

A neural surveyor in somatosensory cortex

Luke E. Miller¹⁻⁴, Cécile Fabio²⁻⁴, Rob van Beers^{1,6}, Alessandro Farnè^{2-5,7}, W. Pieter Medendorp^{1,7}

Affiliations:

1. Radboud University, Donders Institute for Brain, Cognition and Behaviour; Nijmegen, The Netherlands
2. Integrative Multisensory Perception Action & Cognition Team - ImpAct, Lyon Neuroscience Research Center, INSERM U1028, CNRS U5292; Bron, France
3. University of Lyon 1; Villeurbanne, France
4. Hospices Civils de Lyon, Neuro-immersion ; Bron, France
5. Center for Mind/Brain Sciences, University of Trento; Rovereto, Italy
6. Department of Human Movement Sciences, Vrije Universiteit Amsterdam; Amsterdam, The Netherlands
7. Equal contribution

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

Summary

Perhaps the most recognizable sensory map in all of neuroscience is the somatosensory homunculus. Though it seems straightforward, this simple representation belies the complex link between an activation in somatosensory Area 3b and the associated touch location on the body. Any isolated activation is spatially ambiguous without a neural decoder that can read its position within the entire map, though how this is computed by neural networks is unknown. We propose that somatosensory cortex implements multilateration, a common computation used by surveying and GPS systems to localize objects. Specifically, to decode touch location on the body, the somatosensory system estimates the relative distance between the afferent input and the body's joints. We show that a simple feedforward neural network which captures the receptive field properties of somatosensory cortex implements a Bayes-optimal multilateral decoder via a combination of bell-shaped (Area 3b) and sigmoidal (Areas 1/2) tuning curves. Simulations demonstrated that this decoder produced a unique pattern of localization variability between two joints that was not produced by other known neural decoders. Finally, we identify this neural signature of multilateration in actual psychophysical experiments, suggesting that it is a candidate computational mechanism underlying tactile localization.

Introduction

In the 18th century, surveyors in France completed the world's first topographically accurate map of an entire country. To do so, they relied on the computation of *triangulation*; given a precisely known distance between two baseline landmarks, the location of a third landmark could be computed from its angles of intersection with the baseline landmarks. Countries could utilize this simple geometric computation to accurately map the location of all landmarks in their borders (Figure 1A). This is also possible using *multilateration* (or *trilateration*, more specifically) where the known distance between multiple baseline landmarks is used to compute the location of another landmark. These computations are simple yet robust ways to localize objects and therefore still used in modern surveying and global position systems.

Geometric computations involving manipulating distances and angles are also employed by the nervous system of animals to localize and interact with objects in the environment. When navigating an environment, mammals readily return to their starting location by taking into account all computed distances and heading directions travelled (Figure 1B), a phenomenon known as path integration (Mittelstaedt and Mittelstaedt, 1980). Reaching to grasp a visible target is another behavior involving geometric computations (Figure 1C). To do so, the brain must compute a reach vector from distances derived from hand and target position signals (Beurze et al., 2006; Buneo et al., 2002; Flanders et al., 1992), involving transformations that take place in the frontal and parietal cortices (Burnod et al., 1999; Crawford et al., 2004; Medendorp et al., 2005; Pesaran et al., 2006).

Equally crucial to localizing objects in the environment is localizing objects on the personal space of the body. Despite over 180 years of research on the sense of touch (Weber, 1834), the computations underlying tactile localization remain largely unknown. Recent accounts have suggested that tactile localization requires two computational steps (Longo et al., 2010; Medina and Coslett, 2010). First, afferent input must be localized within a topographic map in

somatosensory cortex (Penfield and Boldrey, 1937). However, an activation within this map is not sufficient for localization since it alone explicitly represents little-to-no information about its associated position on the body surface. Localizing touch on the body therefore requires a neural decoder (Seung and Sompolinsky, 1993) that can “read” the topographic landscape of the population response within the map (Nicoletis et al., 1998) and reference this information to stored spatial representations of the body (Head and Holmes, 1911). However, given that the nature of these computations—and how they might be implemented by neural circuits—remain largely unknown, it is unclear whether the brain uses geometric computations to localize objects touching the body.

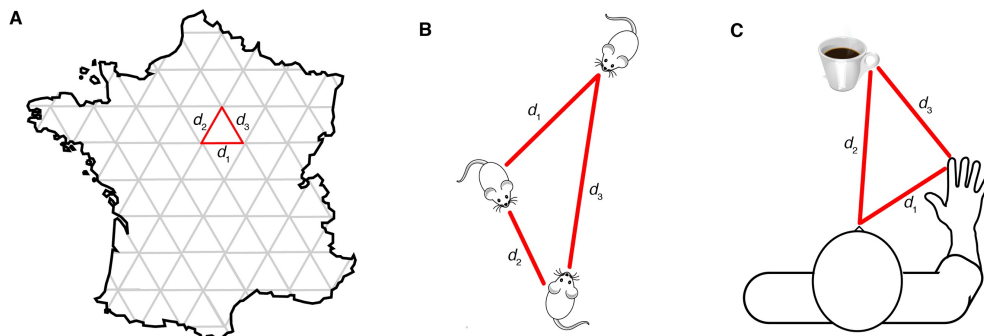


Figure 1. Examples of geometric computations

(A) Idealized example of how multi/trilateration can be used to map a country. After establishing a distance between two baseline landmarks (d_1), a third landmark can be localized by calculating its distance from each landmark individually (d_2 and d_3). Given these initial results (red triangle), a trilateration network (gray) is derived across the rest of the country. (B) Path integration: By computing over distances traveled (d_1 and d_2) a rat can calculate how much it needs to travel (d_3) to return to its starting position. (C) Visuomotor reaching: The magnitude of the reach vector (d_3) can be computed by evaluating over the eye-centered hand position (d_1) and object position (d_2).

We propose that, like a surveyor, the human brain employs *multi/trilateration* to localize an object in body-centered coordinates. To do so, this ‘neural surveyor’ uses simple arithmetic to calculate the relative distance between the location of the afferent input and the joints (the baseline landmarks; Figure 2A). Each landmark’s position in the coordinate system is likely

represented via both online proprioceptive feedback (Proske and Gandevia, 2012) and stored knowledge about the body's geometry (Longo and Haggard, 2010). In the present study, we provide multiple lines of evidence that the brain uses multi/trilateration to localize touch on the limb. We first develop a Bayesian formulation of it in the nervous system. We then develop a population-based neural network model that implements this computation, thus allowing us to identify its neural signatures. Simulations revealed that trilateration produces a unique pattern of localization variability across a limb. Finally, we identify this pattern in actual psychophysical experiments.

Results

A Bayesian formulation of trilateration

In multilateration, the distances between known locations are used to compute an unknown location. In the present paper we will focus mainly on *trilateration*, which requires calculating the distance between *three* unique locations in a common coordinate system. If we consider only a single dimension x , this simply amounts to subtracting each location from one another:

$$\begin{aligned}d_1 &= x_3 - x_1 \\d_2 &= x_2 - x_3 \\d_3 &= x_2 - x_1\end{aligned}\tag{1}$$

in which d_3 is the distance between two known locations, x_1 and x_2 , which serves as a baseline for calculating the unknown third location x_3 .

When applied to localizing a point of touch L on the limb (the x_3 in Eq. 1) the baseline d_3 corresponds to an internal representation of limb size, and x_1 and x_2 are the boundaries of a limb-centered coordinate system. For many limbs (e.g., the forearm), these boundaries—or landmarks—are represented by the position of the proximal and distal joints. Given peripheral input from mechanoreceptors, d_1 and d_2 must be measured via a neural surveyor that can

‘read’ a population response in a central somatotopic map Assuming noiseless signals (e.g., $x_1 = 0, x_2 = 100, x_3 = 75$) and decoding computations, we can re-write Equation 1 to produce two estimates of location:

$$\tilde{L}_1 = x_1 + d_1 \quad (2)$$

$$\tilde{L}_2 = x_2 - d_2$$

Because these estimates are defined within the same limb-centered coordinate system, a final estimate of location can be derived by taking their average, though in the case of noiseless signals both estimates are equal and therefore redundant (i.e., both \tilde{L}_1 and \tilde{L}_2 equal 75).

In the nervous system, however, noise is ubiquitous (Faisal et al., 2008). Sensory encoding is corrupted by receptor noise (Barlow, 1956; Lillywhite and Laughlin, 1979) which is compounded by computational operations performed by the nervous system (McGuire and Sabes, 2009; Shadlen and Newsome, 1998). Hence, an internal estimate of location cannot be taken as one specific estimate but rather as a probability distribution of locations with a variance that is often assumed to be Gaussian. The locations along dimension x are therefore specified as Gaussian random variables and the landmark-centered estimates \tilde{L} are best approximated as independent Gaussian likelihoods with distinct means and variances. Note that even though both estimates share a common feature (i.e., x_3), because they are derived from distinct landmark-specific distance estimates their noise is not necessarily correlated (see next section). Following Bayes’ theorem, touch location L given estimate \tilde{L} , that is the posterior distribution $p(L|\tilde{L})$, relates to:

$$p(L|\tilde{L}) \propto p(\tilde{L}|L)p(L) \quad (3)$$

in which $p(\tilde{L}|L)$ denotes the likelihood, representing probability density of the estimate \tilde{L} given the true location L , and $p(L)$ represents prior information about the location. Given that there

are two independent likelihoods, and assuming the prior over L is flat, the integrated posterior $p(L|\tilde{L}_1, \tilde{L}_2)$ corresponds to the product the two likelihood functions:

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

If the likelihoods are Gaussian distributions, the mean (μ_{INT}) and variance (σ_{INT}^2) of the integrated limb-centered posterior distribution depend on the means (μ_1 and μ_2) variances (σ_1^2 and σ_2^2) of the individual estimates:

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

The integrated posterior thus reflects the maximum-likelihood estimate of touch location L . The integrated variance is always smaller than the variance of either individual estimate; its mean can also be reformulated as the precision-weighted average of each estimate:

$$\mu_{INT} = w_1 \mu_1 + w_2 \mu_2 \quad (5)$$

whose weights depend upon their variances:

$$w_1 = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \quad , \quad w_2 = (1 - w_1)$$

This reformulation will be important below (see Eq. 10 in the next section) where we show that it can be used to extract a near-optimal location estimate on individual trials. Bayesian inference of this form has been demonstrated in a range of behaviors, such as visual object recognition (Kersten et al., 2004), multisensory integration (van Beers et al., 1999; Ernst and Banks, 2002), sensorimotor learning (Kording and Wolpert, 2004), and coordinate transformations (Clemens et al., 2011; McGuire and Sabes, 2009).

