the relevant size

presented in either horizontal or vertical orientation. Participants were first given a trial instruction, which could be of three possible types: Generic Intention (“Explore the object”), Specific Intention of Width (“Perceive object width”) or Specific Intention of Height (“Perceive object height”). They then had 10 seconds to perceptually explore the objects. Active participants were free to move the robot arm as they wanted and resulting trajectories were recorded. Passive participants were asked to relax their arm and were moved by the robot by replaying the trajectory of the matched active participant. After the exploration period, the robot returned to the starting position and participants were asked to estimate a particular object Dimension (Width or Height) by drawing a line on a touch screen that corresponded to their estimated length, using their left (non-dominant) hand.² Participants' estimated object size was computed as the Euclidean distance between the initial and final drawing points in xy-coordinates.

Data preprocessing

Accelerometer and EMG recording Acceleration data (recorded in units of g forces) was bandpass filtered at 0.2-20 Hz with a 5th-order Butterworth filter and smoothed with a 5th-order Savitzky-Golay filter with a window length of 250 data points. EMG signal was filtered at 20-350 Hz with a 4th-order Butterworth filter, at 60 Hz (and its multiples) with a notch filter, detrended and rectified. EMG activity was obtained by (1) applying a Teager-Kaiser (TK) transformation and computing a linear envelope, (2) normalizing the envelope by dividing by maximum value in the whole experimental session (familiarization task for Passive participants). Finally, the data was epoched into resting baseline and trials. We decided not to exclude passive trials based on muscle activity level (as it would introduce an exclusion principle that varies between conditions) but we analyzed its distribution across conditions and its effect on performance level.

Movement trajectory clustering In order to check whether intention manipulation led to distinct movement patterns, recorded trajectories were clustered automatically using a heuristic procedure. For each trial, we classified each movement segment (vector between positions in two subsequent timestamps) based on its direction: “horizontal” ($\{0^\circ, 180^\circ\} \pm 5^\circ$) or “vertical” ($\{90^\circ, 270^\circ\} \pm 5^\circ$). Then, we computed the proportion of each direction within the total number of segments and assigned category of “horizontal” or “vertical” when these labels were present over 75% of the time. Otherwise, a “mixed” category was assigned.

Results

We removed trials based on raw response data abnormalities and not based on response accuracy. Specifically, trials were removed in which the response was not a straight line (10 trials), when indicated length was 2 MAD below the sample median (26), when last response velocity was not 0 (3), response

¹<http://seniam.org/bicepsbrachii.html>

²The dominant hand was attached to the TA-ET interface with a soft Velcro strap and could not be easily used to provide responses.

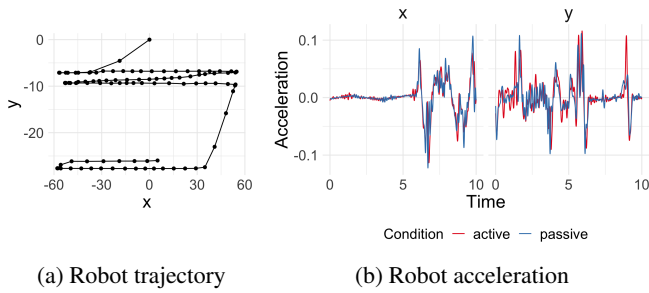


Figure 2: Robot arm movement in Active and Passive Condition.




Figure 3: Movement trajectory examples and cluster distribution by Intention type.

velocity profile was markedly different from other velocities in the sample (based on manual inspection, 2). This resulted in the removal of 2% of the data in total.

First, we checked our experimental manipulations. Figure 2a shows an example trajectory for a particular representative trial while Figure 2b acceleration data for this trial in Active and Passive Conditions. As can be seen, ignoring a small amount of noise, the acceleration profile in the two conditions looks very aligned, attesting to the similarity of the sensori-motor pattern experienced by active and passive participants within each pair. Figure 3a shows movement trajectories plotted separately for different automatically assigned cluster types and Figure 3b shows a distribution of trajectory clusters for the whole sample based on trial instruction. It appears that, indeed, intention manipulation led to different types of trajectories and the pattern is consistent with intuitive predictions: knowing that the task is to estimate the object’s width or height produced trajectories that explored mostly one of these two dimensions (horizontal and vertical). On the other hand, Generic Intention led participants to move in both dimensions.

