A somatosensory computation that unifies limbs and tools

Luke E. Miller[1–4], Cécile Fabio[1–3], Frédérique de Vignemont[7], Alice Roy[6], W. Pieter Medendorp[4,8] & Alessandro Farnè[1–3,5,8]

1. Integrative Multisensory Perception Action & Cognition Team - ImpAct, Lyon Neuroscience Research Center, INSERM U1028, CNRS U5292; Bron, France
2. University of Lyon 1; Villeurbanne, France
3. Hospices Civils de Lyon, Neuro-immersion ; Bron, France
4. Radboud University, Donders Institute for Brain, Cognition and Behaviour; Nijmegen, The Netherlands
5. Center for Mind/Brain Sciences, University of Trento; Rovereto, Italy
6. Laboratoire Dynamique Du Langage, CNRS UMR 5596, Lyon, France
7. Institut Jean Nicod, Department of cognitive studies, Ecole Normale Superieure, PSL University, Paris, France
8. These authors contributed equally to the study

* Lead contact. Luke E. Miller: Luke.Miller@donders.ru.nl

Keywords

Tool use, tactile localization, computation, trilateration, extended sensing

Abstract

It is often claimed that tools are embodied by the user, but whether the brain actually repurposes its body-based computations to perform similar tasks with tools is not known. A fundamental body-based computation used by the somatosensory system is *trilateration*. Here, the location of touch on a limb is computed by integrating estimates of the distance between sensory input and its boundaries (e.g., elbow and wrist of the forearm). As evidence of this computational mechanism, tactile localization on a limb is most precise near its boundaries and lowest in the middle. If the brain repurposes trilateration to localize touch on a tool, we should observe this computational signature in behavior. In a large sample of participants, we indeed found that localizing touch on a tool produced the signature of trilateration, with highest precision close to the base and tip of the tool. A computational model of trilateration provided a good fit to the observed localization behavior. Importantly, model selection demonstrated that trilateration better explained each participant's behavior than an alternative model of localization. These results have important implications for how trilateration may be implemented by somatosensory neural populations. In sum, the present study suggests that tools are indeed embodied at a computational level, repurposing a fundamental spatial computation.

Introduction

The proposal that the brain treats a tool as if it were an extended limb (tool embodiment) was first made over a century ago (1). From the point of view of modern neuroscience, embodiment would entail that the brain reuses its sensorimotor computations when performing the same task with a tool as with a limb. There is indirect evidence that this is the case (2), such as the ability of tool-users to accurately localize where a tool has been touched (3) just as they would on their own body. Several studies have highlighted important similarities between tool-based and body-based tactile spatial processing (3-5), including on a neural level (6). Tool use also modulates somatosensory perception and action processes (7). While these findings suggest functional similarities between tools and limbs, direct evidence that body-based computational mechanisms are repurposed to sense and act with tools is lacking. The present study uses tool-extended sensing as a case study to provide the first computational test of embodiment.

Tactile localization on the body is often characterized by greater precision near body-part boundaries (e.g., its joints), a phenomenon called perceptual anchoring (8, 9). We recently found converging evidence that perceptual anchors are the signature of *trilateration* (10), a computation used by surveyors to localize an object within a map. To do so, a surveyor estimates the object's distance from multiple landmarks of known positions. When applied to

body maps (Figure 1A, bottom), a 'neural surveyor' localizes touch on a body part by estimating the distance between sensory input and body-part boundaries (e.g., the wrist and elbow for the forearm) and then integrating these estimates (11) to produce a Bayes-optimal location percept. Consistent with Weber's Law (12), we found that the noise in each distance estimate increased linearly as a function of distance (Figure 1B). Integrating them resulted in an inverted U-shaped noise profile across the surface, with the lowest noise near the boundaries and highest noise in the middle (i.e., perceptual anchoring).