Trilateration, as formulated above, provides a computational mechanism to localize touch in intrinsic coordinates. In the next section, we explore how it could be implemented by the somatosensory cortex. We describe a simple population-based neural network model that can

trilaterate the location of touch from the population activity within the somatotopic map of Area 3b (Figure 2a). Importantly, determining the neural signatures of a trilateral computation for tactile localization will allow us to make predictions that can be validated using actual psychophysical data.

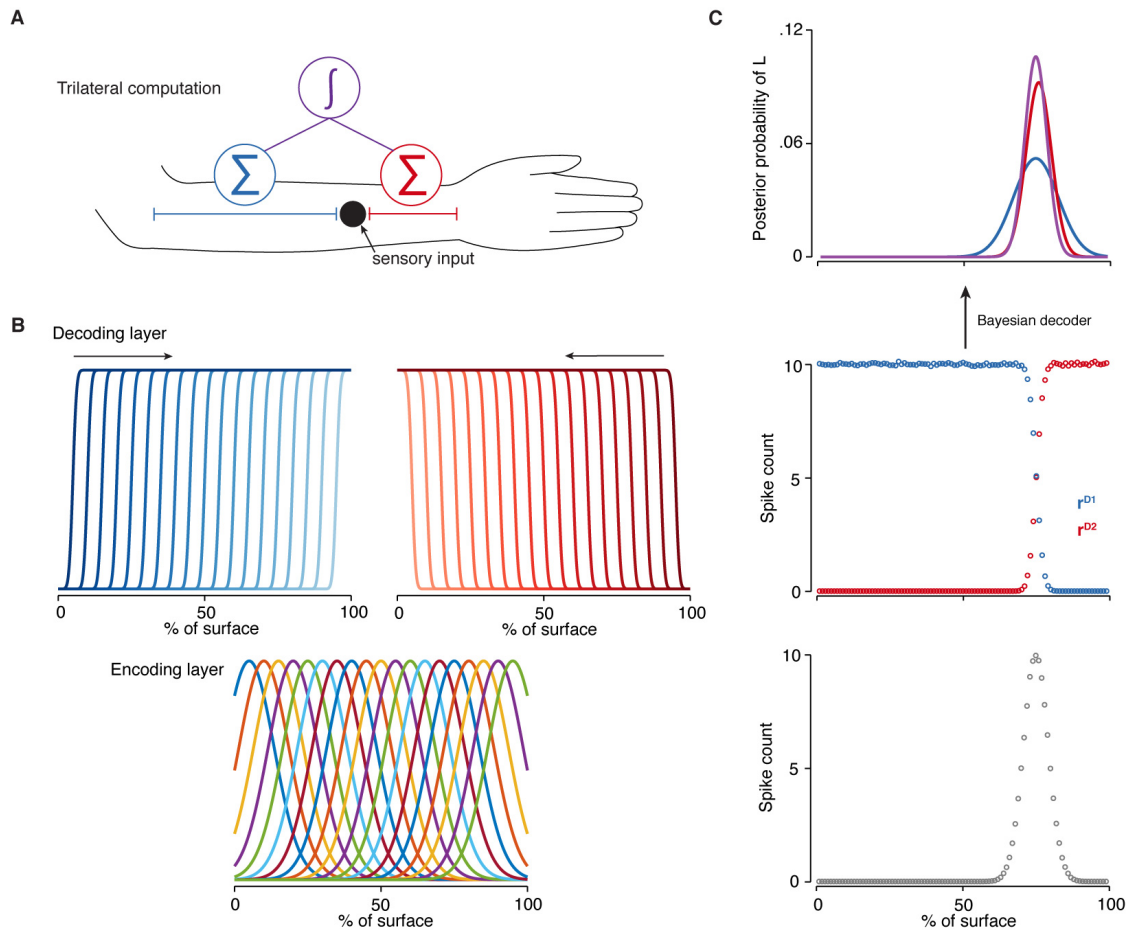


Figure 2. Neural network implementing trilateration

(A) Trilateral computation for tactile localization: The location of touch on the arm is computed by integrating two estimates of the distance of sensory input from each joint. (B) Neural network implementation of trilateration: (lower panel) the encoding layer is composed of bell-shaped tuning curves evenly spaced across the sensory surface, (upper panel) the decoding layer is composed of two subpopulation of evenly spaced sigmoidal tuning curves. The facing direction of each population is denoted by an arrow. How much of the surface is within each neuron's response field is coded by the luminance, with darker colors corresponding to neurons that cover more of the surface. (C) Activations for each layer of the network averaged over 5000 simulations. (lower panel) encoding layer; (middle panel) decoding layer; (upper panel) posterior probabilities of localization for each decoding subpopulation (blue and red) and their integration by the Bayesian decoder (purple).

Neural network model of a trilateration process in somatosensory cortex

We created a fully connected feedforward network composed of an encoding layer, a decoding layer, and a Bayesian decoder. The encoding layer was composed of 100 artificial neurons with Gaussian (bell-shaped) tuning curves f^E (Figure 2B), with likelihood functions $p(r_i^E|L)$ denoting the probability that location L caused r_i^E spikes in encoding neuron i . The likelihood function $p(r_i^E|L)$ can be modeled as a Poisson probability distribution, according to:

$$p(r_i^E|L) = \frac{e^{-f_i^E(L)} f_i^E(L)^{r_i^E}}{r_i^E!} \quad (6)$$

in which f_i^E is the tuning curve of neuron i . All Gaussian tuning curves had identical widths and were evenly spaced across the sensory surface, forming a somatotopic map of a limb. This configuration approximates the receptive field properties of cutaneous neurons in the forelimb representations in Area 3b (Dicarlo et al., 1998; Mountcastle and Powell, 1959). The population response of the encoding neurons is denoted by a vector $\mathbf{r}^E \equiv \{r_1^E, \dots, r_{100}^E\}$, where r_i^E is the spike count of neuron i whose mean and variance are both equal (i.e., the Fano factor is equal to 1). Following previous work (Pitkow and Angelaki, 2017), the population response \mathbf{r}^E can be thought of as representing a probability distribution over the stimulus.

The function of the decoding layer is to estimate the location of L in limb-centered coordinates given the population response \mathbf{r}^E in the somatotopic map. This was implemented using two independent subpopulations of 100 neurons that were fully connected to the encoding layer via synaptic weights whose values corresponded to a sigmoidal distribution with slopes facing away from their respective landmark (see Methods). These subpopulations were therefore composed of evenly-spaced sigmoidal tuning curves f^D , where some neurons were sensitive only to a small portion of the limb surface whereas others were sensitive to its entirety (Figure 2B). Given the relatively on-or-off nature of their response profiles, each neuron can be thought of as contributing a fixed amount of spikes to a final location estimate (see Eq. 7) and

therefore as coding for a single unit of tactile space. Each decoding subpopulation is therefore organized to implicitly encode the distance between input and a specific landmark.

Somatosensory areas 1 and 2 are good candidates for such distance computations since their cutaneous receptive fields match the necessary organization for their implementation. Both regions show a wide variety of receptive field sizes (Hyvarinen and Poranen, 1978), with the selectivity of neighboring neurons often spanning a continuous space from small to the full-limb coverage (Favorov and Whitsel, 1988). These receptive fields often cluster around or span one-or-more joint segments (Iwamura et al., 1983), suggesting that they are best characterized as sigmoidal. Further consistent with the function of the decoding layer, both Areas 1 and 2 receive direct projections from Areas 3a and 3b (Garraghty et al., 1990; Vogt and Pandya, 1978) and therefore integrate cutaneous and proprioceptive signals (Iwamura et al., 1983).

Trilateration requires calculating the distance between L and each landmark (Eq. 2). Given the sigmoidal shapes of the tuning curves in the decoding layer, this information is implicitly coded by the overall population response of the respective decoding populations $\mathbf{r}^D \equiv \{r_1^D, \dots, r_{100}^D\}$. Take an example where L is located at position 75 within the somatotopic map of the encoding layer (Figure 2A,C). For one decoding subpopulation, this point falls inside the maximal response field of neurons $i = [1 \dots 75]$ and outside for neurons $i = [76 \dots 100]$; this is reflected in the shape of its response \mathbf{r}^{D1} . Conversely, the subpopulation response \mathbf{r}^{D2} reflects the fact that this point falls inside the maximal response field of neurons $i = [1 \dots 25]$ and outside for neurons $i = [26 \dots 100]$. As can be seen in Figure 2C, the shape of each population response is a sigmoid that intersects at L . Since each neuron likely codes for a single unit of tactile space, calculating the distance between L and each landmark amounts to pooling the spike count of each subpopulation. We can therefore reformulate the location estimates \tilde{L} of Equation 2 as follows:

$$\tilde{L}_1^{(n)} = x_1 + g^{-1} \sum_{i=1}^{100} r_i^{D1} \quad (7)$$

$$\tilde{L}_2^{(n)} = x_2 - g^{-1} \sum_{i=1}^{100} r_i^{D2}$$

where x is the location of the corresponding landmark (in our network, $x_1 = 0$; $x_2 = 100$), g is a constant that converts spike count into units of tactile space, and ‘ (n) ’ denotes that this estimate comes from trial n . Furthermore, the quantity r_i^D corresponds to the spike count of decoder neuron i resulting from the weighted integration of all activity in the encoding layer: $r_i^{D1} = \mathbf{w}_i^D \cdot \mathbf{r}^E + \varepsilon$, where \mathbf{w}_i^D is the vector of synaptic weights connecting neuron i to the encoding layer, ‘ \cdot ’ is the dot product, and ε is the Poisson noise (Eq. 6).