Estimation accuracy

In order to analyze size estimation accuracy we computed two measures from raw length data. Both measures were chosen over a more direct measure of accuracy such as absolute size error to account for the fact that participants were not given any reference for how the vibration extent perceived via the ET corresponds to actual object size. That is, since the vibration extent depends both on the object size and its distance

to the ET, different participants could have assumed a different correspondence scale. To eliminate this aspect of inter-individual variability, not relevant to the present hypothesis, we focused on the measures that capture the quality and precision of perceptual scale that each participant applied.

First, we looked at **accuracy as Actual-Estimate Correlation (AEC)**, which expresses how distinct the different object sizes appeared (e.g., a correlation close to 0 would indicate that all object sizes were perceived as similar or the answers were given randomly). Each person’s estimated sizes were grouped by within-person independent variables of interest (Intention and object Dimension), a Pearson correlation coefficient was obtained between estimated and actual object sizes and Fisher transformation ($z = \frac{1}{2} \ln(\frac{1+r}{1-r})$) was applied. Comparing resulting coefficients with a mixed ANOVA showed a significant main effect of Intention ($F(1, 24) = 7.94, p < 0.01$) with a large effect size (partial $\eta^2 = 0.25$) but no significant effects of Condition or Dimension and no interaction (see Fig. 4a for an illustration omitting the Dimension variable).

Second, we examined accuracy as **Percent Variable Error (%VE)**, which expresses how consistently each particular size was perceived. This measure is calculated as

$$\%VE = \left[\frac{\sqrt{\sum(x_i - M)^2/n}}{M} \right] * 100 \quad (1)$$

where x_i is the perceived size on a given trial i , M is the mean perceived size and n is the number of trials for a given combination of within-subject variables (real Object Size, Intention, Dimension). Again, a mixed ANOVA showed a significant main effect of Intention ($F(1, 24) = 5.29, p < 0.05$) with a large effect size (partial $\eta^2 = 0.18$) but no significant effects of Condition or Dimension and no interaction (Fig. 4b).

It must be noted that there was considerable individual variability between different pairs of participants in terms of how Active or Passive Condition affected their perceptual accuracy (Fig. 5). In particular, while for two of the pairs (pair 1 and 13), we found that passive participants exhibited significantly higher %VE (at significance threshold of $p < 0.0038$ corrected for multiple comparisons; $t(43.10) = -3.58, p < .001$, Cohen’s $d = -1.09$ and $t(36.12) = -3.69, p < .001$, Cohen’s $d = -1.23$ respectively), there were also pairs in which the active person showed higher error (e.g., pair 6 and 8, non-significant) and pairs in which there was no difference (e.g., pair 3 and 10). This variability likely contributed to the lack of group-level significant effect of exploration condition.

Additional analyses

We examined EMG activity expressed as normalized Teager-Kaiser area across Active and Passive Conditions. As can be seen from Figure 6, passive participants exhibited lower levels of muscle activation during trials than active participants (confirmed with the Welch two-sample t-test, $t(918.41) = 44.26, p < .001$; Cohen’s $d = 2.92$), attesting to the validity of our experimental setup. Furthermore, we




Figure 4: Perceptual accuracy across conditions and intention types. Error bars represent 95% confidence intervals.

found no significant correlation between the average amount of trial EMG activity in the passive group and their average perceptual accuracy, measured either as AEC ($\rho = 0.53$, $S = 172.00$, $p = 0.067$) or %VE ($\rho = -0.26$, $S = 458.00$, $p = 0.394$). It is possible that a more trial-level approach could uncover such correlations within each participant data. However, it would require a different experimental design that allows for a trial-level (as opposed to aggregated) measure of perceptual accuracy.




Figure 5: Difference in %VE between active and passive participant for each pair.

Furthermore, we examined the dependence of size estimate on the time participants spent in contact with the object and average exploration movement velocity. Previous research found that when estimating linear extent of two-dimensional objects by touch, the judgments are based on movement velocity and movement time across the object surface (Armstrong & Marks, 1999; Hollins & Goble, 1988). Time information was available equally to both groups, which could explain a null effect of Condition. In the current setup, data from the ET was not synchronized with robot movement data and participants were not instructed to execute any particular movements. Therefore, we could not extract clean sweeps across the object surface to obtain precise measures of sweep time or velocity. However, recorded trajectories showed that

when given a Specific Intention for a particular direction, participants did in fact move in a sweeping fashion. Therefore, to perform our analysis, we focused on trials in which the motion trajectory was classified as unambiguously “horizontal” or “vertical” and took total contact time in a trial (indicated by the ET motor being switched on) as a proxy for the time that indexes movement across the object.

