



# Prediction, postdiction, and perceptual length contraction: a Bayesian low-speed prior captures the cutaneous rabbit and related illusions

Daniel Goldreich\* and Jonathan Tong

Department of Psychology, Neuroscience & Behaviour, McMaster University, Hamilton, ON, Canada

## Edited by:

Yuki Yamada, Yamaguchi University, Japan

## Reviewed by:

Iris M. D. Vilares, Northwestern University and Rehabilitation Institute of Chicago, USA

Robert Van Beers, VU University Amsterdam, Netherlands

## \*Correspondence:

Daniel Goldreich, Department of Psychology, Neuroscience & Behaviour, McMaster University, 1280 Main Street West, Hamilton, ON L8S 4K1, Canada.  
e-mail: goldrd@mcmaster.ca

Illusions provide a window into the brain's perceptual strategies. In certain illusions, an ostensibly task-irrelevant variable influences perception. For example, in touch as in audition and vision, the perceived distance between successive punctate stimuli reflects not only the actual distance but curiously the inter-stimulus time. Stimuli presented at different positions in rapid succession are drawn perceptually toward one another. This effect manifests in several illusions, among them the startling cutaneous rabbit, in which taps delivered to as few as two skin positions appear to hop progressively from one position to the next, landing in the process on intervening areas that were never stimulated. Here we provide an accessible step-by-step exposition of a Bayesian perceptual model that replicates the rabbit and related illusions. The Bayesian observer optimally joins uncertain estimates of spatial location with the expectation that stimuli tend to move slowly. We speculate that this expectation – a Bayesian prior – represents the statistics of naturally occurring stimuli, learned by humans through sensory experience. In its simplest form, the model contains a single free parameter,  $\tau$ : a time constant for space perception. We show that the Bayesian observer incorporates both pre- and post-dictive inference. Directed spatial attention affects the prediction-postdiction balance, shifting the model's percept toward the attended location, as observed experimentally in humans. Applying the model to the perception of multi-tap sequences, we show that the low-speed prior fits perception better than an alternative, low-acceleration prior. We discuss the applicability of our model to related tactile, visual, and auditory illusions. To facilitate future model-driven experimental studies, we present a convenient freeware computer program that implements the Bayesian observer; we invite investigators to use this program to create their own testable predictions.

**Keywords: probabilistic inference, sensory saltation, motion illusions, tactile spatial attention, optimal percepts, Kalman smoothing, somatosensory spatiotemporal perception, sensory uncertainty**

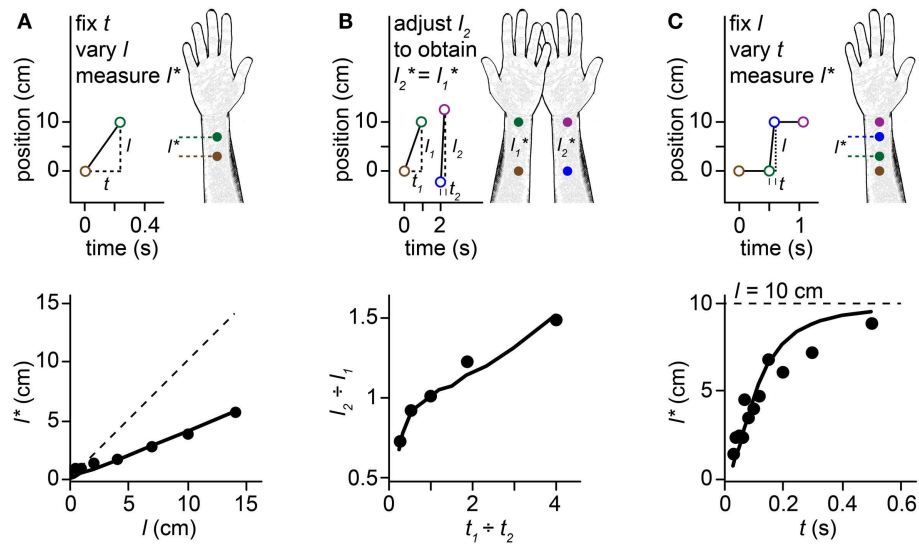
## INTRODUCTION

Illusions provide investigators a window into the brain's unconscious perceptual strategies. In a particularly interesting category of illusions, an ostensibly task-irrelevant stimulus feature strongly influences the perception of a target feature. Here we consider one group of such illusions, characterized by the curious influence of time on the tactile perception of space (Figure 1).

When humans are asked to judge the distance between two brief taps delivered in rapid succession to the skin, they consistently underestimate the true distance. Indeed, the perceived distance between taps shortens systematically as the time between taps is reduced. This *perceptual length contraction* occurs even when the participant is explicitly instructed to attend only to the distance between stimuli, and to ignore the time. The phenomenon is particularly pronounced on the forearm and other body areas that have poor spatial acuity. Several striking illusions result from this puzzling compressive effect of time on space perception (Figures 1A–C). For instance, a stimulus sequence consisting of

two-taps delivered at one position followed by two taps at another, with a short inter-stimulus interval (ISI) separating the second and third taps, is perceived as four taps hopping progressively along the arm: the second and third taps are perceptually displaced from their true positions, as if attracted toward one another (Figure 1C). This phenomenon is known as sensory saltation, or more famously, the cutaneous rabbit illusion (Geldard and Sherrick, 1972; Geldard, 1982). Analogous phenomena occur in vision (Geldard, 1976; Lockhead et al., 1980; Khoo et al., 2011) and audition (Bremer et al., 1977; Shore et al., 1998; Getzmann, 2009).

Why does time influence space perception in this manner? Much research supports the view that perception works out a probabilistic best guess. An optimal probabilistic (i.e., Bayesian) observer interprets the current sensory input, not in isolation, but rather within the context of the structure and statistics of the natural world (Knill and Pouget, 2004; Vilares and Kording, 2011). By exploiting its knowledge of the world, the observer achieves a more



**FIGURE 1 | Perceptual length contraction.** Perception underestimates the distance between successive taps to the skin. Stimuli on the forearm are illustrated in the upper panels, along with their perception (forearm sketches). Corresponding human data and Bayesian model fits are plotted in the lower panels. In this and subsequent figures, we illustrate stimulus sequences that progress distally on the arm; the illusions occur also for stimuli in the opposite direction. **(A) Top:** at short ISI ( $t$ ), the perceived length ( $l^*$ ) between two taps to the forearm is less than the actual length ( $l$ ). **Bottom:** perceived length grows linearly with actual length, but with a slope less than 1. Filled circles: human perceptual data from Marks et al. (1982) for electrocutaneous stimuli delivered at  $t = 0.24$  s. Solid line: fit of the Bayesian model. Dashed line:  $l = l^*$ . **(B) Top:** a pair of taps delivered to the right forearm at short ISI ( $t_2$ ) is perceived to have the

same spacing as a more closely spaced pair of taps ( $l_1 < l_2$ ) delivered to the left forearm at longer ISI ( $t_1 > t_2$ ). **Bottom:** the spacing ratio,  $l_2$ -to- $l_1$ , resulting in perceived equality of spacing on the two arms, as a function of the ISI ratio,  $t_1$ -to- $t_2$ . Filled circles: human perceptual data from Lechelt and Borchert (1977). Curve: fit of the Bayesian model. Data points from left to right had  $t_1 = 0.2, 0.35, 0.5, 0.65,$  and  $0.8$  s, with  $t_2 = 1.0$  s  $- t_1$ , and  $l_1 = 10$  cm. **(C) Top:** 4 taps delivered to two skin sites are perceived as hopping sequentially along the arm, because the short ISI ( $t$ ) between taps 2 and 3 results in contraction of the perceived distance between them ( $l^* < l$ ). **Bottom:** the perceived length from taps 2–3 asymptotically approaches the actual length ( $l = 10$  cm, dashed line) as ISI is increased. Filled circles: human perceptual data from Kilgard and Merzenich (1995). Curve: fit of the Bayesian model.