In the present study, we investigated whether the somatosensory system *repurposes* trilateration to localize touch a tool (Figure 1A, top). In contrast with the body, tools are not innervated with sensory receptors and their central representation is unlikely to be somatotopic. Instead, the somatosensory system must 'tune into' a tool's mechanical response in order to use it as a sensor. Where a rod is touched is encoded in its vibratory response (3), which forms a multidimensional feature space that is general across rods and is isomorphic with tool-centered space. The brain could therefore repurpose trilateration by computing over this feature space (Figure 1C). In this formulation (see Methods for more details), its boundaries stand in for the boundaries of tool-centered space and distance estimates (Figure 1A) are computed within a neural representation of the feature space. If this is indeed the case, localizing touch on a tool would be characterized by an inverted U-shaped pattern of variable errors (i.e., perceptual anchors) across its surface.

Results

In the current experiment (n=38), we investigated whether tactile localization on a 60-cm hand-held rod is characterized by the U-shaped pattern of variability (Figure 1B) that is characteristic of trilateration. In two tasks, we measured participants' ability to localize an object that was actively contacted with a hand-held tool. In the *image-based task*, participants indicated the point of touch on a downsized drawing of the tool. In the *space-based task*, participants indicated the point of touch in external space. The latter task ensured that localization was not truncated by boundaries in the range of possible responses.

Consistent with prior results (3), we found that participants were generally quite accurate at localizing touch on the tool. Linear regressions comparing perceived and actual hit location found slopes near unity both the image-based task (mean slope: 0.93, 95% CI [0.88, 0.99]) and the space-based task (mean slope: 0.89, 95% CI [0.82, 0.95]). Analysis of the variable errors (2x6 repeated measures ANOVA) found a significant main effect of hit location ($F(5,185)=36.1$, $p<.001$) but no main effect of task ($F(1,37)=0.39$, $p=.54$) or an interaction

($F(5,185)=0.21$, $p=.96$). Crucially, the pattern of variable errors (Figure 1D) in both tasks displayed the hypothesized inverted U-shape, which was of similar magnitude to what we observed for touch on the arm (8, 10).