We fitted a linear mixed model to predict response length with contact time and average movement speed and included random effects of real object size (to eliminate its trivial influence on contact time) and participant ID. The effect of contact time was statistically significant and positive ($\beta = 8.80$, $t(772) = 2.50$, $p = 0.013$) while the effect of average movement speed was statistically significant and negative ($\beta = -1.24$, $t(772) = -6.07$, $p < .001$). The model’s total explanatory power was substantial (conditional $R^2 = 0.69$) but the part related to the fixed effects alone (marginal R^2) was small (0.07). Including an interaction term between contact time and condition did not improve the fit of the model. Therefore, we have tentative evidence that both groups of participants could be relying on time and velocity information but further study with more precise trajectory data is required.

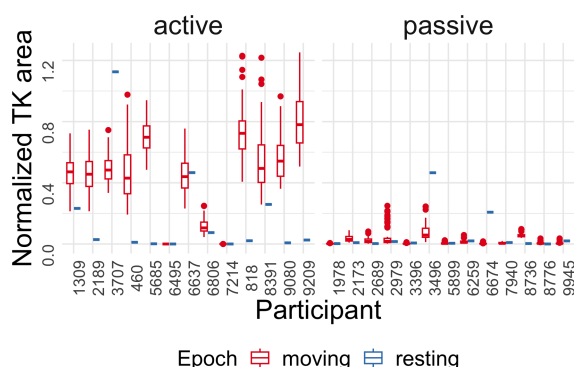


Figure 6: EMG activity as normalized Teager–Kaiser area in Active and Passive Conditions.

Discussion

We have designed a novel experimental interface that allows for testing a variety of perceptual tasks in precisely matched Active and Passive Conditions. Although the attempt to create such an interface is not new (e.g., Perrotta et al., 2020; Richardson et al., 2000), the advantage of our solution is in its potential ability to test a variety of tasks and stimuli (size, shape, affordances perception in 1D-3D) and its integration with SSDs and brain and body physiological recording.

Using the TA-ET interface, we examined size estimation accuracy as a function of exploration mode (Active vs Passive Condition) and intention (Generic vs Specific). We found that perceiving the object under Specific Intention conditions led to higher accuracy, in line with our Hypothesis 2. However, we found no difference between Active and Passive exploration mode, contrary to our Hypothesis 1 and some previous studies (Heller, 1986; Lepora et al., 2013; Smith et al., 2009) but in line with others (Lederman, 1981, 1983; Lederman & Klatzky, 1987). We also found no evidence for interaction between independent variables, contrary to our Hypothesis 3.

Our findings on the significance of perceptual intention are interesting as they highlight the interactions between higher-level cognitive influences, such as goal-orientation, and what is typically thought of as lower level perceptual activity. The potential linkages seen between action-related intentionality and perceptual quality can be explained by recent theories of active inference, which posit that intentions, or the initial goals of perceptual exploration, not only guide the manner in which percepts are cognitively represented but how perceptual information is dynamically sampled from the world (Parr & Pezzulo, 2021). It may be worth further examining, within the active inference framework – or broader embodied cognitive science approach – how the specificity of intentions in perceptual exploration may influence behavioral performance, by either guiding actions such that they result in appropriate stimulation, or by modulating attention to the relevant input (Bermejo et al., 2020; Parr & Friston, 2019), even when the agent is not controlling its acquisition.

Regarding the null effect of Condition, we can consider several explanations. First, it is possible that we did not have sufficient statistical power to find an effect for the group variable, especially given the variability between pairs of participants. This variability could stem from differences in motor abilities required for the task: executing line drawings with a non-dominant hand as a way of assessing object dimensions. Alternatively, it could point to the diversity of perceptual strategies in acquiring required information. Both issues might interact with a coarse-grained active versus passive distinction. Future studies could address these issues in a number of ways. The motor abilities could be assessed explicitly and accounted for in the analysis or, alternatively, a response modality could be replaced with a more discrete type (forced choice). Perceptual strategies could be distinguished based on a more systematic characterization of movement trajectories and related to task performance. Finally, a within-subject

design could be adopted. However, this requires careful consideration of the possible confounds. Preserving an exact match between active and passive trajectories would prevent counter-balancing the order of conditions (as active trajectories have to be recorded first) and introduce a possible learning effect. Alternatively, passive trajectories could be generated in advance, thereby allowing for order counter-balancing but breaking the exact match in sensorimotor patterns.