Given that the spike count r for each neuron will vary from trial-to-trial, this equation is only valid for estimating the location of a single instance of touch. In an experiment, we wish to estimate behavioral performance over multiple trials. We can therefore rewrite Equation 7 to account for the firing rate statistics of each neuron (Eq. 6) as well as the uncertainty inherent in the location of each landmark x . The location estimate \tilde{L} can thus be written as a likelihood function:

$$p(\tilde{L}_1 | L, \mathbf{r}^E, s_p) = p(x_1 | s_p) + g^{-1} \sum_{i=1}^{100} p(r_i^{D1} | \mathbf{r}^E, L) \quad (8)$$

$$p(\tilde{L}_2 | L, \mathbf{r}^E, s_p) = p(x_2 | s_p) - g^{-1} \sum_{i=1}^{100} p(r_i^{D2} | \mathbf{r}^E, L)$$

where $p(r_i^D | \mathbf{r}^E, L)$ is the probability distribution of each decoding neuron’s response profile given touch at L and encoding population response \mathbf{r}^E , and $p(x | s_p)$ is the probability distribution of the location of each landmark given proprioceptive signals s_p (i.e., online feedback and stored representation).

Equation 8 demonstrates that calculating the distance between L and x amounts to summing the entire population response r^D . The variance of each estimate is therefore equivalent to the sum of the individual variances of each neuron's response to the input, divided by g . Because responses are restricted to neurons with non-zero response fields between L and x (Figure 2C), a greater distance means more active neurons in the population response and hence a location estimate with higher variance. A consequence of this pooling is therefore that the variance in each location estimate increases approximately linearly as a function of distance from the landmark. It is important to note that because this reflects a property of variance, distance-dependent noise is independent of the type of statistics that governs neural spiking (e.g., Poisson, Gaussian) or the Fano factor that links the mean and variance of the spike count.

The final step in trilateration amounts to deriving an optimal estimate of location from the two decoding distributions. Because of the linear relationship between noise and distance ($r=0.99$ in our simulations), a Bayesian decoder could theoretically perform optimal integration on each trial by using population spike count as a proxy for the variability in the estimates (Figure 2C). Because this is done on a single trial, we refer to this estimate as \tilde{L}_{INT} . Given Equation 5, averaging each estimate weighted by the inverse of their overall activity would approximate the maximum-likelihood estimate \tilde{L}_{INT} of location L and therefore:

$$\tilde{L}_{INT}(n) = \frac{\sum r^{D2}}{\sum r^{D1} + \sum r^{D2}} \left(x_1 + g^{-1} \sum_{i=1}^{100} r_i^{D1} \right) + \frac{\sum r^{D1}}{\sum r^{D1} + \sum r^{D2}} \left(x_2 - g^{-1} \sum_{i=1}^{100} r_i^{D2} \right) \quad (9)$$

It is important to note that integrating estimates of location (Eq. 4) assumes that they have independent sources of noise. However, given that both decoding layers are connected to the same encoding layer, they will both inherit its noise and could therefore be highly correlated. Unless the noise in the encoding layer can be removed from location estimates, the assumption of independence is violated.

In the present network, the problem of correlated noise is actually taken care of by Equation 9. Given Equation 8, any estimation bias due to noise in the encoding layer will have an opposite effect on both estimates (i.e., it's added to \tilde{L}_1 but subtracted from \tilde{L}_2). For example, let's imagine that the true value of \tilde{L} is 75 but both decoding subpopulations inherit a bias of 5 from the encoding layer. \tilde{L}_1 would be equal to 80 whereas \tilde{L}_2 would equal 70—thus, when these two values are integrated, the bias largely cancels out. Trilateration of the form in Equation 9 therefore removes any inherited biases from the encoding layer. As can be seen in Figure 2C the output of the trial-specific output of Bayesian decoder laid out in Equation 9 produces an estimate whose variance is lower than both landmark-specific estimates and is consistent with maximum likelihood integration.

Simulations identify plausible neural signatures of trilateration

So far we have provided a Bayesian formulation of trilateration and presented a plausible model of how this computation could be implemented in a simple feedforward network. We next investigated the localization behavior of this model by simulating single points of touch at each position within a limb-centered coordinate system (5000 iterations per location). The parameters of our initial simulations were based on known properties of somatosensory neurons (see Methods). These simulations included only two landmarks and therefore best reflect body parts with two joints, such as the forearm; more complicated body parts, such as the fingers, may require additional landmarks to accurately simulate.

Both subpopulations in the decoding layer (Figure 2B) were able to localize touch with minimal constant error (Figure 3A, upper panel), demonstrating that each produced unbiased estimates of location from the sensory input. However, as predicted from Equation 8, the variance in their estimates rapidly increased as a function of distance from each landmark (Figure 3B, upper panel). The pattern of location-dependent variability for each landmark-specific subpopulations was almost completely anti-correlated ($r = -.99$), forming an X-shaped pattern across

the surface of the limb. Noise thus renders the estimate of each subpopulation unreliable for most of the limb.

We next examined the output of the Bayesian decoder from Equation 9 (Figure 2C). As expected, integrating both estimates increased the reliability (Figure 3B, lower panel; for accuracy: Figure 3A, lower panel) of localization. Intriguingly, the variance of the Bayesian decoder's estimate formed an inverted U-shaped curve across the surface of the limb (Figure 3B, lower panel), with the lowest decoding variance near the landmarks and the highest decoding variance in the middle. This exact pattern of variance was also found when we computed the integration directly from the two simulated likelihoods (Eq. 4), demonstrating that our network optimally combines both estimates. These results demonstrate that near-optimal trilateration is possible on a single trial of touch—a necessary pre-requisite for our network to be plausible. It should be noted that by the term 'near-optimal' we do not mean being near the Cramér-Rao lower bound; rather, our Bayesian decoder is near-optimal in the sense of performing maximum-likelihood integration.

We then explored whether the observed inverted U-shaped variance profile is a signature of neural trilateration. To do so, we compared our network to three networks with a single encoding layer and distinct decoders (see Methods). The purpose of this comparison was not to compare overall decoding variance, but to assess whether the inverted U-shaped pattern of variance is specific to trilateration. In these networks, the location of touch was determined either by (1) the peak of the population response in the encoding layer (Riesenhuber and Poggio, 1999); (2) the minimal squared-error between the population response and location-specific templates (Deneve et al., 1999); or (3) the log-likelihood of the encoding population (Jazayeri and Movshon, 2006). In all cases, decoding variance was better than trilateration but did not vary as a function of stimulus location (Figure 3C). Furthermore, it is unclear whether these decoders could disambiguate the location of touch from the activation in the encoding layer (see Introduction).

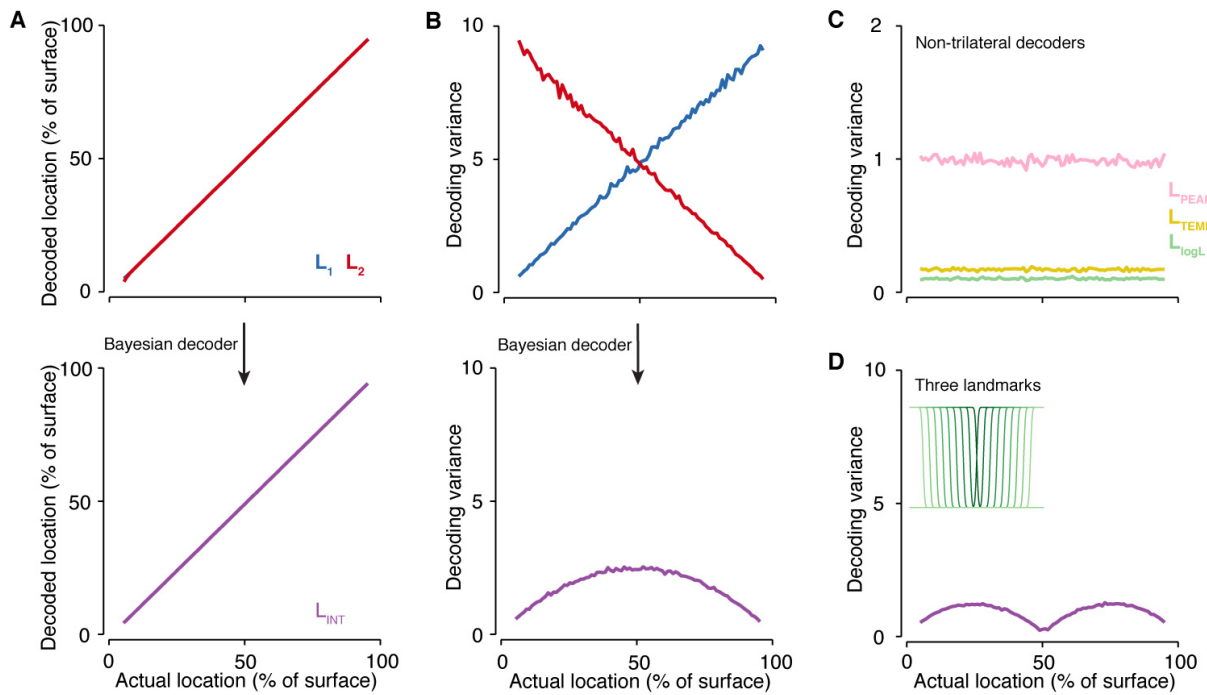


Figure 3. Simulation results and predictions

(A) Localization accuracy for the estimates of each decoding subpopulation (upper panel; L_1 , blue; L_2 , red) and after integration by the Bayesian decoder (lower panel; L_{INT} , purple). (B) Decoding variance (units: % of surface) for each decoding subpopulation (upper panel) increased as a function of distance from each landmark. Interestingly, integration by the Bayesian decoder (lower panel) led to an inverted U-shaped pattern across the surface. In these simulations, g was set to the maximum synaptic drive between the encoding and decoding layers in order to decorrelate both estimates (see Methods). (C) Decoding variance for three other decoders: a peak decoder (L_{PEAK}), a template-matching decoder (L_{TEMP}), and a log-likelihood decoder (L_{logL}). Decoding variance did not vary as a function of touch location for either of these decoders. (D) Results of the first prediction: Adding a third landmark in the middle of sensory surface led to an inverted W-shaped pattern of decoding variance. Inset: The receptive fields of the decoder subpopulation centered on this third landmark.