accurate perceptual inference. Following the Bayesian model of Goldreich (2007), we hypothesize that perception interprets successive taps to the skin as arising from a moving object that touches down intermittently, and that perception expects slowly moving objects to occur more often than rapidly moving ones. We speculate that the expectation for slow movement results from a lifetime of experience with tactile stimuli that are primarily stationary (e.g., the pressure of clothing against the skin) or – somewhat less frequently – slowly moving (e.g., grooming, movement of clothing during walking, etc.). Thus, in the observer's experience, stimuli separated by large distances at short ISI are uncommon. Faced with such a stimulus sequence, and somewhat uncertain as to the true locations of the taps, the brain concludes that the sensory measurements were caused by a stimulus sequence that was more probable *a priori*: one that moved at a slower speed (i.e., shorter distance) on the skin. Under this view, the influence of time over space perception, far from reflecting a design flaw in our perceptual machinery, is a consequence of optimal probabilistic inference under conditions of sensory uncertainty.

Here, we present and elaborate on the Bayesian observer model introduced by Goldreich (2007). We show that our model is compatible with the view that the rabbit illusion – and perceptual length contraction generally – involves concomitant pre- and post-diction. By prediction, we mean an inference process in which

earlier sensory events influence the perception of later ones. By postdiction, we mean an inference process in which later sensory events influence the perception of earlier ones (Eagleman and Sejnowski, 2000). We show interestingly that pre- and post-diction emerge naturally from our model, even though the model does not explicitly represent these processes. We show further that directed spatial attention shifts the Bayesian observer's percept by modulating the prediction-postdiction balance. Finally, we apply our Bayesian model to the perception of spatiotemporal stimulus patterns that are more complex than those depicted in Figure 1.

## THE FUNDAMENTALS OF THE BAYESIAN OBSERVER

Stochastic variability in stimulus-evoked neural activity presents one of many challenges to perception. An identical repeated stimulus – such as a tap to a particular location on the skin – will evoke a different neural response on each trial (Sripati et al., 2006). Consequently, a given response could have been caused by a stimulus at any one of many locations. The spatial uncertainty caused by stochastic variability is lessened, but not eliminated, when a stimulus activates a larger number of neurons. On the forearm, where receptor density is relatively low, humans can localize a stimulus to within about  $\pm 1$  cm of its true location; on the fingertip, where receptor density is much higher, localization improves to about  $\pm 1$  mm (Weinstein, 1968).

To model stochastic neural variability, we assume that a single tap to the skin evokes an internal position *measurement* that is randomly sampled from a Gaussian distribution centered at the true tap position, with a standard deviation,  $\sigma_s$ , that depends on the receptor density (the subscript *s* signifies “spatial”)<sup>1</sup>. On repeated trials with an identical tap position, the measurement will vary stochastically, but on average will equal the true position. In the absence of any other perceptual influence, the measurement is the location the observer perceives. Consequently, on average the perception of an isolated single tap to the skin is veridical. However, unlike an isolated single tap, a rapid spatiotemporal tap sequence is not veridically perceived (**Figure 1**). To understand why, we explore a probabilistic model – a Bayesian observer that makes a perceptual best guess.

We begin by considering sequences of two taps, which result in two uncertain spatial measurements ( $x_{1m}, x_{2m}$ ) and a detected time,  $t$ , between them<sup>2</sup>. The Bayesian observer (**Figure 2**) attempts

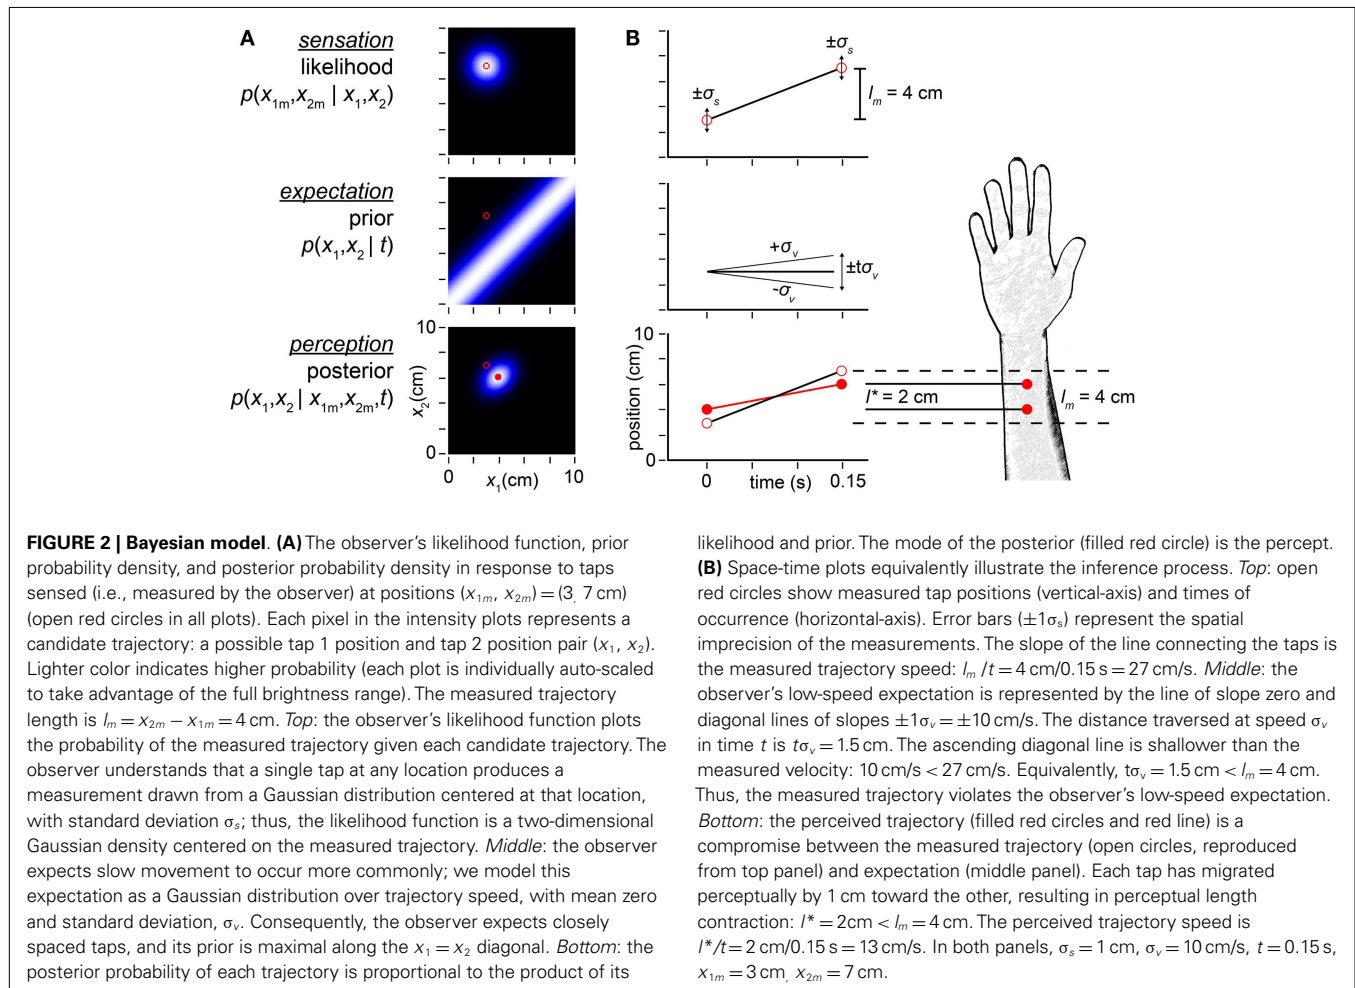
to infer the actual tap positions ( $x_1, x_2$ ) that produced the measurements ( $x_{1m}, x_{2m}$ ). We refer to each possible ( $x_1, x_2$ ) pair as a *candidate trajectory*, and to the measured positions ( $x_{1m}, x_{2m}$ ) as the *measured trajectory*. The Bayesian observer considers both the *likelihood* and the *prior probability* of every candidate trajectory. A trajectory’s likelihood is the probability that the trajectory would give rise to the measured trajectory. The plot of trajectory likelihoods – the likelihood function – is a cloud of uncertainty centered on the measured trajectory (**Figure 2A**, top). We analogize the likelihood function to a (typically unconscious) *sensation* – a precursor to the conscious percept.