An alternative explanation for the null effect is that passive participants had sufficient information to adequately perform the given task, without explicit behavioural control during the exploration phase. One source of motor information could have derived from surreptitious muscle activation among passive participants despite the instruction to relax their arm. However, this possibility can be excluded based on generally lower EMG activity in the passive group and a lack of correlation between muscle activation and accuracy. Another possible source of information is task-specific variables. Based on previous studies, it appears that movement velocity and time are sufficient for judging linear extent (Hollins & Goble, 1988). Since proprioception was still available to passive participants (i.e., they were not anesthetized), they could perceive these variables and estimate size accordingly. This possibility ties in with a more theoretical point, according to which *active* perception is not exhausted by the mere presence of motor activity (Barandiaran et al., 2009). Rather, it is perception modulated by the agent's goals and attention to features relevant to the task. In this case, passive participants could still be active in this more covert aspect, since they had the same access to their own movement and time information, and intentional knowledge as active participants.

Additionally, we could consider Active Condition features as a reason for the null effect. It could be that active participants were at a disadvantage because they had a higher cognitive load (Van Doorn et al., 2012) in having to not just perceive but also plan and execute movements with the unfamiliar robot-based apparatus. Furthermore, the TA-ET interface enabled movement that was rich but still more restricted than being allowed to move completely naturally if the Enactive Torch were to be held in participants' hand and manipulated directly. Further study should implement a completely free baseline condition to assess the potential effect of the interface on exploratory strategies and behavioral performance.

In our future work we plan to investigate a more complex perceptual task that allows for more individual variability in exploration strategies, which could enhance the specificity of movement patterns to particular perceivers and thereby also enhance the active-passive contrast. We also plan to examine additional indices of perceptual performance, such as response confidence or the experiential quality of perception with an SSD device, i.e., the extent to which the current distance-to-touch device enables vision-like capacity. Finally, looking at the effect of Active and Passive exploration in a context of learning to perceive would be an avenue to explore.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number 22K18163 awarded to Ekaterina Sangati.