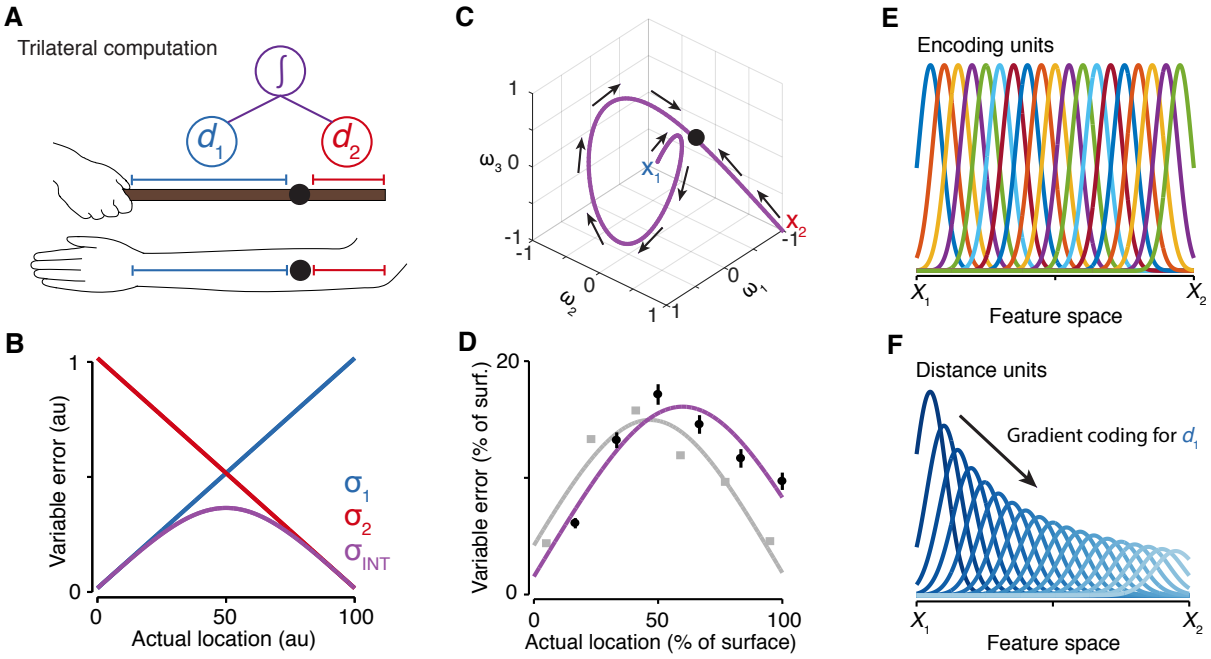

Figure 1. Model, results, and possible neural implementation

(A) The trilateral computation applied to the space of the arm (bottom) a hand-held rod (top). Distance estimates from sensory input (black) and each boundary (D1 and D2) are integrated (purple) to form a location estimate. (B) In our model, the noise in each distance estimate (D1, D2) increases linearly with distance. The integrated estimate forms an inverted U-shaped pattern. (C) The vibration-location feature space (purple) from handle (X_1) to tip (X_2). ω corresponds to a resonant frequency, the black dot corresponds to the hit location (as in panel A) within the feature space, and the arrows are the gradients of distance estimation. (D) The group-level variable error (black dots) is fit well by our probabilistic model of trilateration. Error bars are the 95% confidence intervals. These results are similar to what we previously observed for localization on the arm (gray line and squares) in (10). (E) Broadly tuned bell-shaped units encoding a feature space, which could be body space or vibration space. (F) The tuning of distance units forms a gradient (black arrow) away from one boundary of the feature space (here, X_1) with firing rate decreasing towards X_2. Higher firing rate corresponds to a more precise estimate (13). The other subpopulation's gradient goes in the other direction (not shown).

We next used computational modelling to confirm that the observed pattern of variable errors was indeed due to trilateration. We fit each participant's variable errors with a probabilistic model of optimal trilateration (Figure 1A-B) that was derived from its theoretical formulation (see, Methods). We compared the trilateration model to an alternative hypothesis: The

inverted U-shaped pattern is due to truncation at the boundaries of localization (12), which cuts off the range of possible responses and thus produces lower variability at these boundaries. We fit a boundary truncation model to directly compare to our trilateration model. Given the lack of a main effect of task and to increase statistical power, we collapsed across both tasks in this analysis.

Our computational model of trilateration provided a good fit to the variable errors observed during tactile localization on a tool (Figure 1D). We observed a high coefficient of determination at the level of individual participants (mean R^2 ±sem: 0.71±0.03; range: 0.29–0.95); indeed, 30 out of 38 participants had an R^2>0.6. In contrast, the R^2 of the boundary truncation model was substantially lower than the trilateration model (mean: 0.29±0.04; range: -0.19–0.71).

We next compared each model directly using the Bayesian Information Criteria (BIC). The BIC score for the trilateration model was lower in all 38 participants (mean±sd; Trilateration: 11.88±5.88; Truncation: 18.74±4.70). Statistically, 32 participants showed moderate evidence (ΔBIC>2) and 20 participants showed substantial evidence (ΔBIC>6) in favor of trilateration. In total, our results strongly suggest that, as with the body, touch on a tool is localized via trilateration.

Discussion

If tools are embodied by the sensorimotor system (2), we would expect that the brain repurposes its body-based sensorimotor computations to perform similar tasks with tools. Consistent with this hypothesis, we observed that localizing touch on the surface of a tool is characterized by perceptual anchors (9), suggesting that trilateration is the spatial computation employed by the somatosensory system to localize touch on body parts and tools alike (10). These findings go well-beyond prior research on embodiment (2) by directly identifying the computations that functionally unify tools and limbs. However, as discussed in the Introduction, there are important representational differences between limbs and tools. Our results therefore have important implications for how trilateration would be repurposed at a neural level for tool-extended sensing.