A trajectory’s prior probability is the frequency with which the observer expects the trajectory to occur; this may be the prevalence of the trajectory in nature, which the observer has learned from experience. The plot of prior probabilities – the prior density – represents the observer’s *expectation* regarding trajectory occurrence. Crucially, our Bayesian observer believes that slow trajectories are more common than fast ones. We model this low-speed prior

<sup>1</sup>Neuroscientists may find it useful to conceive of the measurement as the location of the peak of evoked activity in the underlying receptor population (or its cortical equivalent), or more precisely as the maximal likelihood estimate of stimulus location, based on the neural response.

<sup>2</sup>We assume here that the observer veridically perceives the time between taps, such that temporal uncertainty is zero. Goldreich (2007) showed that temporal

uncertainty exerts a negligible effect on the percept when stimuli occur on a skin region with poor spatial acuity, such as the forearm. Accordingly, here we confine ourselves to modeling stimuli on the forearm, which is also the skin region most often tested in experimental studies of the cutaneous tau and rabbit illusions.



as a Gaussian density over trajectory speed, with mean zero and standard deviation  $\sigma_v$  (the subscript  $v$  signifies “velocity”). Thus, trajectories in which the two taps are spaced closer together (i.e., lower-speed trajectories) have greater prior probability than those in which the taps are spaced farther apart (Figure 2A, middle).

Using Bayes’ rule, the observer multiplies each trajectory’s likelihood by its prior probability to obtain its posterior (final) probability. In essence, the observer combines *sensation* with *expectation* to achieve *perception*. The mode of the posterior distribution – the most probable trajectory – is the observer’s percept (Figure 2A, bottom). Because of the low-speed prior, the percept underestimates the distance between rapidly presented stimuli. In the example illustrated, whereas the measured tap positions were (3, 7 cm), the percept was (4, 6 cm). The perceived distance between taps ( $l^* = 2$  cm) was thus half the measured distance ( $l_m = 4$  cm) (Figures 2A,B).

How, exactly, does the time between taps influence perceptual length contraction? This question is answered in Figure 3. Because speed is distance divided by time, the prior probability falls off more sharply with distance when the time between taps is short.

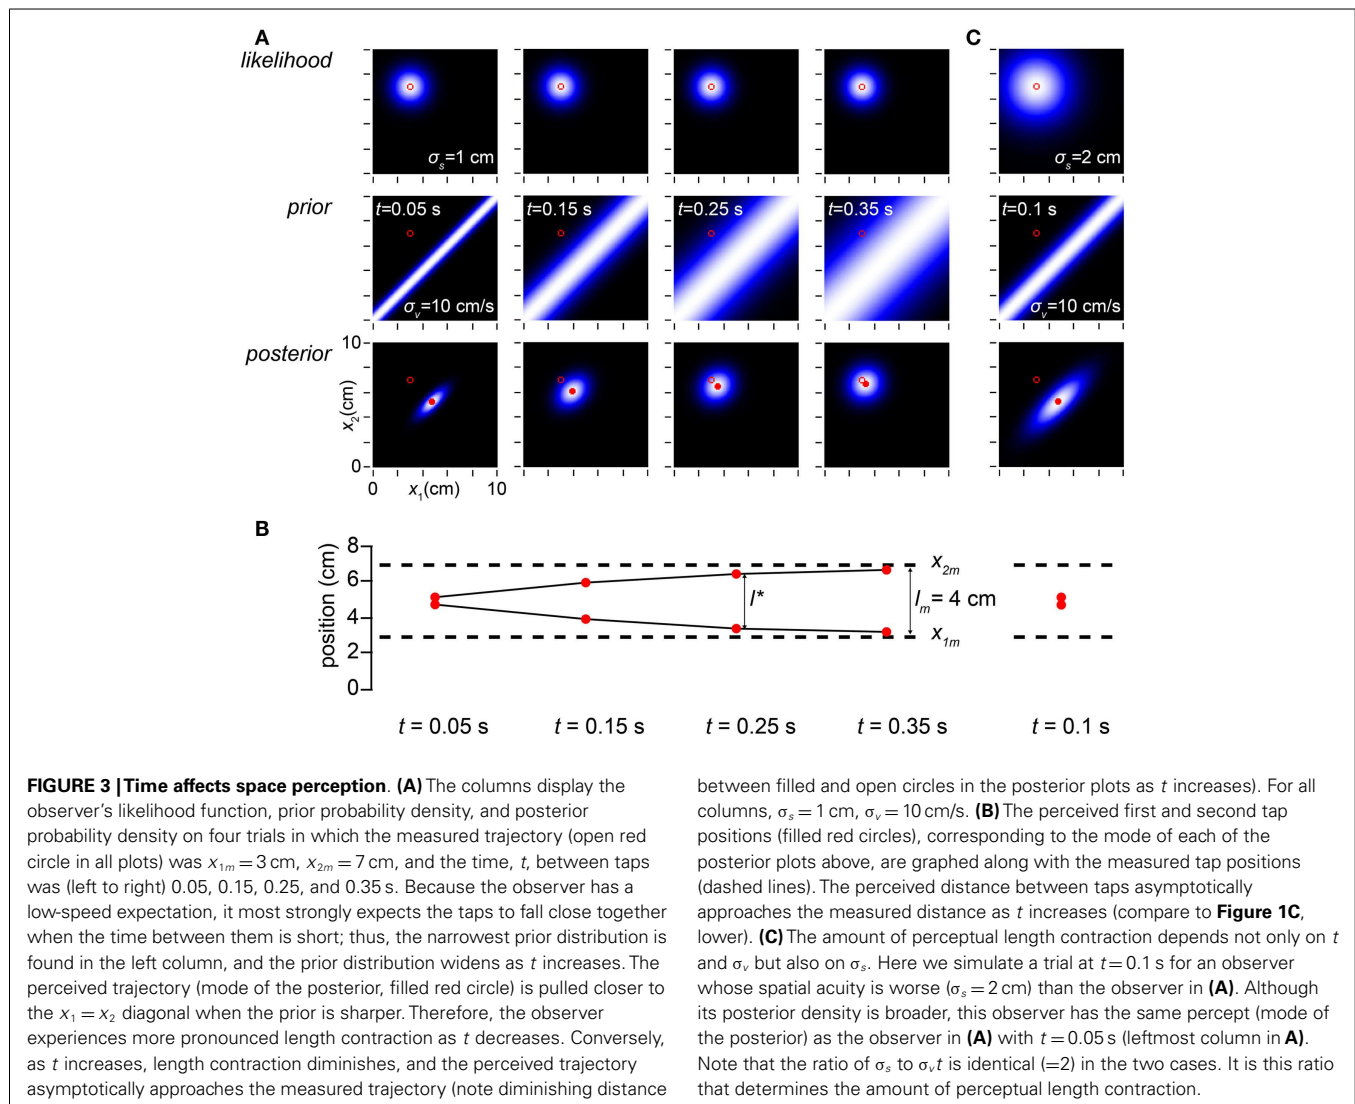
While always maximal along the  $x_1 = x_2$  diagonal, the prior widens as ISI increases (Figure 3A, left to right). As a consequence, perceptual length contraction is most pronounced at shorter ISIs; as ISI increases, the perceived distance between taps asymptotically approaches the measured distance (Figure 3B).