Why might the somatosensory cortex implement trilateration if decoding variance is not statistically optimal? One possibility has to do with the need for rapid decoding of location. Tactile localization in body-centered coordinates is completed with as little as 40 milliseconds of processing in primary somatosensory cortex (Miller et al., 2019a). Even with the high firing rate of somatosensory neurons (i.e., 50–100Hz; Bensmaia et al., 2008; Chapman and Ageranioti-Bélangier, 1991; Nicoletis et al., 1998; Reed et al., 2010), this corresponds to only one or two

spikes per layer. Indeed, we found that our trilateration network could decode touch location with a mean firing rate of a single spiker per layer, which corresponds to roughly 20-30 ms of processing in total. Decoding accuracy was as in Figure 3A and variance formed the inverted U-shaped pattern seen in Figure 3B, though it was higher. On the other hand, statistically optimal decoders—such as those implementing template matching—require recurrent processing and would therefore need much longer to implement (Deneve et al., 1999). Trilateration provides a means to tradeoff optimality with speed.

It is important to note that the inverted U-shaped pattern of variability is not a byproduct of our chosen network architecture but is due to the implemented computations. Per Equation 8, variance in an estimate will increase linearly as a function of the number of spikes contributing to that estimate. Any architecture where firing rate linearly codes for the distance from a landmark will therefore observe distance-dependent noise. Given that the optimal behavior is to integrate multiple noisy estimates (Equation 9), our observed pattern of variability will always be observed when trilateration computes tactile location on surfaces with two landmarks.

In all, our simulations demonstrate that observed location-specific pattern of variance is a signature of trilateration. We next conducted psychophysical experiments to confirm this pattern of variability in behavioral data. This is a necessary validation that our model is capturing something real in the computations underlying tactile localization. Importantly, observing an inverted U-shaped pattern of localization variability in our experiments would suggest that humans are ideal observers that trilaterate touch location near-optimally.

Trilateration explains tactile localization on the arm

In two psychophysical experiments, we investigated patterns of perceptual variability during tactile localization on the arm (see Methods). In Experiment 1 (n=11), participants localized points of touch passively applied to the volar surface of their forearm. In Experiment 2 (n=14), participants localized touch after they actively contacted an object with their forearm. Each

participant's responses were fit with a linear regression and the slope was taken as a measure of localization accuracy. Participants in Experiment 1 were highly accurate at localizing passive touches (slope: 1.04, 95% CI [1.00, 1.08]; Figure 4A, upper row). Similarly, participants in Experiment 2 were highly accurate at localizing active touches (slope: 1.06, 95% CI [0.99, 1.12]; Figure 4B, upper row). Importantly, in both experiments, we observed the expected inverted U-shaped pattern of variability (Figure 4A-B, bottom row). Thus, perceptual variability was dependent upon where the touch occurred, as predicted by trilateration.

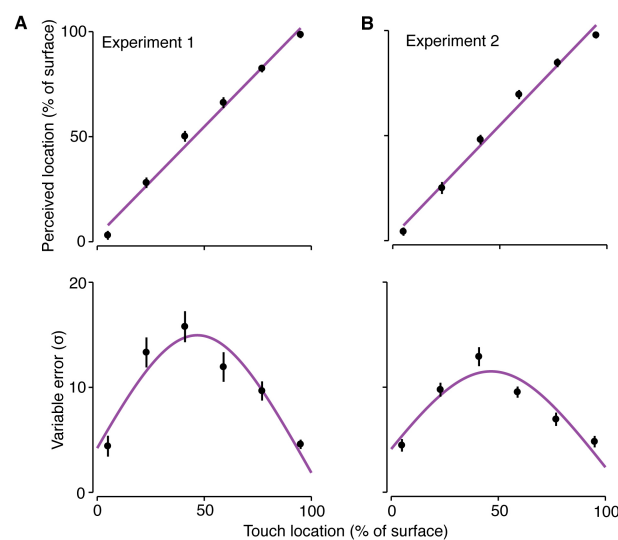


Figure 4. Results of behavioral experiments

The results of (A) Experiment 1 and (B) Experiment 2. Perceived location and perceptual variability as a function of touch location (0 = elbow; 100 = wrist). Tactile localization in both experiments was very accurate (upper rows). The fit line corresponds to a linear regression fit to the group-level data. The variable errors (lower rows) exhibited the expected signature of trilateration. The fit line corresponds to the computational model fit to the group-level data.

We used a reverse engineering approach (Clemens et al., 2011) to validate that the observed perceptual variability was due to trilateration. Because we cannot measure the parameters of trilateration directly, we inferred them by using least-squares regression to model each participant's variable error as a function of location (see Methods). Our regression model had three free parameters: one parameter that quantified the distance-dependent noise ($\hat{\sigma}$) and one

intercept parameter per landmark ($\hat{\epsilon}$). As in Equation 4, the model consisted of integrating landmark-specific patterns of noise to form a final pattern based on an optimal estimate of location.

Trilateration explained a large portion of the location-specific patterns of variability in each experiment. We found good fits for the group-level variable errors for both passive (Figure 4A, lower panel; $R^2=0.89$) and active touch (Figure 4B, upper panel; $R^2=0.86$). Importantly, trilateration provided a good fit ($R^2 > 0.5$) for every participant in Experiment 1 (mean \pm sem: 0.81 ± 0.04 ; range: $0.54-0.94$) and Experiment 2 (mean \pm sem: 0.80 ± 0.04 ; range: $0.57-0.98$). Figure 5 displays five randomly selected participants per experiment. The model fits for each participant in Experiments 1 and 2 are listed in Supplementary Tables 1 and 2, respectively.

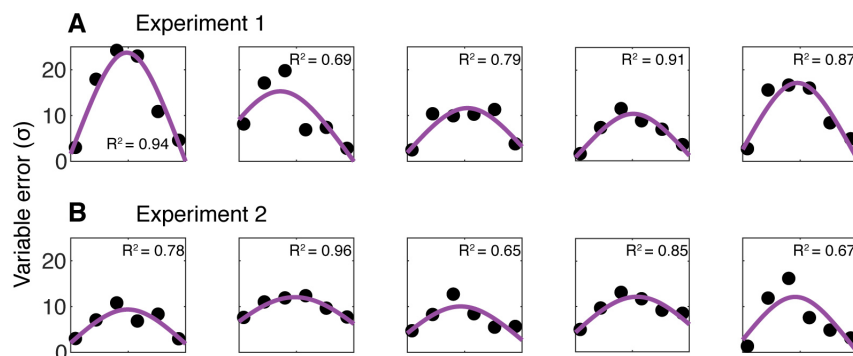


Figure 5. Model fitting confirms trilateration in our behavioral experiments

The variable errors in (A) Experiment 1 and (B) Experiment 2 were well-fit by a computational model of trilateration. Each rows shows five randomly chosen participants from each experiment and the corresponding goodness of fit.

Finally, we statistically compared the fit parameters for each experiment (Supplementary Tables 1–2). We found that the intercept terms ($\hat{\epsilon}_1$ and $\hat{\epsilon}_2$) did not significantly differ in Experiments 1 (paired t-test: $P > .2$, corrected) or 2 (paired t-test: $P > .2$, corrected). As the intercept terms likely reflect uncertainty about the location of landmarks, this is consistent with findings of similar proprioceptive sensitivity for the wrist and elbow (Fuentes and Bastian, 2010; Marini et al., 2016). It is also interesting to note that the noise parameter ($\hat{\sigma}$) in the active touch

experiment was lower than the passive touch experiment. However, given that the conditions of the stimulation were not closely matched across experiments, we chose not to statistically compare them in the present study.

In sum, we found that our participants' behavior reflected what is expected from ideal localizers. That is, their localization was on average unbiased and showed the pattern of variable error consistent with near-optimal trilateration. More generally, our behavioral experiments reveal broad agreement with the predictions of our population-based neural network model.

Model predictions

Thus far, we have shown that a neurally plausible implementation of trilateration accurately explains patterns of tactile localization in humans. We will now focus on several predictions made by our model of trilateration that can be tested empirically.

First, an inverted U-shaped pattern of variability will always be observed on either side of a landmark. This is somewhat trivial for individual limbs that share a single joint—such as the upper and lower arms—as trilateration is presumably implemented separately on either surface. However, it is non-trivial in cases where a salient object (e.g., jewelry) might serve as an artificial landmark. These objects can become an integrated part of the wearer's body representation (Aglioti et al., 1997) and may alter tactile perception. Figure 3D shows the effects of a third landmark centered in the middle of a sensory surface. As predicted, decoding variability now exhibits two hills, matching what has been observed when a vibrating object is placed in the middle of the arm (Cholewiak and Collins, 2003). How artificial landmarks might be instantiated in the plastic reorganization of somatosensory cortex is an open question, though continuous stimulation does lead to the reorganization somatosensory receptive fields (Dinse et al., 2003).