The computational goal of tactile localization—for bodies and tools—is to derive a spatial code from spatially-ambiguous neural activity (14), a problem that is solved by trilateration. Our previously neural network modelling (10) suggests that trilateration on the body can be implemented using two types of neural populations—perhaps located in S1—with distinct tuning properties: (i) An encoding population whose units are broadly tuned to a specific stimulus feature (e.g., arm space; Figure 1E); (ii) Two decoding subpopulations that compute distances

within this feature space. The units of these subpopulations have tuning properties with gradients that are anchored to the representational boundaries of the encoding population (e.g., wrist and elbow; Figure 1F). For example, the peak firing rate of one subpopulation might be highest near the elbow and lowest near the wrist. The proportionality between firing rate and perceptual precision (13) means that the activity of each subpopulation computes a distance estimate whose variance is distance-dependent (Figure 1B).

Our results suggest that this implementation is repurposed to perform the necessary transformations during tool-extended sensing. Given that the decoding population (Figure 1F) computes distance over the represented feature space in the encoding layer (Figure 1E), we propose that the somatosensory system applies them to compute distances within vibration space (Figure 1C). This requires that the human brain contains neural populations that encode for the full feature space of rod vibration. While very little is known about how these types of vibrations are represented by the somatosensory system, our modeling results and prior research (3, 6) suggest that there are neural populations that encode their properties. Indeed, recent evidence suggests that neurons in human S1 are tuned to the natural statistics of vibrations (15), consistent with the finding that S1 performs the initial stages of localization on a rod (6). The repurposing of the neural network architecture for trilateration may explain tool embodiment and the emergence of a shared spatial code between tools and skin.

Author Contributions

L.E.M., C.F., and A.F. conceived of the behavioral experiments; C.F. performed the behavioral experiments; L.E.M. and C.F. analyzed the behavioral data; L.E.M., A.F., and W.P.M. conceived of the computational model; F.V. and A.R. provided conceptual input; L.E.M., A.F. and W.P.M. wrote the original draft of the manuscript; All authors provided feedback on the manuscript and approved its final form.

Competing interests

The authors declare no competing interests.

Material and Methods

Theoretical formulation of trilateration

In the present section, we provide a theoretical basis for how trilateration can be applied to tactile localization on a tool, which will aid in our modelling (See next section). The *computational goal* of trilateration is to estimate the location of an object via distance estimates taken from vantage points of known positions. Applied to tactile localization, this amounts to estimating the location of touch by averaging over distance estimates taken from the boundaries of a sensory surface (Figure 1A), which are assumed to be internally represented. For a body part (e.g., forearm), these boundaries are often its joints (e.g., wrist and elbow) and lateral sides. For simple tools such as rods, this corresponds to their handle and tip.

For simplicity, let us consider only a single dimension and localization between the two boundaries of a tool-centered coordinate system (i.e., handle and tip). We propose that the somatosensory system only needs three spatial variables, $\{x_1, x_2, x_3\}$, to derive an estimate of the touch location $\tilde{L}$ on the tool. The variables x_1 and x_2 correspond to the proximal and distal boundaries, respectively. The variable x_3 corresponds to the sensory input.

The true representational nature of these variables depends on the type of sensory information that encodes where a tool was touched. We have previously argued that touch location is encoded in rod's resonant frequencies (3). Contact on its surface excites a location-specific combination of these resonant modes (between ~50–1000 Hz), called *vibratory motifs*. Together, these motifs can be mapped within a multidimensional feature space that forms a vibration-to-location isomorphism (Figure 1C). The task of the nervous system is therefore to learn the isomorphic mapping between the feature space of the vibrations and tool-centered space. We have previously found that disrupting the expected vibrations of a rod prevents accurate localization (3).The present model therefore assumes that the sensorimotor system of the user has indeed internalized this mapping.