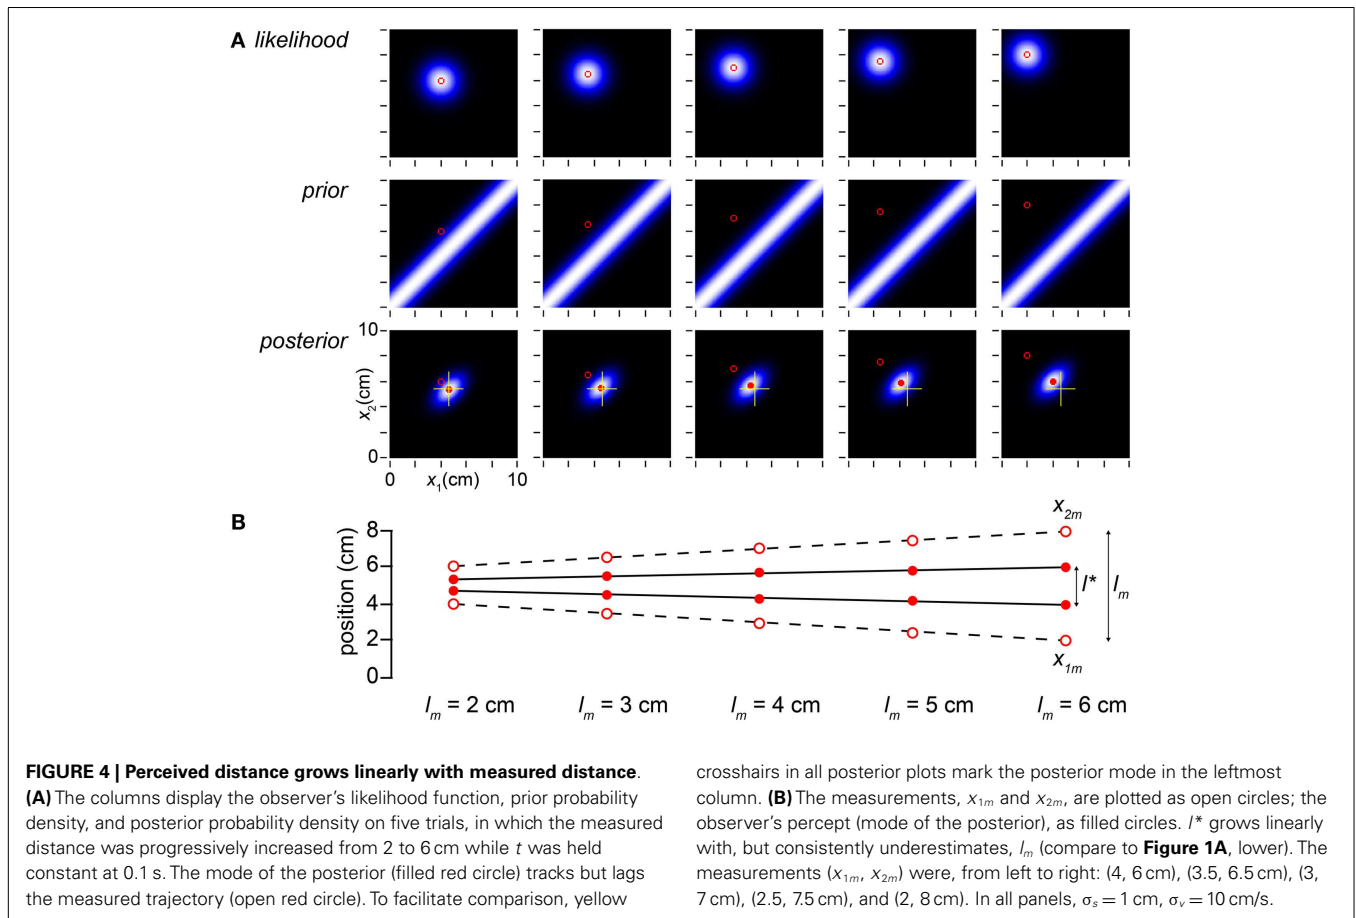
We have explained the influence of time on the Bayesian observer’s perception of space, but what of the influence of space itself on space perception? In Figure 4, we find reassuringly that  $l^*$  varies linearly with  $l_m$ , although length contraction ensures that the slope of the relationship is less than one.

### THE PERCEPTUAL LENGTH CONTRACTION FORMULA

In the Section “The Bayesian model” in Appendix, we show that the Bayesian observer’s posterior density is a two-dimensional Gaussian distribution. The mode of the posterior reveals a relationship between  $l^*$  and  $l_m$ :

$$l^* = \frac{l_m}{1 + 2 \left( \frac{\sigma_s}{\sigma_v t} \right)^2} \quad (1)$$





Equation 1 is the perceptual length contraction formula, first reported by Goldreich (2007). Notice that, as we have seen, this formula predicts that  $l^*$  asymptotically approaches  $l_m$  in the limit that  $t$  approaches infinity (**Figures 3A,B**), that the degree of length contraction is determined by the ratio of  $\sigma_s$  to  $\sigma_v t$  (**Figure 3C**), and that, at fixed  $t$ ,  $l^*$  relates linearly to, but underestimates,  $l_m$  (**Figure 4**).

Because  $\sigma_s$  and  $\sigma_v$  occur only as a ratio in the length contraction formula, it is convenient to rewrite the formula as:

$$l^* = \frac{l_m}{1 + 2\left(\frac{\tau}{t}\right)^2} \tag{2}$$

where tau ( $\tau$ ), defined as  $\sigma_s/\sigma_v$ , has units of time, and is the model's single free parameter<sup>3</sup>. From Eq. 2 we see that tau is a time constant for space perception. The smaller the value of tau, the more the perceived length increases toward the measured length as inter-stimulus time increases:  $l^* = (1/3) l_m$  when  $t = \tau$ , and  $l^* = (2/3) l_m$  when  $t = 2\tau$  (**Figure 5A**). Thus, the larger the value of  $\tau$ , the more susceptible the observer is to perceptual length contraction: for a given  $t$  and  $l_m$ , an observer with a larger  $\tau$  will perceive a shorter trajectory (**Figures 5A,B**).