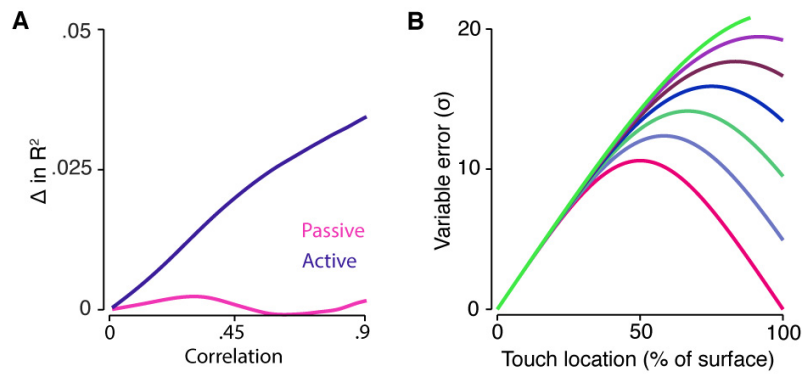


Figure 6. Model predictions

(A) Second prediction: For active touch (Experiment 2; purple), we observe an increase in the goodness of fit as a function of correlation between each location estimate. This is not observed for passive touch (Experiment 1; pink).
(B) Third prediction: Effect of increased variability in the second landmark location (The mark 100% on the surface). As variability increases (from 0% to 30% of tactile space; in steps of 5%), the inverted U-shaped pattern becomes more linear.

Correlated noise can have detrimental effects on population coding (Zohary et al., 1994) and the optimality of integration (Oruç et al., 2003). Our second prediction is that, due to differences in the magnitude of noise correlations, the integration of location estimates \tilde{L} during active touch will be suboptimal compared to passive touch. Though Equation 9 accounts for shared noise in each estimate arising in the encoding layer, other sources of noise may affect the decoding layer and therefore the optimality of the integration. Indeed, arm movements modulate the cutaneous responses of neurons in Area 1 (Jiang et al., 1991) and increases the magnitude of its noise correlations (Song and Francis, 2013). We tested this prediction in our own experiments by modelling trilateration with varying levels of correlation between estimates (see Methods). As can be seen in Figure 6A, accounting for correlated noise improved the fit of our model for the active touch experiment (Experiment 2) but not passive touch (Experiment 1).

Third, we predict that increasing the localization variability of a single landmark will modify the shape of the perceptual variability across the limb. Specifically, as the variability of a landmark increases, the inverted U-shaped pattern of variability will become less symmetrical

(Figure 6B). The variability of joint-based feedback can be modified, for example, by adding noise into the system via tendon vibration (Lackner, 1988). In the most extreme case of completely deafferenting a joint, variability would become linear. This might not be feasible, however, given that stored offline representations of body size also play a role in the position of a landmark within a body-centered coordinate system.

Fourth, we predict that a similar pattern of perceptual variability will be found when localizing touch on a hand-held tool. Indeed, humans can accurately localize where a tool has been touched (Miller et al., 2018; Yamamoto and Kitazawa, 2001). We recently found evidence that mechanisms in somatosensory cortex for localizing touch on an arm are re-used to localize touch on a tool (Miller et al., 2019a). Furthermore, tool use leads to lasting changes in somatosensory perception (Canzoneri et al., 2013; Cardinali et al., 2011; Miller et al., 2014, 2017) that are likely driven by plasticity in somatosensory cortex (Miller et al., 2019b; Schaefer et al., 2004). Given the high-degree of flexibility in the somatosensory system, we propose that the computation of trilateration is also used to localize touch on tools. Moreover, whether trilateration during tool sensing could involve mechanisms in somatosensory cortex—or is only implemented in higher-level frontoparietal regions—should be addressed using neurophysiology or functional neuroimaging.

Discussion

We proposed and tested the computation of multilateration as a candidate mechanism underlying tactile localization. Neural network modeling showed that this computation can be simply implemented in feedforward circuits of the somatosensory cortex, which integrate multiple location estimates into a single optimal surface-centered estimate. Simulations further indicated a location-dependent pattern of perceptual variability that reflects a signature of near-optimal trilateration. This signature was then found in two psychophysical experiments involving touch

on the arm. We conclude that multilateration is an important computation for localizing touch in the intrinsic coordinates of a sensory surface.

Multilateration provides a unified account of tactile ‘perceptual anchors’

Tactile perception varies across the surface of an individual body part. Perhaps the most striking example is the increased perceptual precision near the joints of the body (Brooks et al., 2019; Cholewiak and Collins, 2003; Cody et al., 2008; Knight et al., 2014; Longo, 2017; Pillsbury, 1895; De Vignemont et al., 2009), a phenomenon termed ‘perceptual anchoring’. Despite being first observed over 180 years ago by the psychophysicist E.H. Weber (Weber, 1834), the underlying reason why perception is tied to the proximity to joints is unknown. It is unlikely that ‘perceptual anchors’ have a peripheral origin since the receptive fields of mechanoreceptors are not more densely distributed near joints (Vallbo et al., 1995). Instead, they likely have a central explanation, such as the fact that joints function as category boundaries between somatosensory representations (Shen et al., 2018; De Vignemont et al., 2009). How this could be instantiated computationally has never been made explicit.

The present study suggests that the perceptual anchoring of tactile localization is a consequence of Bayesian trilateration in the somatosensory cortex. In our neural network, each decoding subpopulation is organized in reference to a specific landmark (i.e., a joint), consistent with their role as boundaries between different body-centered coordinate systems (i.e., category boundaries). Because these subpopulations represent the distance between touch and a specific landmark using a rate code, decoding noise increases linearly as a function of distance (Eqns. 8 & 9)—the closer touch is to a landmark, the more precisely it will be decoded. Therefore, integrating estimates with distance-dependent noise naturally leads to higher perceptual precision near landmarks. Our findings thus provide a unified computational explanation of perceptual anchoring in touch.

A spatial function for large somatosensory receptive fields

The shape of a neuron's tuning curve not only reflects principles of optimality (Smith and Lewicki, 2006) and statistics of the environment (Schwartz and Simoncelli, 2001) but also the functional role of the population they are embedded in (Salinas, 2006). Whereas bell-shaped tuning curves—as in our encoding layer—may be optimal for representing locations within a topographic map, sigmoid-shaped tuning curves—as in our decoding layer—are optimal for representing specific values (Sanger, 2003). Consistent with this, the Fisher information of sigmoidal tuning curves is centered around a single point in feature space (Yarrow and Series, 2015). We therefore reasoned that sigmoid-shaped tuning curves provide a natural means by which the somatosensory cortex could represent a unit of tactile space.

Indeed, our neural network model demonstrated that, given a specific organization, a population of sigmoid-shaped neurons can implicitly calculate the distance between touch and a landmark. While it is unclear whether such an organization actually exists in the somatosensory cortex, it is well known that receptive field sizes of cutaneous neurons in somatosensory Areas 1 and 2 vary on a continuum from small to large (Hyvarinen and Poranen, 1978). For example, two neighboring neurons in the forearm region of Area 1 may code for a small region by the wrist and the entire forearm (Favorov and Whitsel, 1988). Given the broad tuning of these neurons, sigmoid-shaped tuning curves provide a good approximation for their coding properties. Despite the popular belief that large receptive fields lack spatial discrimination, our results are consistent with previous evidence to the contrary (Foffani et al., 2008; Nicolelis et al., 1998) in demonstrating that they can precisely code spatial information at the population level.

Neural implementation of the Bayesian decoder

In our neural network, the distance between touch and each landmark (Eqs. 1 & 2) is represented by the pooled activity of two subpopulations of decoding neurons (Eq. 8). By weighting

each subpopulation by its overall activity, the Bayesian decoder could estimate the location of touch near-optimally (Eq. 9). Unlike the encoding and decoding layers, we left the implementation of the Bayesian decoder largely unspecified. There are therefore several open questions about the nature of this computational step.

First, it is unclear whether the pooling of activity in each subpopulation would be implemented by single neurons or an entire neuronal population. Single neurons in area LIP (Shadlen and Newsome, 2001) are known to integrate information from an entire sensory population (Beck et al., 2008). While this has never been directly demonstrated for somatosensory processing, several somatosensory regions have neurons with receptive fields covering an entire limb (Favorov and Whitsel, 1988; Iwamura et al., 1983; Mountcastle et al., 1975; Sakata et al., 1973), suggesting that they pool across a population of tactile neurons as formulated in Equations 8 and 9. Alternatively, pooling could be implemented by an entire population of neurons (Seung and Sompolinsky, 1993). Indeed, it is often argued that Bayesian inference is best implemented at the population level (Ma et al., 2006), such as with basis functions with multi-dimensional attractors (Deneve et al., 2001). This should be addressed in future research.

Second, it is unclear whether the Bayesian decoding would be implemented in somatosensory cortex or higher-order associative regions, such as posterior parietal cortex. Low-level sensory regions can implement Bayesian inference. For example, auditory spatial cues are optimally integrated by the owl midbrain during sound source localization (Cazettes et al., 2016; Fischer and Peña, 2011). It is therefore possible that a third subpopulation of neurons in somatosensory Areas 1 or 2 could optimally integrate signals from both decoding subpopulations. Instead, Bayesian decoding might be performed by somatosensory regions in the posterior parietal cortex (Breveglieri et al., 2008; Duhamel et al., 1998; Mountcastle et al., 1975; Seelke et al., 2012), which are known to play a role in tactile localization (Reed et al., 2005). These regions are likely important for referencing sensory signals to stored knowledge of body size (Ehrsson et al., 2005), an important component of multilateration (Eqns. 1 & 2). Most likely,

Bayesian decoding during tactile localization is implemented by both feedforward signals in somatosensory cortex and feedback signals from posterior parietal cortex (Jones et al., 2007).

Is multilateration a general spatial computation?

Whether multilateration is involved in other forms of spatial cognition is unclear. However, its equations map onto other known distance-based geometric computations implemented in the nervous system. Consider the example of reaching to grasp a coffee mug (Figure 1C). Per Equation 1, the magnitude of the reach vector (d_3) can be computed simply by subtracting the distance between the eyes and the coffee mug (d_2) from the distance between the hand and the eyes (d_1). This operation is thought to be performed by neurons in posterior parietal cortex (Buneo, et al., 2002, Beurze, et al., 2006).