We can therefore re-conceptualize x_1 and x_2 as estimates about the proximal and distal boundaries of the feature space. The variable x_3 can be re-conceptualized as an activation within a population of cortical neurons that are differentially sensitive to the statistics of the vibratory motifs, perhaps as early as S1 (15). Each variable x_i can be thus thought of as a Gaussian likelihood

$$p(x_i|X_i) = N\left(X_i, \sigma_i^2\right) \tag{1}$$

where the mean X_i corresponds to its true spatial position and the variance σ_i^2 corresponds to the uncertainty in its internal estimate. Here, X_1 and X_2 are internal estimates of the boundaries of the feature space and X_3 is the activation within the feature space, as encoded by a neural

population response. Crucially, the relationship between X_3 and L remains ambiguous without a decoder that knows the appropriate isomorphism and can place the population activity within the tool-centered coordinate system.

Trilateration solves this problem by calculating the distance between the proximal and distal boundaries of the coordinate system (x_1 and x_2, respectively) and the activation x_3 within the population of cortical neurons.

$$\begin{aligned} p(d_1|x_1, x_3) &= N(X_3 - X_1,\ \sigma_1^2(d_1)) \\ p(d_2|x_2, x_3) &= N(X_2 - X_3,\ \sigma_2^2(d_2)) \end{aligned} \tag{2}$$

where each distance estimate d_i corresponds to a Gaussian likelihood with a mean equal to the distance between X_3 and the respective landmark and a variance that scales with distance. That is, localization estimates are more precise when the touch is physically closer to a boundary than when it is farther away. As discussed in the Main Text, these estimates could be the result of populations of neurons tuned to distances within a vibration-related feature space. This distance-dependent noise is consistent with the Weber-Fechner law and is a consequence of how distance computation is implemented by a neural decoder.

Given the above distance estimates (Eq. 2), we can derive two estimates of touch location $\tilde{L}_i$ that are aligned within a common tool-centered coordinate system:

$$\begin{aligned} p(\tilde{L}_1|L) &= p(x_1|X_1) + p(d_1|x_1, x_3) \\ p(\tilde{L}_2|L) &= p(x_2|X_2) - p(d_2|x_2, x_3) \end{aligned} \tag{3}$$

These two location estimates can be used to derive a final estimate. However, given the presence of distance-dependent noise, the precision of each estimate will vary across the sensory surface (Figure 1B). Assuming the absence of a prior for touch location, the statistically optimal solution (i.e., maximum likelihood) is to integrate both estimates:

$$p(L|\tilde{L}_1, \tilde{L}_2) \propto p(\tilde{L}_1|L)p(\tilde{L}_2|L) \tag{4}$$

Here, the mean (μ_{INT}) and variance (σ_{INT}^2) of the integrated tool-centered posterior distribution depend on the means (μ_1 and μ_2) variances (σ_1^2 and σ_2^2) of the individual estimates:

$$\mu_{INT} = \left(\frac{\mu_1}{\sigma_1^2} + \frac{\mu_2}{\sigma_2^2}\right)\sigma_{INT}^2 \quad , \quad \sigma_{INT}^2 = \frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \tag{5}$$

The integrated posterior $p(L|\tilde{L}_1, \tilde{L}_2)$ thus reflects the maximum-likelihood estimate of touch location L. Given that the noise in each individual estimate scales linearly with distance, inte-

gration has the consequence of producing an inverted U-shaped pattern of variance (Figure 1B). This pattern of variability serves as a computational signature of trilateration, which we have observed for tactile localization on the arm and fingers (10). Our computational analyses implement this probabilistic model of trilateration (see below).

Participants

Forty right-handed participants (24 females, 23.65 ± 2.48 years of age) in total completed our behavioral experiments. Two participants were removed due to inability to follow task instructions, leaving thirty-eight in total to be analyzed. All participants had normal or corrected-to-normal vision and no history of neurological impairment. Every participant gave informed consent before the experiment. The study was approved by the ethics committee (CPP SUD EST IV, Lyon, France).