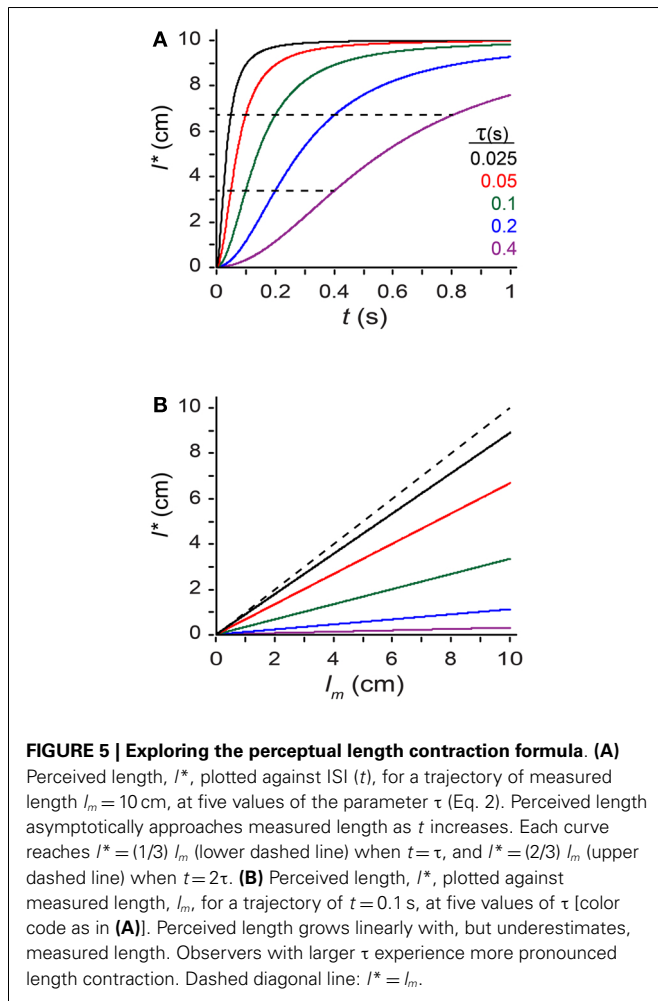
<sup>3</sup>We note for reference that Goldreich (2007) defined the model's free parameter as  $\lambda = \sigma_v/\sigma_s$ ; thus, the lambda parameter in that paper is simply the reciprocal of the tau parameter.

To develop an intuition for these effects of tau, consider that the parameter can be rewritten:

$$\tau = \frac{\sigma_s}{\sigma_v} = \frac{1/\sigma_v}{1/\sigma_s} = \frac{\text{strength of low-speed expectation}}{\text{spatial acuity}} \tag{3}$$

Thus, tau reflects the strength of the observer's low-speed expectation relative to the observer's spatial acuity. Tau is large in an observer with poor spatial acuity (large  $\sigma_s$ ) and a strong expectation for slow movement (small  $\sigma_v$ ). This observer places trust in the low-speed expectation; the observer's perception is considerably length contracted. Tau is small in an observer with excellent spatial acuity (small  $\sigma_s$ ) and little expectation regarding movement speed (large  $\sigma_v$ ). This observer places trust in the measurement; the observer's perception is only modestly length contracted.

The perceptual length contraction formula closely fits human data from a variety of experiments (**Figure 1**; see also Goldreich, 2007 for additional data fits). The fit is particularly satisfying given that the formula has just a single free parameter. The best-fit  $\tau$ -values for the data displayed in **Figures 1A–C** were 0.21, 0.11, and 0.08 s. The larger  $\tau$  for the **Figure 1A** fit may reflect the use of electrocutaneous stimuli by Marks et al. (1982), the source of the data plotted in **Figure 1A**. Electrical pulses tend to be more difficult to localize (larger  $\sigma_s$ ) than mechanical taps (Higashiyama



and Hayashi, 1993), which were used to generate the data in **Figure 1B** (Lechelt and Borchert, 1977) and **Figure 1C** (Kilgard and Merzenich, 1995). Measures of point localization suggest that  $\sigma_s$  is on the order of 1 cm in response to light mechanical stimuli on the forearm (Weinstein, 1968; Martikainen and Pertovaara, 2002; Cody et al., 2008); thus, taking  $\tau = 0.1$  s as a nominal value for mechanosensory perception on the forearm, we infer that  $\sigma_v$  is on the order of 10 cm/s.

### BAYESIAN PERCEPTION IS OPTIMAL BECAUSE IT IS BENEFICIALLY BIASED

Before developing our model further, we pause to consider an important conceptual question: we have described the Bayesian observer as achieving an optimal perceptual inference, but we have also shown that the observer consistently underestimates the measured distance between taps. How can an observer be both biased and optimal? This important question applies to any Bayesian observer with a non-uniform prior distribution.

The short answer to the question is that bias is optimal when it accurately reflects the stimulus statistics. In a world in which slow trajectories are more common than fast ones (and, therefore, among trajectories with any given inter-stimulus time,  $t$ , short lengths are more common than long ones), an observer is justified

in perceiving trajectories as shorter than measured. Paradoxically, then, the Bayesian observer is optimal precisely because it is biased.

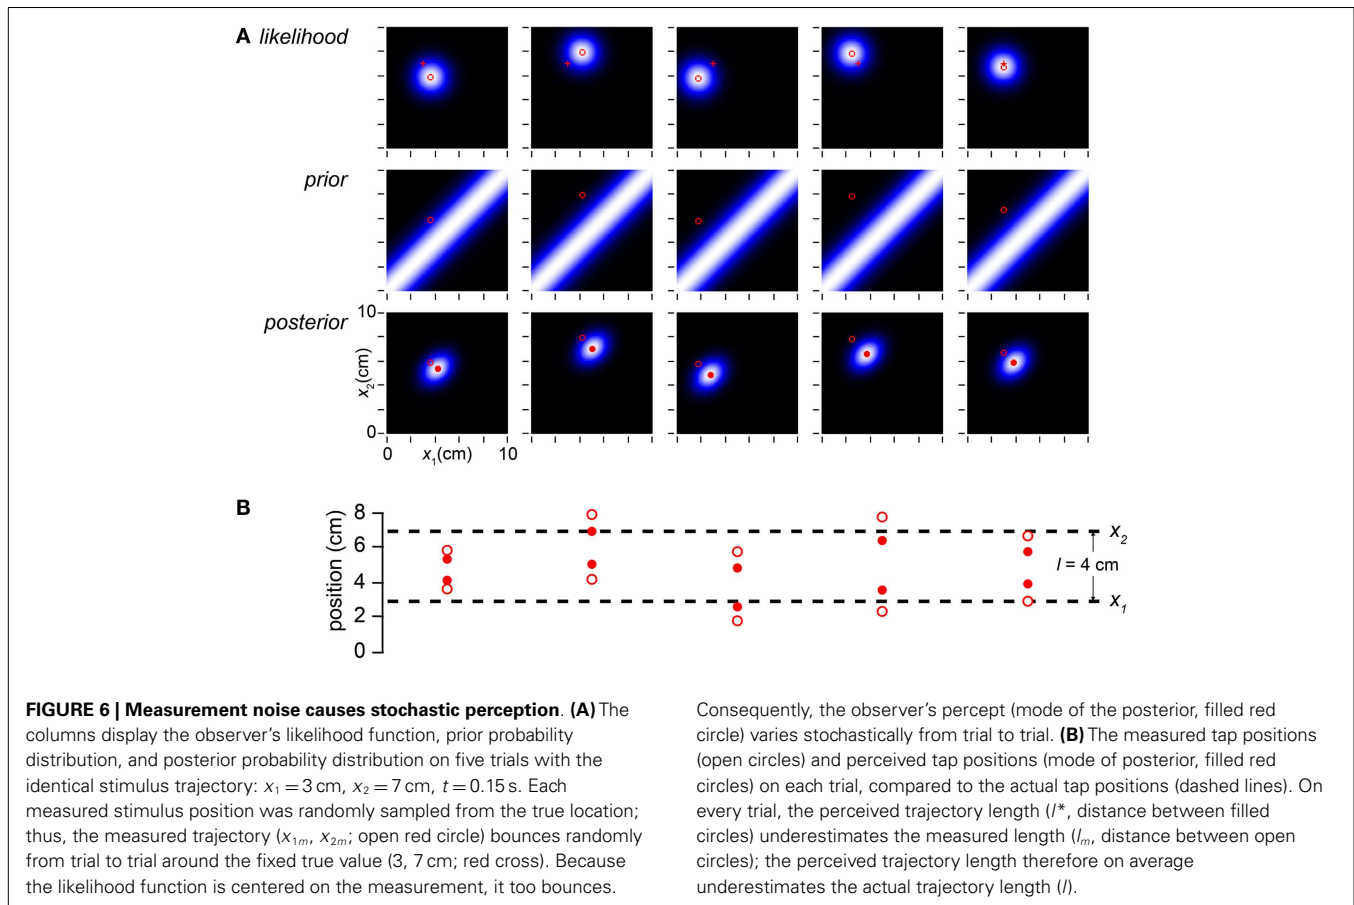
To understand this thoroughly, we must appreciate the consequences of both measurement and stimulus variability. In **Figures 2–5** we artificially specified  $(x_{1m}, x_{2m})$ . In a laboratory experiment, however, the investigator can control only the stimulus, not the measurements. As explained, we conceive of each measured tap location as sampled from a Gaussian distribution of standard deviation  $\sigma_s$ , centered on the actual tap location. Thus, if the skin is stimulated repeatedly with the identical trajectory, the measurement and consequently the percept will vary stochastically from trial to trial (**Figure 6**).