As shown in our study, when distance is encoded in the overall firing rate of a population, noise in each location estimate scales linearly with distance (Equation 8). Given that distance is treated as a magnitude in this computation, distance-dependent noise is consistent with Weber's Law and may therefore be a general feature of multilateral computations. This suggests that patterns of distance-dependent noise can serve to identify multilateration in other domains. This appears to be the case in allocentric vision, where noise in the estimated location of an object is dependent upon its distance from landmarks in the scene (Aagten-Murphy and Bays, 2019). Further, the dominant source of error during path integration is noise that accumulates with distance travelled (Stangl et al., 2020). Interesting, similar errors are found for path integration in the tactile domain (Fardo et al., 2018; Moscatelli et al., 2014), suggesting that they may involve multilateration as well.

Conclusion

In sum, our results suggest that, like a surveyor, the somatosensory system employs near-optimal multilateration to localize a tactile stimulus. This computation is likely implemented, at least partially, in somatosensory cortex. Future work should address how multilateration can

be extended to cases of localization in two (Mancini et al., 2011) or three dimensions (Azañón et al., 2016), as well as when touch occurs under more dynamic contexts (Maij et al., 2013). Furthermore, it remains to be seen to what extent other spatial behaviors—such as path integration, allocentric vision, and reaching—should be reformulated as implementing multilateration.

Material and Methods

Neural network modeling

Network parameters

We devised a simple two layer feedforward neural network that implements trilateration to localize touch on a sensory surface. Each layer of the network was composed of 100 artificial neurons whose preferred locations were evenly spaced across the sensory surface (Figure 2). The space of the surface was always modelled in terms of percentage (i.e., 0-100% of the surface). The properties of the ‘neurons’ in each layer approximated important aspects of actual neurons found in the somatosensory cortex (Delhaye et al., 2018).

In the encoding layer, neurons were modelled as narrowly tuned Gaussian tuning curves f^E (FWHM: 5% of the surface). Each neuron in the decoding layer was fully connected to each neuron in the encoding layer via a synaptic weight vector w^D . The synaptic weights formed a sigmoidal distribution with values ranging from 0 to 0.2, a standard deviation of 1% of the surface, and a central point that was neuron-dependent and evenly spaced across the space of the encoding layer. The maximum value of 0.2 was chosen because, given our network configuration, it produced near-identical peak spike counts in neurons in the encoding and decoding layers. Given this connectivity, the tuning curves in each decoding sub population f^D were sigmoidal in shape (Figure 2).

As discussed at length in the Main Text, neurons in the encoding and decoding layers approximated the shapes of tuning curves found in Areas 3b and 1/2, respectively. In our initial simulations, the mean peak spike count of these neurons was set to 10 spikes. Given that high peak firing rates (50-100Hz) are observed in Areas 3, 1, and 2 (Bensmaia et al., 2008; Chapman and Ageranioti-Bélanger, 1991; Nicolelis et al., 1998; Reed et al., 2010), this corresponds to ~100-200 ms of processing. We made the simplifying assumption that the variability in firing rates corresponds to a Poisson process with a Fano factor of 1. While neurons

measured *in vivo* do exhibit a range of Fano factors, the mean Fano factor in several brain regions is close to 1 following stimulus onset (Churchland et al., 2010) and thus this assumption is not without merit.

To convert units of spikes in the decoding subpopulations into units of distance—and to account for shared bias in the encoding layer—we set g to equal the maximum synaptic drive from the encoding to the decoding layer:

$$g = \operatorname{argmax}_r W^D \cdot \mathbf{r}^E$$

in which W^D is the matrix of synaptic weights for each decoding subpopulation, \mathbf{r}^E is the population response in the encoding layer, and ‘ \cdot ’ is the dot product. This is one possible value of g and is largely used for convenience to uncorrelate both decoding estimates (Supplementary Figure 1B). However, it is important to note that its tuning is largely unnecessary since any bias is removed by the weighted integration of both estimates (see Main Text; Equation 9).

Simulations for a multilateral decoder

To investigate the neural consequences of a trilateral computation, we simulated 5000 instances of touch at each location on the sensory surface using the above network. Our initial simulations used a mean peak spike count of 10 spikes for neurons in both layers. Simulations were also performed across a wide range of mean peak spike counts, ranging from 1 to 100 spikes. In all cases the same decoding accuracy and shape of variability was found. It should also be noted that the results did not depend on identical peak spike counts in both layers, and instead this was done for convenience. The equations underlying trilateration in our network (Eqns. 6–9) and all simulations were implemented using custom code in Matlab.

Simulating alternative decoders

We further simulated the neural consequences of three different decoding schemes in order to compare them with the trilateral decoder. Each decoder contained only a single encoding

layer, which was identical to what is described above. These comparisons were not done to determine which decoder had the highest decoding accuracy but instead to investigate whether their decoding variance varied as a function of touch location. Each decoder estimated the location of touch \tilde{L} from the population response \mathbf{r}^E in the encoding layer (mean peak spike count: 10 spikes). All simulations were performed using custom Matlab code. The equation underlying each decoding scheme is as follows:

Log-likelihood decoder: For this decoder, the estimated location of touch \tilde{L} corresponds to the location in the encoding population that maximizes the log-likelihood $\log \mathcal{L}$ for a given tactile input at L (Jazayeri and Movshon, 2006; Ma et al., 2006), and is calculated via the following equation:

$$\tilde{L} = \underset{L}{\operatorname{argmax}} \log \mathcal{L}(L) = \sum_{i=1}^{100} r_i^E \log f_i^E(L)$$

where \tilde{L} is the location within the encoding layer that maximizes the log-likelihood $\log \mathcal{L}$ for a given tactile input at L , computed by summing the product between each neuron's response r^E and the log of its tuning curve.

Winner-take-all decoder: For this decoder, the estimated location of touch \tilde{L} corresponds to neuron in the encoding population with the highest spike count.

$$\tilde{L} = \underset{L}{\operatorname{argmax}} \mathbf{r}^E(L)$$

in which \mathbf{r}^E is the population response of the encoding neurons following touch at location L .

Template matching decoder: For this decoder, the estimated location of touch \tilde{L} corresponds to the location-specific Gaussian-shaped template f^T (FWHM: 5% of surface) whose shape best matches the shape of the population response \mathbf{r}^E :

$$\tilde{L} = \underset{L}{\operatorname{argmin}} \sum_{i=1}^{100} (r_i^E - f_L^T(i))^2$$

in which r_i^E is the spike count of the i^{th} neuron to the location of touch L and $f_L^T(i)$ is the value of the L^{th} template corresponding to the location i . This equation minimizes the mean squared error between population response and the location-specific templates and therefore reflects maximum likelihood estimation.

Behavioral experiment

Participants

Twenty-seven right-handed participants in total completed our behavioral experiments. Twelve participated in Experiment 1 (8 females, 24 ± 0.63 years of age) and fifteen in Experiment 2 (9 females, 24.2 ± 0.56 years of age). One participant from both Experiments 1 and 2 was removed due to inability to follow task instructions. 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).

Experiment 1: Passive touch on the forearm

During the task, participants were seated comfortably in a cushioned chair with their torso aligned with the edge of a table and their right arm resting on the table top behind an occluding board. 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 applied passively to their arm. In an experimental session, participants completed two tasks with distinct reporting methods (order counterbalanced across participants; combined in the results of the Main Text). In the 'drawing

task', participants used a cursor to indicate the corresponding location of touch on a downsized drawing of a human arm (12 cm in length; forearm and hand); the purpose of using a downsized drawing was to dissociate it from the external space occupied by the real arm. 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 their real arm. The red cursor (circle, 0.2 cm radius) was constrained to move in the center of the screen occupied by the drawing. In the 'external space 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 proprioceptive information about their arm position as no other sensory cues were available to do so.

In each experiment, unknown to the participant, there were six evenly-spaced touch locations between 5% to 95% the length of the arm (18% intervals; elbow-to-wrist; mean arm length: 23.9 ± 0.4 cm). 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 trial structure for each task was as follows: In the 'Pre-touch 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 cue (tap on the right shoulder) indicated an impending touch on the volar surface of the forearm. Touch was applied with a von Frey microfilament at a suprathreshold level of stimulation (180 g of force) for approximately one second. In the 'Localization phase', participants made their task-relevant judgment with the cursor, controlled by the trackball. Participants never received feedback about their performance.

Experiment 2: Active touch with the forearm

The experimental procedures were identical to Experiment 1 with the following exceptions. Throughout the experiment, the participant's right elbow was placed upright in a padded support with the entire arm hidden from view behind a long occluding board. The task of participants was to localize touches that resulted from active contact between their right arm (mean arm length: 23.5 ± 0.5 cm) and an object (rounded tip plastic cylinder; 1 mm radius). The arm was placed at a height necessary for a 1 cm separation between the object and the forearm at a posture parallel with the table. To minimize auditory cues during the task, pink noise was played continuously over noise-cancelling headphones. During each trial, a 'go' cue (tap on the right shoulder) indicated that they should actively bring their forearm from to its upright posture into contact with the object, placed at one of six locations (5% to 95% of forearm length, evenly spaced). Participants were instructed to attempt to hit the object with the same speed and force across trials, though this was not measured. The number of trials and reporting methods were as in Experiment 1.

Data analysis

Localization accuracy

We used least-squares linear regression to analyze the localization accuracy of each task in each experiment. 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 to zero and one (Miller et al., 2018). To standardize the data for each participant and each surface, all judgments were converted to a percentage of total surface length. For the drawing task, this amounted to converting judged locations on the drawing into percentage of drawing length. For the external space task, this amounted to converting judged locations in the space of the screen into percentage of actual surface or perceived rod length. In the Main Text, we collapsed localization analysis across both localization

tasks for the passive and active datasets. This is because performance on the drawing and external tasks was nearly identical: passive dataset (drawing vs. external; slope: 1.03 vs. 1.06) and active dataset (drawing vs. external; slope: 1.04 vs. 1.08).