References

- Armstrong, L., & Marks, L. E. (1999). Haptic perception of linear extent. *Perception & psychophysics*, *61*(6), 1211–1226.
- Arzamarski, R., Isenhower, R. W., Kay, B. A., Turvey, M. T., & Michaels, C. F. (2010). Effects of intention and learning on attention to information in dynamic touch. *Attention, perception & psychophysics*, *72*(3), 721–735.
- Austin, T. R., & Sleight, R. B. (1952). Accuracy of tactual discrimination of letters, numerals, and geometric forms. *Journal of experimental psychology*, *43*(3), 239–247.
- Bach-y-Rita, P., Collins, C. C., Saunders, F. A., White, B., & Scadden, L. (1969). Vision substitution by tactile image projection. *Nature*, *221*, 963–964.
- Bajcsy, R., Aloimonos, Y., & Tsotsos, J. K. (2018). Revisiting active perception. *Autonomous Robots*, *42*, 177–196.
- Barandiaran, X. E., Di Paolo, E., & Rohde, M. (2009). Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive behavior*.
- Bermejo, F., Hüg, M. X., & Di Paolo, E. A. (2020). Rediscovering richard held: Activity and passivity in perceptual learning. *Frontiers in psychology*, *11*, 844.
- Díaz, A., Barrientos, A., Jacobs, D. M., & Travieso, D. (2012). Action-contingent vibrotactile flow facilitates the detection of ground level obstacles with a partly virtual sensory substitution device. *Human Movement Science*, *31*(6), 1571–1584.
- Froese, T., & Ortiz-Garin, G. U. (2020). Where is the action in perception? an exploratory study with a haptic sensory substitution device. *Frontiers in psychology*, *11*, 809.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of comparative and physiological psychology*, *56*, 872–876.
- Heller, M. A. (1986). Active and passive tactile braille recognition. *Bulletin of the psychonomic society*, *24*(3), 201–202.
- Heller, M. A. (1989). Texture perception in sighted and blind observers. *Perception & Psychophysics*, *45*(1), 49–54.
- Heller, M. A., & Boyd, M. E. (1984). Touching with a wand. *Perceptual and motor skills*, *58*(2), 390.
- Heller, M. A., Rogers, G. J., & Perry, C. L. (1990). Tactile pattern recognition with the optacon: Superior performance with active touch and the left hand. *Neuropsychologia*, *28*(9), 1003–1006.
- Heller, M. A., Nesbitt, K. D., & Scrofano, D. K. (1991). Influence of writing style and categorical information on identification of tactile numerals and letters. *Bulletin of the Psychonomic Society*, *29*(4), 365–367.
- Hollins, M., & Goble, A. K. (1988). Perception of the length of voluntary movements. *Somatosensory Research*, *5*(4), 335–348.
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in psychtoolbox-3. *Perception 36 ECVF Abstract Supplement*.
- Lederman, S. J. (1981). The perception of surface roughness by active and passive touch. *Bulletin of the Psychonomic Society*, *18*(5), 253–255.
- Lederman, S. J. (1983). Tactual roughness perception: Spatial and temporal determinants. *Canadian Journal of Psychology/Revue canadienne de psychologie*, *37*(4), 498–511.
- Lederman, S. J., & Klatzky, R. L. (1987). Hand movements: A window into haptic object recognition. *Cognitive psychology*, *19*(3), 342–368.
- Lepora, N. F., Martinez-Hernandez, U., & Prescott, T. J. (2013). *Active touch for robust perception under position uncertainty*.
- Loo, C. K., Hall, L. A., McCloskey, D. I., & Rowe, M. J. (1983). Proprioceptive contributions to tactile identification of figures: Dependence on figure size. *Behavioural brain research*, *7*(3), 383–386.
- Loomis, J. M. (1985). Tactile recognition of raised characters: A parametric study. *Bulletin of the Psychonomic Society*, *23*(1), 18–20.
- Magee, L. E., & Kennedy, J. M. (1980). Exploring pictures tactually. *Nature*, *283*(5744), 287–288.
- Mossio, M., & Taraborelli, D. (2008). Action-dependent perceptual invariants: From ecological to sensorimotor approaches. *Consciousness and cognition*, *17*(4), 1324–1340.
- Myin, E. (2016). Perception as something we do. *Journal of Consciousness Studies*, *23*(5-6), 80–104.
- O'Regan, J., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*, 939–1031.
- Palatinus, Z., Kelty-Stephen, D. G., Kinsella-Shaw, J., Carello, C., & Turvey, M. T. (2014). Haptic perceptual intent in quiet standing affects multifractal scaling of postural fluctuations. *Journal of experimental psychology. Human perception and performance*, *40*(5), 1808–1818.
- Parr, T., & Friston, K. J. (2017). The active construction of the visual world. *Neuropsychologia*, *104*, 92–101.
- Parr, T., & Friston, K. J. (2019). Attention or salience? *Current opinion in psychology*, *29*, 1–5.
- Parr, T., & Pezzulo, G. (2021). Understanding, explanation, and active inference. *Frontiers in systems neuroscience*, *15*.
- Perrotta, A., Pais-Vieira, C., Allahdad, M. K., Bicho, E., & Pais-Vieira, M. (2020). Differential width discrimination task for active and passive tactile discrimination in humans. *MethodsX*, *7*, 100852.
- Prescott, T. J., Diamond, M. E., & Wing, A. M. (2011). Active touch sensing. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *366*(1581), 2989–2995.
- Richardson, B., Symmons, M., & Accardi, R. (2000). The TDS: A new device for comparing active and passive-

- guided touch. *IEEE transactions on rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, 8(3), 414–417.
- Smith, A. M., Chapman, C. E., Donati, F., Fortier-Poisson, P., & Hayward, V. (2009). Perception of simulated local shapes using active and passive touch. *Journal of neurophysiology*, 102(6), 3519–3529.
- Symmons, M. A. (2004). Active versus passive touch: Superiority depends more on the task than the mode. *Touch, blindness, and neuroscience*, 179–185.
- Van Doorn, G. H., Dubaj, V., Wuillemin, D. B., Richardson, B. L., & Symmons, M. A. (2012). Cognitive load can explain differences in active and passive touch. *Haptics: Perception, Devices, Mobility, and Communication*, 91–102.
- Yarbus, A. L. (1967). Eye movements during perception of complex objects. *Eye movements and vision*, 171–211.