By incorporating measurement variability, the simulation shown in **Figure 6** is a more realistic representation of a laboratory experiment than are the simulations shown in the earlier Figures. Crucially for our understanding of the paradox of bias and optimality, however, **Figure 6** would be an unrealistic portrayal of the Bayesian observer's experience in the real-world. In the real-world, not only the measurements but also the trajectories themselves are drawn from a distribution. In **Figure 7**, we more closely simulate what we envision to be real-world tactile experience. The figure plots the lengths of one million trajectories sampled from a zero-mean velocity distribution (for clarity of illustration, all with  $t = 0.15$  s), from each of which spatial measurements were sampled and processed into a percept.

A comparison of the statistics of the measured length,  $l_m$  (**Figure 7A**) with those of the perceived length,  $l^*$  (**Figure 7B**) reveals that, although the observer's perception is biased, it is more accurate than the measurement. In fact, the observer's perception is optimal precisely because it is biased. To understand why, consider that the majority of these real-world trajectories have very short lengths ( $l$  close to zero). Because short trajectories are more common, any measured length,  $l_m$ , most often originates from a trajectory of shorter true length,  $l$ . The Bayesian observer's percept is biased by the prior to take this crucial knowledge into account; consequently, over the course of many trials, the percept more closely reflects the true stimulus than the measurement does. This is indicated by the smaller vertical scatter of the percept (**Figure 7B**, left) than of the measurement (**Figure 7A**, left) around the diagonal line.

Further inspection of the scatterplot in **Figure 7A** reveals that, for any true trajectory length,  $l$ , the measurement,  $l_m$ , occurs with equal frequency above and below the diagonal line. Thus, the histogram of  $l_m$  samples is centered on  $l$  (**Figure 7A**, center). For this reason, the measured length is termed an "unbiased estimator" of the true length. Despite this lofty denomination, however, it is clear from the same scatterplot that for any magnitude  $l_m$  other than 0, the distribution of true lengths has a smaller average magnitude (when  $l_m > 0$ ,  $l$  tends to lie to the left of the diagonal line; when  $l_m < 0$ ,  $l$  tends to lie to the right of the diagonal line). Thus,  $l_m$  is an inaccurate estimator in the sense that the stimuli that result in a particular  $l_m$  are on average offset from that  $l_m$  (**Figure 7A**, right). If an observer were to report  $l_m$  as the estimate of trajectory length, the observer would be found to systematically report trajectories as being longer than they actually are.

**Figure 7B** shows that the statistics of the perceived length,  $l^*$ , are opposite in character to those of the measured length. For any



true trajectory length,  $l$ , the perceived length,  $l^*$ , systematically underestimates the magnitude of  $l$  (Figure 7B, left and center). Thus, the perceived length is termed a “biased estimator.” This bias is beneficial, however: because of it, at any  $l^*$ , the distribution of true lengths is centered on a mean of  $l^*$  (the values of  $l$  are symmetrically distributed around the diagonal line in the scatterplot). Thus,  $l^*$  is an accurate estimator in the sense that the stimuli that result in a particular  $l^*$  indeed on average have length equal to that  $l^*$  (Figure 7B, right). The observer's report of  $l^*$  can be trusted as accurately reflecting, on average, the true trajectory length. Importantly, the variance of  $l$  given  $l^*$  (Figure 7B, right) is smaller than the variance of  $l_m$  given  $l$  (Figure 7A, center). This again reveals that the percept is more accurate than the measurement.

### SELECTIVE SPATIAL ATTENTION SHIFTS THE PERCEIVED TRAJECTORY

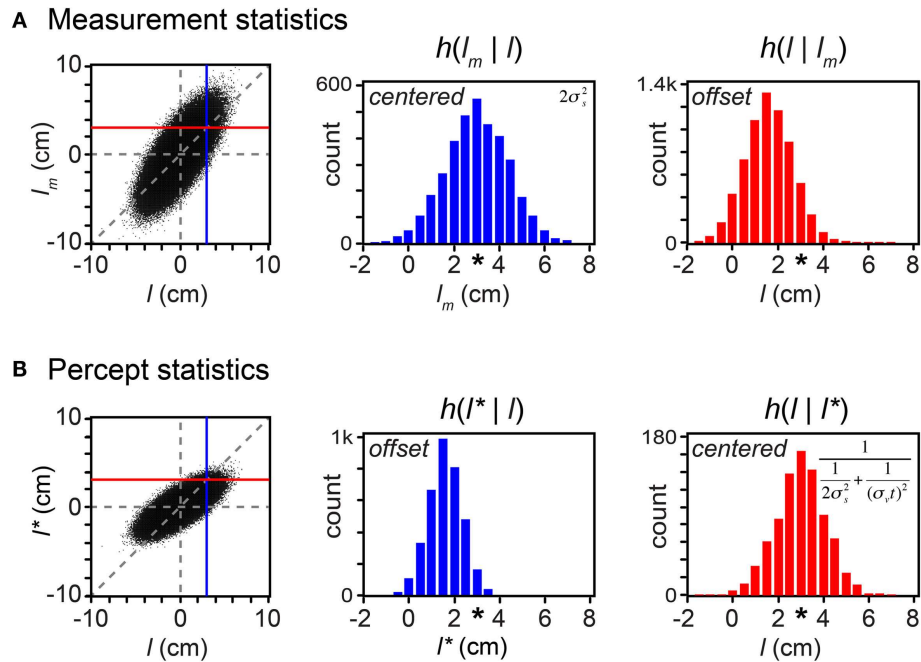
Up to this point, we have assumed that the observer's spatial uncertainty,  $\sigma_s$ , is uniform within the tested area ( $\sigma_s$  will, of course, differ between body areas, such as forearm and finger). However, spatial attention is associated with cortical receptive field recruitment and sharpening within the attended area (Anton-Erxleben and Carrasco, 2013). Thus, if an observer were to focus attention preferentially on one location, we might expect  $\sigma_s$  to decrease there while plausibly increasing at unattended locations. Indeed, on the arm, the spatial error of localization decreases by as much as 30%

when attention is directed to the stimulated skin region (Moore et al., 1999; O'Boyle et al., 2001).

If spatial acuity is modulated by selective attention, how might length contraction percepts be affected? In a cutaneous rabbit experiment, Kilgard and Merzenich (1995) found that when participants were not asked to focus their attention to any particular area of the arm, the midpoints of the perceived and actual trajectories tended to coincide (Figure 8A, left). In contrast, when participants were instructed to direct their attention either distally or proximally, the midpoint of the perceived trajectory shifted toward the attended location (Figure 8A, center, right). This occurred because the tap within the attended skin area migrated less perceptually than did the tap within the unattended area, an effect confirmed by Flach and Haggard (2006).