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 Bayesian formulation of trilateration (Equations 4-5) as well as our neural network implementation (Equations 8-9). As discussed extensively 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 8), corresponding to the elbow and wrist for the arm (Experiments 1 and 2). 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:

$$\sigma_1^2 = (\hat{\varepsilon}_1 + d_1 \hat{\sigma})^2$$

$$\sigma_2^2 = (\hat{\varepsilon}_2 + d_2 \hat{\sigma})^2$$

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 1 and 2), and $\hat{\sigma}$ is the magnitude of noise per unit of distance. Note that because the noise term $\hat{\sigma}$ corresponds to a general property of the underlying neural network (Equation 8), it is the same for each landmark. The distance-dependent noise for the integrated estimate is therefore:

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

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 (Clemens et al., 2011), where they serve as free parameters that are fit to each participant's variable errors. The equations from (Oruç et al., 2003) were used to account for correlated variability in the integration process (see Model Predictions). 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' (Clemens et al., 2011; McGuire and Sabes, 2009; Niehof et al., 2019). 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. The values of the fit parameters were compared within experiment using paired t -tests.

References

- Aagten-Murphy, D., and Bays, P.M. (2019). Independent working memory resources for egocentric and allocentric spatial information. *PLoS Comput. Biol.* *15*, 1–20.
- Aglioti, S., Smania, N., Manfredi, M., and Berlucchi, G. (1997). Disownership of left hand and objects related to it in a patient with right brain damage. *Neuroreport*.
- Azañón, E., Mihaljevic, K., and Longo, M.R. (2016). A three-dimensional spatial characterization of the crossed-hands deficit. *Cognition* *157*, 289–295.
- Barlow, H.B. (1956). Retinal noise and absolute threshold. *J. Opt. Soc. Am.*
- Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., and Pouget, A. (2008). Probabilistic Population Codes for Bayesian Decision Making. *Neuron* *60*, 1142–1152.
- van Beers, R.J., Sittig, A.C., and Denier Van Der Gon, J.J. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *J. Neurophysiol.* *81*, 1355–1364.
- Bensmaia, S.J., Denchev, P. V., Dammann, J.F., Craig, J.C., and Hsiao, S.S. (2008). The representation of stimulus orientation in the early stages of somatosensory processing. *J. Neurosci.* *28*, 776–786.
- Beurze, S.M., Van Pelt, S., and Medendorp, W.P. (2006). Behavioral reference frames for planning human reaching movements. *J. Neurophysiol.* *96*, 352–362.
- Breviglieri, R., Galletti, C., Monaco, S., and Fattori, P. (2008). Visual, somatosensory, and bimodal activities in the macaque parietal area PEc. *Cereb. Cortex* *18*, 806–816.
- Brooks, J., Seizova-Cajic, T., and Taylor, J.L. (2019). Biases in tactile localization by pointing: Compression for weak stimuli and centering for distributions of stimuli. *J. Neurophysiol.* *121*, 764–772.
- Buneo, C.A., Jarvis, M.R., Batista, A.P., and Andersen, R.A. (2002). Direct visuomotor transformations for reaching. *Nature* *416*, 632–636.
- Burnod, Y., Baraduc, P., Battaglia-Mayer, A., Guigon, E., Koechlin, E., Ferraina, S., Lacquaniti, F., and Caminiti, R. (1999). Parieto-frontal coding of reaching: An integrated

framework. *Exp. Brain Res.*

Canzoneri, E., Ubaldi, S., Rastelli, V., Finisguerra, A., Bassolino, M., and Serino, A. (2013). Tool-use reshapes the boundaries of body and peripersonal space representations. *Exp. Brain Res.*

Cardinali, L., Brozzoli, C., Urquizar, C., Salemme, R., Roy, A.C., and Farnè, A. (2011). When action is not enough: Tool-use reveals tactile-dependent access to Body Schema. *Neuropsychologia.*

Cazettes, F., Fischer, B.J., and Peña, J.L. (2016). Cue reliability represented in the shape of tuning curves in the owl's sound localization system. *J. Neurosci.* *36*, 2101–2110.

Chapman, C.E., and Ageranioti-Bélanger, S.A. (1991). Discharge properties of neurones in the hand area of primary somatosensory cortex in monkeys in relation to the performance of an active tactile discrimination task - I. Areas 3b and 1. *Exp. Brain Res.* *87*, 319–339.

Cholewiak, R.W., and Collins, A.A. (2003). Vibrotactile localization on the arm: Effects of place, space, and age. *Percept. Psychophys.* *65*, 1058–1077.

Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nat. Neurosci.* *13*, 369–378.

Clemens, I.A.H., de Vrijer, M., Selen, L.P.J., van Gisbergen, J.A.M., and Medendorp, W.P. (2011). Multisensory processing in spatial orientation: An inverse probabilistic approach. *J. Neurosci.* *31*, 5365–5377.

Cody, F.W.J., Garside, R.A.D., Lloyd, D., and Poliakoff, E. (2008). Tactile spatial acuity varies with site and axis in the human upper limb. *433*, 103–108.

Crawford, J.D., Medendorp, W.P., and Marotta, J.J. (2004). Spatial transformations for eye-hand coordination. *J. Neurophysiol.* *92*, 10–19.

Delhaye, B.P., Long, K.H., and Bensmaia, S.J. (2018). Neural Basis of Touch and Proprioception in Primate Cortex. *Compr. Physiol.*

Deneve, S., Latham, P.E., and Pouget, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nat. Neurosci.* *2*, 740–745.

Deneve, S., Latham, P.E., and Pouget, A. (2001). Efficient computation and cue integration

with noisy population codes. *Nat. Neurosci.* 4, 826–831.

Dicarlo, J.J., Johnson, K.O., and Hsiao, S.S. (1998). Structure of receptive fields in area 3b of primary somatosensory cortex in the alert monkey. *J. Neurosci.* 18, 2626–2645.

Dinse, H.R., Ragert, P., Pleger, B., Schwenkreis, P., and Tegenthoff, M. (2003). Pharmacological modulation of perceptual learning and associated cortical reorganization. *Science* (80-.).

Duhamel, J.R., Colby, C.L., and Goldberg, M.E. (1998). Ventral intraparietal area of the macaque: Congruent visual and somatic response properties. *J. Neurophysiol.*

Ehrsson, H.H., Kito, T., Sadato, N., Passingham, R.E., and Naito, E. (2005). Neural substrate of body size: Illusory feeling of shrinking of the waist. *PLoS Biol.*

Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433.

Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303.

Fardo, F., Beck, B., Cheng, T., and Haggard, P. (2018). A mechanism for spatial perception on human skin. *Cognition* 178, 236–243.

Favorov, O., and Whitsel, B.L. (1988). Spatial organization of the peripheral input to area 1 cell columns. I. the detection of “segregates.” *Brain Res. Rev.* 13, 25–42.

Fischer, B.J., and Peña, J.L. (2011). Owl’s behavior and neural representation predicted by Bayesian inference. *Nat. Neurosci.* 14, 1061–1066.

Flanders, M., Tillery, S.I.H., and Soechting, J.F. (1992). Early stages in a sensorimotor transformation. *Behav. Brain Sci.*

Foffani, G., Chapin, J.K., and Moxon, K.A. (2008). Computational role of large receptive fields in the primary somatosensory cortex. *J. Neurophysiol.*

Fuentes, C.T., and Bastian, A.J. (2010). Where is your arm? Variations in proprioception across space and tasks. *J. Neurophysiol.* 103, 164–171.

Garraghty, P.E., Florence, S.L., and Kaas, J.H. (1990). Ablations of areas 3a and 3b of monkey somatosensory cortex abolish cutaneous responsivity in area 1. *Brain Res.*

- Head, H., and Holmes, G. (1911). Sensory disturbances from cerebral lesions. *Brain*.
- Hyvarinen, J., and Poranen, A. (1978). Receptive field integration and submodality convergence in the hand area of post-central gyrus of the alert monkey. *J. Physiol.* *283*, 539–556.
- Iwamura, Y., Tanaka, M., Sakamoto, M., and Hikosaka, O. (1983). Converging patterns of finger representation and complex response properties of neurons in area 1 of the first somatosensory cortex of the conscious monkey. *Exp. Brain Res.* *51*, 327–337.
- Jazayeri, M., and Movshon, J.A. (2006). Optimal representation of sensory information by neural populations. *Nat. Neurosci.* *9*, 690–696.
- Jiang, W., Chpman, C., and Lamarre, Y. (1991). Modulation of the cutaneous responsiveness of neurones in the primary somatosensory cortex during conditioned arm movements in the monkey. *Exp. Brain Res.* *84*, 342–354.
- Jones, S.R., Pritchett, D.L., Stufflebeam, S.M., Hämäläinen, M., and Moore, C.I. (2007). Neural correlates of tactile detection: A combined magnetoencephalography and biophysically based computational modeling study. *J. Neurosci.* *27*, 10751–10764.
- Kersten, D., Mamassian, P., and Yuille, A. (2004). Object Perception as Bayesian Inference. *Annu. Rev. Psychol.* *55*, 271–304.
- Knight, F.L.C., Longo, M.R., and Bremner, A.J. (2014). Categorical perception of tactile distance. *Cognition* *131*, 254–262.
- Kording, K., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. *Nature* *427*, 1–4.
- Lackner, J.R. (1988). Some proprioceptive influences on the perceptual representation of body shape and orientation. *Brain*.
- Lillywhite, P.G., and Laughlin, S.B. (1979). Transducer noise in a photoreceptor. *Nature* *277*, 569–572.
- Longo, M.R. (2017). Expansion of perceptual body maps near – But not across – The wrist. *Front. Hum. Neurosci.*
- Longo, M.R., and Haggard, P. (2010). An implicit body representation underlying human position sense. *Proc. Natl. Acad. Sci. U. S. A.* *107*, 11727–11732.

Longo, M.R., Azañón, E., and Haggard, P. (2010). More than skin deep: Body representation beyond primary somatosensory cortex. *Neuropsychologia* *48*, 655–668.

Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* *9*, 1432–1438.

Maij, F., Wing, A.M., and Pieter Medendorp, W. (2013). Spatiotemporal integration for tactile localization during arm movements: A probabilistic approach. *J. Neurophysiol.*

Mancini, F., Longo, M.R., Iannetti, G.D., and Haggard, P. (2011). A supramodal representation of the body surface. *Neuropsychologia*.