Experimental procedure: Active touch with a tool

During the task, participants were seated comfortably in a cushioned chair with their torso aligned with the edge of a table and their right elbow placed in a padded arm rest. The entire arm was hidden from view with a long occluding board. A 60 cm-long rod (handle length: 12-cm; cross-sectional radius: 0.75 cm) was placed in their right hand. This rod was either wooden (twenty-five participants) or PVC (thirteen participants). The arm was placed at a height necessary for a 1 cm separation between the object (see below) and the rod at a posture parallel with the table. On the surface of the table, an LCD screen (70 x 30 cm) lay backside down in the length-wise orientation; the edge of the LCD screen was 5 cm from the table's edge. The center of the screen was aligned with the participant's midline.

The task of participants was to localize touches resulting from active contact between the rod and an object (foam-padded wooden block). In an experimental session, participants completed two tasks with distinct reporting methods (order counterbalanced across participants). In the *image-based task*, participants used a cursor to indicate the corresponding location of touch on a downsized drawing of a rod (20 cm in length; handle to tip); the purpose of using a downsized drawing was to dissociate it from the external space occupied by the real rod. The drawing began 15 cm from the edge of the table, was raised 5 cm above the table surface, and was oriented in parallel with the real rod. The red cursor (circle, 0.2 cm radius) was constrained to move in the center of the screen occupied by the drawing. In the *space-based task*, participants used a cursor to indicate the corresponding location of touch within in an empty LCD screen (white background). The cursor was constrained to move along the vertical bisection of the screen and could be moved across the entire length of the screen.

It is critical to note that in this task, participants were forced to rely on somatosensory information about tool length and position as no other sensory cues were available to do so.

The trial structure for each task was as follows: In the 'Pre-contact phase', participants sat facing the computer screen with their left hand on a trackball. A red cursor was placed at a random location within the vertical bisection of the screen. A 'go' cue (brief tap on the right shoulder) indicated that they should actively strike the object with the rod. In the 'Localization phase', participants made their task-relevant judgment with the cursor, controlled by the trackball. Participants never received feedback about their performance. To minimize auditory cues during the task, pink noise was played continuously over noise-cancelling headphones.

The object was placed at one of six locations, ranging from 10 cm from the handle to the tip (10–60 cm from the hand; steps of 10 cm). The number of object locations was unknown to participants. In each task, there were ten trials per touch location, making 60 trials per task and 120 trials in total. The specific location for each trial was chosen pseudo-randomly. The entire experimental session took approximately 45 minutes.

The experiment started with a five-minute sensorimotor familiarization session. Participants were told to explore, at their own pace, how the tool felt to contact the object at different locations. They were instructed to pay attention to how the vibrations varied with impact location. Visual and auditory feedback of the tool and tool-object contact was prevented with a blindfold and pink noise, respectively. Participants were, however, allowed to hold the object in place with their left hand while contacting it with the tool but were not allowed to explore the rod with their right hand.

At the end of the space-based task, participants used the cursor to report where they felt the tip of the rod (aligned in-parallel to the screen). The judged location of the tip (mean: 56.5 cm; SEM: 1.62 cm) was very similar to the rod's actual length (i.e., 60 cm). It is critical to reiterate here that participants had never seen the rod prior up to this point of the experiment, and likely relied on proprioceptive feedback about its dimensions.

Data Analysis

Regression analysis

Prior to analysis, all judgments in the image-based task were converted from pixels of drawing space to percentage of tool space. All judgments in the space-based task were normalized such that their estimated tip location corresponded to 100% of tool space. We then used least-squares linear regression to analyze the localization accuracy. The mean localization judgment for each touch location was modelled as a function of actual object location. Accuracy

was assessed by comparing the group-level confidence intervals around the slope and intercept.

Modelling perceptual variability

Our model of trilateration in the somatosensory system assumes that the perceived location of touch is a consequence of the optimal integration of two independent location estimates, $\tilde{L}_1$ and $\tilde{L}_2$. This is exemplified in our formulation of trilateration (Equations 1-5). As discussed in the Main Text, trilateration predicts that noise in each estimate varies linearly as a function of the distance of touch from two landmarks (Equation 2; Figure 1B), corresponding to the handle and tip. For any location of touch L along a tactile surface, the variance in each landmark-specific location estimate $\tilde{L}$ can therefore be written as follows:

$$\sigma_1^2 = (\hat{\varepsilon}_1 + d_1\hat{\sigma})^2$$
$$\sigma_2^2 = (\hat{\varepsilon}_2 + d_2\hat{\sigma})^2 \tag{6}$$

in which $\hat{\varepsilon}$ is a landmark-specific intercept term that likely corresponds to uncertainty in the location of each landmark, d is the distance of touch location L from the landmark (Equations 2–3), and $\hat{\sigma}$ is the magnitude of noise per unit of distance. We assume that the noise term $\hat{\sigma}$ corresponds to a general property of the underlying neural network and therefore model it as the same value for each landmark. The distance-dependent noise for the integrated estimate is therefore:

$$\sigma_{INT} = \sqrt{\frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2}} \tag{7}$$

The three parameters in the model ($\hat{\sigma}$, $\hat{\varepsilon}_1$, and $\hat{\varepsilon}_2$) are properties of the underlying neural processes that implement trilateration and are therefore not directly observable. They must therefore be inferred using a reverse engineering approach, where they serve as free parameters that are fit to each participant's variable errors. We simultaneously fit the three free parameters to the data using non-linear least squares regression. Optimal parameter values were obtained through maximum likelihood estimation using the Matlab routine 'fmincon'. All modelling was done with the combined data from both localization tasks. R^2 values for each participant in each experiment were taken as a measure of the goodness-of-fit between the observed and predicted pattern of location-dependent noise.

References

1. H. Head, G. Holmes, Sensory disturbances from cerebral lesions. *Brain* **34**, 102-254 (1911).
2. M. Martel, L. Cardinali, A. C. Roy, A. Farne, Tool-use: An open window into body representation and its plasticity. *Cognitive Neuropsychology* **33**, 82-101 (2016).
3. L. E. Miller *et al.*, Sensing with tools extends somatosensory processing beyond the body. *Nature* **561**, 239-242 (2018).
4. S. Yamamoto, S. Kitazawa, Sensation at the tips of invisible tools. *Nature neuroscience* **4**, 979-980 (2001).
5. K. Kilteni, H. H. Ehrsson, Sensorimotor predictions and tool use: Hand-held tools attenuate self-touch. *Cognition* **165**, 1-9 (2017).
6. L. E. Miller *et al.*, Somatosensory Cortex Efficiently Processes Touch Located Beyond the Body. *Curr Biol* **29**, 4276-4283 e4275 (2019).
7. L. Cardinali *et al.*, Tool-use induces morphological updating of the body schema. *Current Biology* **19**, R478-479 (2009).
8. R. W. Cholewiak, A. A. Collins, Vibrotactile localization on the arm: Effects of place, space, and age. *Percept Psychophys* **65**, 1058-1077 (2003).
9. F. de Vignemont, A. Majid, C. Jola, P. Haggard, Segmenting the body into parts: evidence from biases in tactile perception. *Q J Exp Psychol (Hove)* **62**, 500-512 (2009).
10. L. E. Miller *et al.*, A neural surveyor to map touch on the body. *Proceedings of the National Academy of Sciences* (in press).
11. I. A. Clemens, M. De Vrijer, L. P. Selen, J. A. Van Gisbergen, W. P. Medendorp, Multisensory processing in spatial orientation: an inverse probabilistic approach. *J Neurosci* **31**, 5365-5377 (2011).
12. F. H. Petzschner, S. Glasauer, K. E. Stephan, A Bayesian perspective on magnitude estimation. *Trends Cogn Sci* **19**, 285-293 (2015).
13. W. J. Ma, J. M. Beck, P. E. Latham, A. Pouget, Bayesian inference with probabilistic population codes. *Nat Neurosci* **9**, 1432-1438 (2006).
14. M. R. Longo, E. Azanon, P. Haggard, More than skin deep: body representation beyond primary somatosensory cortex. *Neuropsychologia* **48**, 655-668 (2010).
15. L. Wang, J. M. Yau, Signatures of vibration frequency tuning in human neocortex. *bioRxiv* (2021).