The Bayesian observer replicates this attention effect: when  $\sigma_s$  decreases in one skin area relative to the other, the perceived trajectory midpoint shifts toward the attended location (Figures 8B,C). The relatively precise measurement of the “attended tap” impedes its perceptual migration, while the relatively imprecise measurement of the “unattended tap” facilitates its perceptual migration. In this situation, length contraction is accomplished primarily by the perceptual displacement of the unattended tap.

In the Section “Generalization to inhomogeneous spatial uncertainty” in Appendix, we derive a generalization of the length contraction formula that incorporates separate  $\sigma_{s1}$  and  $\sigma_{s2}$  values representing spatial uncertainty around the two tap locations.



**FIGURE 7 | Bayesian perception is optimal because it is biased.** On each of 1 million trials, a first tap position ( $x_1$ ) was drawn from a uniform distribution, and a second tap position ( $x_2$ ) was drawn from a Gaussian distribution centered on the first tap position, with standard deviation  $t\sigma_v = 1.5$  cm (i.e.,  $\sigma_v = 10$  cm/s,  $t = 0.15$  s; see Eq. A8 in Appendix). Measured positions,  $x_{1m}$  and  $x_{2m}$ , were then drawn independently from Gaussian distributions of standard deviation  $\sigma_s = 1$  cm, centered on the corresponding tap positions ( $x_1$  and  $x_2$ ). **(A)** *Left*: scatterplot of measured trajectory length ( $l_m = x_{2m} - x_{1m}$ ) against actual trajectory length ( $l = x_2 - x_1$ ) for each of the trials (dots); negative lengths indicate trajectories in which  $x_2 < x_1$ . Dashed vertical and horizontal lines:  $l = 0$  and  $l_m = 0$ . Diagonal dashed line:  $l_m = l$ . Vertical blue line:  $l = 3$  cm. Horizontal red line:  $l_m = 3$  cm. *Center*: histogram ( $h$ ) of  $l_m$  values that occurred when  $l$  was between 2.95 and 3.05 cm (i.e.,  $l_m$

samples that fell along the blue vertical line in the scatterplot). The histogram is a Gaussian distribution centered at  $l_m = 3$  cm (asterisk). *Right*: histogram of  $l$  values of trajectories that gave rise to  $l_m$  between 2.95 and 3.05 cm (i.e.,  $l$  samples that fell along the red horizontal line in the scatterplot). The histogram represents the observer’s posterior density over  $l$ . It is a Gaussian distribution centered at  $l = 1.6$  cm, not 3 cm (asterisk). **(B)** *Left*, center, and right panels as in **(A)**, but for  $l^*$  rather than  $l_m$ . *Center*:  $l^*$  is a biased estimator. *Right*: on trials in which the observer perceived  $l^* = 3$  cm, the true trajectory length averaged 3 cm. Because the perceived length is a deterministic function of the measurement, this histogram has the same variance as the posterior density over  $l$ . Inset formulas in **(A)** center and **(B)** right show the variances of these histograms (See “One-dimensional reductions” in Appendix). These are equal to the mean-squared error between each estimator and the true length.

In the general equation, the single spatial uncertainty,  $\sigma_s$ , of Eq. 1 is replaced by the root-mean-square uncertainty at the two locations,  $\sigma_{rms}$ :

$$l^* = \frac{l_m}{1 + 2\left(\frac{\sigma_{s(rms)}}{\sigma_v t}\right)^2} = \frac{l_m}{1 + \frac{\sigma_{s1}^2 + \sigma_{s2}^2}{(\sigma_v t)^2}} \quad (4)$$

We show further that the shift,  $\Delta_{midpt}$ , in the perceived trajectory midpoint away from the measured trajectory midpoint is:

$$\Delta_{midpt} = \frac{l_m}{2} \left( \frac{\sigma_{s1}^2 - \sigma_{s2}^2}{(\sigma_v t)^2 + \sigma_{s1}^2 + \sigma_{s2}^2} \right) \quad (5)$$

**THE PREDICTIVE-POSTDICTIVE FORMULATION**

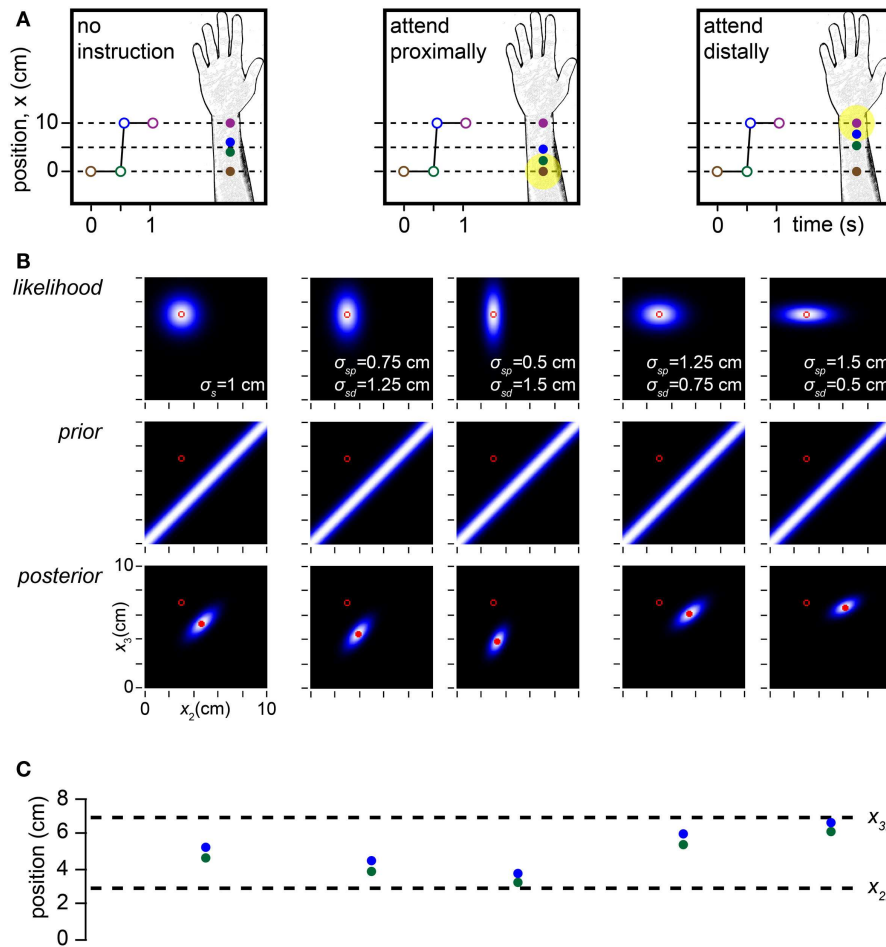
The rabbit illusion is often described as providing compelling evidence for perceptual postdiction, a process whereby the perception of an earlier event is modified by the occurrence of a later one. Postdiction is indeed an attractive explanation for the perceptual migration of tap 2 toward the location of tap 3 in the rabbit illusion (Figure 1C). As shown by Kilgard and Merzenich

(1995), tap 3 also migrates perceptually toward the location of tap 2 (Figure 1C). Therefore, prediction apparently is also at play: the perception of a later event (tap 3) depends upon an earlier one (tap 2).

In light of these considerations, it may seem surprising that our Bayesian observer replicates length contraction illusions without explicitly representing either pre- or postdictive inference. How is this possible? The answer is that pre- and postdiction are implicitly embedded in the model via the action of the low-speed prior. The low-speed prior transforms the observer’s likelihood function into a posterior density by pulling the observer’s perception of each tap position toward the measured position of the other (Figure 2).