Marini, F., Squeri, V., Morasso, P., and Masia, L. (2016). Wrist proprioception: Amplitude or position coding? *Front. Neurobot.* *10*, 1–8.

McGuire, L.M.M., and Sabes, P.N. (2009). Sensory transformations and the use of multiple reference frames for reach planning. *Nat. Neurosci.* *12*, 1056–1061.

Medendorp, W.P., Goltz, H.C., Crawford, J.D., and Vilis, T. (2005). Integration of target and effector information in human posterior parietal cortex for the planning of action. *J. Neurophysiol.* *93*, 954–962.

Medina, J., and Coslett, H.B. (2010). From maps to form to space: Touch and the body schema. *Neuropsychologia* *48*, 645–654.

Miller, L.E., Longo, M.R., and Saygin, A.P. (2014). Tool morphology constrains the effects of tool use on body representations. *J. Exp. Psychol. Hum. Percept. Perform.* *40*, 2143–2153.

Miller, L.E., Longo, M.R., and Saygin, A.P. (2017). Visual illusion of tool use recalibrates tactile perception. *Cognition* *162*, 32–40.

Miller, L.E., Montroni, L., Koun, E., Salemme, R., Hayward, V., and Farnè, A. (2018). Sensing with tools extends somatosensory processing beyond the body. *Nature* *561*, 239–242.

Miller, L.E., Ravenda, V., Bolognini, N., Hayward, V., Farnè, A., Ravenda, V., Bahmad, S., Koun, E., and Salemme, R. (2019a). Somatosensory Cortex Efficiently Processes Touch Located Beyond the Body. *Curr. Biol.* 1–8.

Miller, L.E., Longo, M.R., and Saygin, A.P. (2019b). Tool Use Modulates Somatosensory Cortical Processing in Humans. *J. Cogn. Neurosci.*

Mittelstaedt, M.L., and Mittelstaedt, H. (1980). Homing by path integration in a mammal.

Naturwissenschaften.

Moscattelli, A., Naceri, A., and Ernst, M.O. (2014). Path integration in tactile perception of shapes. *Behav. Brain Res.*

Mountcastle, V.B., and Powell, T.P. (1959). Neural mechanisms subserving cutaneous sensibility, with special reference to the role of afferent inhibition in sensory perception and discrimination. *Bull. Johns Hopkins Hosp.* 105, 201–232.

Mountcastle, V.B., Lynch, J.C., Georgopoulos, A., Sakata, H., and Acuna, C. (1975). Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *J. Neurophysiol.* 38, 871–908.

Nicolelis, M.A.L., Ghazanfar, A.A., Stambaugh, C.R., Oliveira, L.M.O., Laubach, M., Chapin, J.K., Nelson, R.J., and Kaas, J.H. (1998). Simultaneous encoding of tactile information by three primate cortical areas. *Nat. Neurosci.* 1, 621–630.

Niehof, N., Perdreau, F., Koppen, M., and Medendorp, W.P. (2019). Contributions of optostatic and optokinetic cues to the perception of vertical. *J. Neurophysiol.* 122, 480–489.

Oruç, I., Maloney, L.T., and Landy, M.S. (2003). Weighted linear cue combination with possibly correlated error. *Vision Res.* 43, 2451–2468.

Penfield, W., and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain.*

Pesaran, B., Nelson, M.J., and Andersen, R.A. (2006). Dorsal Premotor Neurons Encode the Relative Position of the Hand, Eye, and Goal during Reach Planning. *Neuron* 51, 125–134.

Pillsbury, W.B. (1895). Some Questions of the Cutaneous Sensibility. *Am. J. Psychol.* 7, 42–57.

Pitkow, X., and Angelaki, D.E. (2017). Inference in the Brain: Statistics Flowing in Redundant Population Codes. *Neuron* 94, 943–953.

Proske, U., and Gandevia, S.C. (2012). The proprioceptive senses: Their roles in signaling body shape, body position and movement, and muscle force. *Physiol. Rev.* 92, 1651–1697.

Reed, C.L., Klatzky, R.L., and Halgren, E. (2005). What vs. where in touch: An fMRI study. *Neuroimage.*

Reed, J.L., Qi, H.X., Pouget, P., Burish, M.J., Bonds, A.B., and Kaas, J.H. (2010). Modular

processing in the hand representation of primate primary somatosensory cortex coexists with widespread activation. *J. Neurophysiol.* *104*, 3136–3145.

Riesenhuber, M., and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.* *2*, 1019–1025.

Sakata, H., Takaoka, Y., Kawarasaki, A., and Shibutani, H. (1973). Somatosensory properties of neurons in the superior parietal cortex (area 5) of the rhesus monkey. *Brain Res.*

Salinas, E. (2006). How behavioral constraints may determine optimal sensory representations. *PLoS Biol.* *4*, 2383–2392.

Sanger, T.D. (2003). Neural population codes. *Curr. Opin. Neurobiol.* *13*, 238–249.

Schaefer, M., Rothemund, Y., Heinze, H.J., and Rotte, M. (2004). Short-term plasticity of the primary somatosensory cortex during tool use. *Neuroreport.*

Schwartz, O., and Simoncelli, E.P. (2001). Natural signal statistics and sensory gain control. *Nat. Neurosci.* *4*, 819–825.

Seelke, A.M.H., Padberg, J.J., Disbrow, E., Purnell, S.M., Recanzone, G., and Krubitzer, L. (2012). Topographic maps within brodmann's area 5 of macaque monkeys. *Cereb. Cortex* *22*, 1834–1850.

Seung, H., and Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci.* *90*, 10749–10753.

Shadlen, M.N., and Newsome, W.T. (1998). The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *J. Neurosci.* *18*, 3870–3896.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.*

Shen, G., Smyk, N.J., Meltzoff, A.N., and Marshall, P.J. (2018). Neuropsychology of human body parts: Exploring categorical boundaries of tactile perception using somatosensory mismatch responses. *J. Cogn. Neurosci.*

Smith, E.C., and Lewicki, M.S. (2006). Efficient auditory coding. *Nature* *439*, 978–982.

Song, W., and Francis, J.T. (2013). Tactile information processing in primate hand somatosensory cortex (S1) during passive arm movement. *J. Neurophysiol.* *110*, 2061–2070.

Stangl, M., Kanitscheider, I., Riemer, M., Fiete, I., and Wolbers, T. (2020). Sources of path integration error in young and aging humans. *Nat. Commun.* 11.

Vallbo, A.B., Olausson, H., Wessberg, J., and Kakuda, N. (1995). Receptive field characteristics of tactile units with myelinated afferents in hairy skin of human subjects. *J. Physiol.*

De Vignemont, F., Majid, A., Jola, C., and Haggard, P. (2009). Segmenting the body into parts: Evidence from biases in tactile perception. *Q. J. Exp. Psychol.* 62, 500–512.

Vogt, B.A., and Pandya, D.N. (1978). Cortico-cortical connections of somatic sensory cortex (areas 3, 1 and 2) in the rhesus monkey. *J. Comp. Neurol.* 177, 179–191.

Weber, E. (1834). The sense of touch.

Yamamoto, S., and Kitazawa, S. (2001). Sensation at the tips of invisible tools. *Nat. Neurosci.* 979–980.

Yarrow, S., and Series, P. (2015). The influence of population size, noise strength and behavioral task on Best-Encoded stimulus for neurons with unimodal or monotonic tuning curves. *Front. Comput. Neurosci.*


Zohary, E., Shadlen, M.N., and Newsome, W.T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370, 140–143.

Participant	$\hat{\varepsilon}_1$	$\hat{\varepsilon}_2$	$\hat{\sigma}$	R^2	$RMSE$
1	2.12	2.14	0.28	0.91	1.31
2	1.91	3.23	0.28	0.79	2.28
3	9.46	0.01	0.34	0.69	4.13
4	3.57	2.05	0.29	0.55	4.06
5	0.76	1.17	0.27	0.91	1.36
6	3.50	5.89	0.35	0.72	3.31
7	2.62	0.36	0.45	0.87	2.85
8	7.53	2.09	0.51	0.77	4.34
9	2.08	0.97	0.29	0.93	1.29
10	1.66	0.01	0.65	0.94	2.50
11	13.50	4.98	0.27	0.84	2.03
Mean	4.43±1.21	2.08±0.59	0.36±0.04	0.81±0.04	2.68±0.35

Supplementary Table 1. Best fit parameters for Experiment 1

Participant	$\hat{\varepsilon}_1$	$\hat{\varepsilon}_2$	$\hat{\sigma}$	R^2	$RMSE$
1	4.23	2.68	0.21	0.65	2.25
2	2.27	2.81	0.21	0.87	1.19
3	4.07	4.64	0.15	0.57	1.71
4	5.04	2.69	0.28	0.96	0.89
5	2.75	0.01	0.31	0.68	3.62
6	3.99	6.06	0.28	0.92	1.20
7	4.32	6.21	0.24	0.85	1.42
8	6.14	0.01	0.27	0.80	2.15
9	1.57	4.84	0.29	0.98	0.63
10	1.92	0.41	0.26	0.73	2.47
11	9.86	0.01	0.32	0.87	2.3
12	4.72	0.01	0.34	0.63	5.07
13	6.94	6.39	0.21	0.96	0.55
14	2.09	1.90	0.22	0.78	1.86
Mean	4.28±0.61	2.77±0.66	0.26±0.01	0.80±0.04	1.95±0.33

Supplementary Table 2. Best fit parameters for Experiment 2



Supplementary Figure 1. Removal of correlated bias inherited from encoding layer

(A) When the value of g is set to a fixed value, the biases from the encoding layer are inherited by both location estimates. Therefore, they are highly correlated (10000 simulations). (B) When the value of g is instead set to the maximum synaptic drive between the encoding and decoding layers (see Methods), this removes the inherited bias and thus the correlation between the bias in both location estimates drops near zero.