We can reveal the pre- and postdiction hidden in the Bayesian observer by decomposing the model’s two-dimensional ( $x_1, x_2$ ) calculations (Figure 9A) into a series of one-dimensional inferences regarding each tap’s position individually (Figure 9B). Using its low-speed expectation, the observer can from the first tap’s likelihood function predict a probability distribution over the position of the subsequent, second, tap, and from the second tap’s likelihood function postdict a probability distribution over the





**FIGURE 8 | Modeling the effects of spatial attention. (A)** Depiction of a cutaneous rabbit illusion experiment reported by Kilgard and Merzenich (1995). Participants either received no specific instruction or were instructed to direct their attention (yellow highlight) toward the proximal or distal forearm. The investigators found that in the directed attention conditions, the perceived positions of tap 2 (green) and tap 3 (blue) were shifted toward the attended location (forearm sketches). **(B)** In the Bayesian observer, a reduction in  $\sigma_s$  at the attended relative to the unattended location reproduces the perceptual shift reported by Kilgard and Merzenich (1995). *Left panel:* the Bayesian observer’s likelihood function, prior and posterior density when  $\sigma_s$  does not vary with location, simulating the no-instruction condition in **(A)**. In

this case, the perceived and measured trajectory midpoints coincide. *Center two panels:* effect of  $\sigma_{sp} < \sigma_{sd}$ , where the subscripts  $p$  and  $d$  refer to the proximal and distal arm areas. The greater the reduction of  $\sigma_{sp}$  relative to  $\sigma_{sd}$ , the more the perceived trajectory migrates proximally toward the tap 2 measurement. *Right two panels:* effect of  $\sigma_{sd} < \sigma_{sp}$ . The greater the reduction of  $\sigma_{sd}$  relative to  $\sigma_{sp}$ , the more the perceived trajectory migrates distally toward the tap 3 measurement. For all plots in **(B)**, the measurements ( $x_{2m}$ ,  $x_{3m}$ ) were (3, 7 cm), the time between taps 2 and 3 was 0.06 s, and  $\sigma_v$  was 10 cm/s. **(C)** The perceived (mode of posterior) tap 2 and 3 positions (green and blue circles) for each of the five conditions in **(B)** directly above, compared to the measured tap positions (dashed lines).

position of the previous, first, tap (arrows in **Figure 9B**). We call these two distributions the *predicted prior* and *postdicted prior* densities<sup>4</sup>.

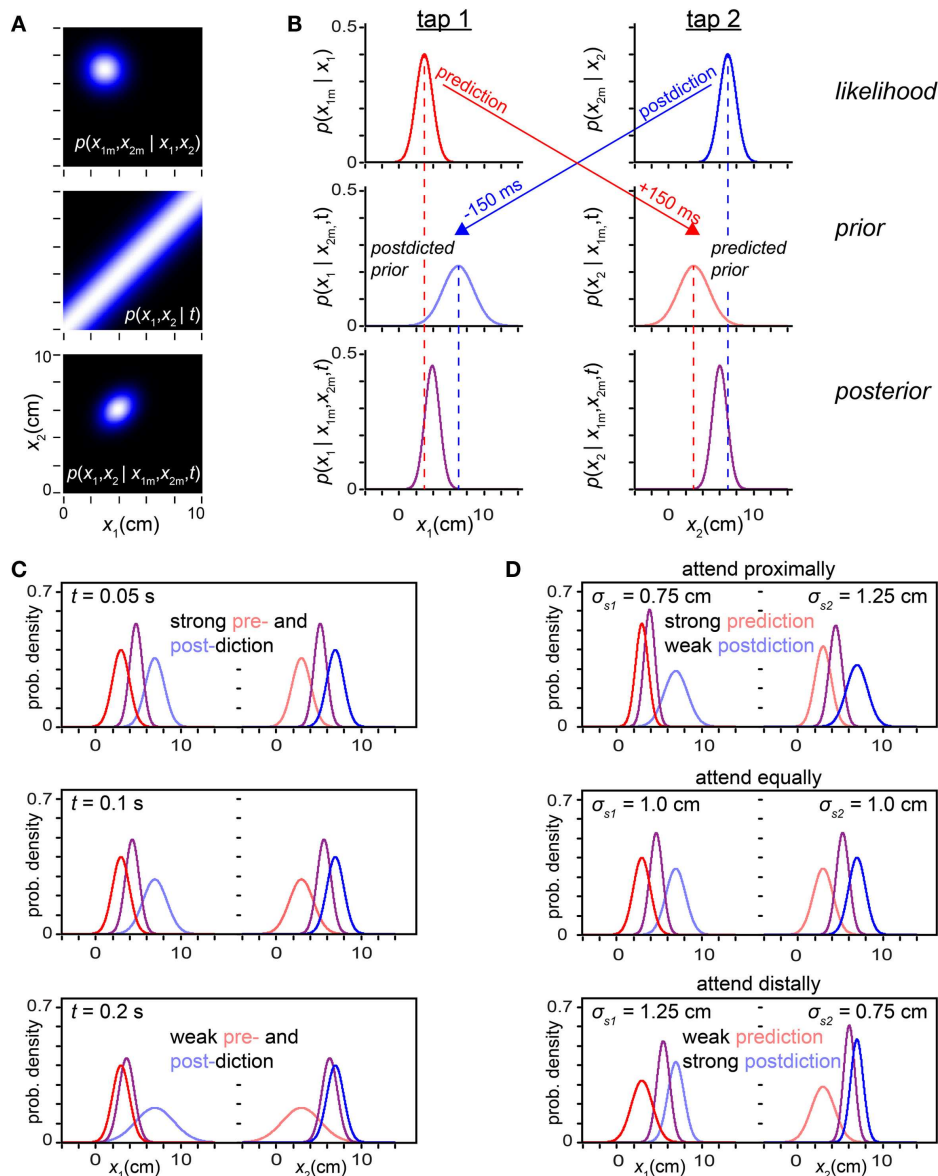
Next, the observer simply multiplies each tap’s likelihood function by that tap’s prior to obtain the posterior density over the tap’s position. We show in the Sections “One-dimensional reductions” and “The prediction-postdiction formulation” in Appendix that

the posteriors so obtained are identical to those that would result from extracting one-dimensional distributions from the joint ( $x_1$ ,  $x_2$ ) posterior: if the joint posterior (**Figure 9A**, bottom) were marginalized (i.e., integrated) vertically, it would yield the posterior over  $x_1$  shown in **Figure 9B**, bottom left; if integrated horizontally, it would yield the posterior over  $x_2$  shown in **Figure 9B**, bottom right.

In the Section “The prediction-postdiction formulation” in Appendix, we show that the predicted and postdicted priors are Gaussian densities, and that their means and variances are:

$$\begin{aligned} \mu_{pre} &= x_{1m} & \mu_{post} &= x_{2m} \\ \sigma_{pre}^2 &= \sigma_{s1}^2 + (\sigma_v t)^2 & \sigma_{post}^2 &= \sigma_{s2}^2 + (\sigma_v t)^2 \end{aligned} \tag{6}$$

<sup>4</sup>Note that “prior” in the Bayesian context does not imply “before” the stimulus occurs, but rather “independent of the measurement.” The predicted prior over tap 2’s position is constructed using all knowledge available to the observer except the tap 2 measurement,  $x_{2m}$ . Similarly, the postdicted prior over tap 1’s position is constructed using all knowledge available to the observer except the tap 1 measurement,  $x_{1m}$ .



**FIGURE 9 | Prediction-postdiction formulation. (A)** The observer's two-dimensional joint  $(x_1, x_2)$  likelihood function, prior and posterior densities. The measured trajectory was  $x_{1m} = 3$  cm,  $x_{2m} = 7$  cm, with  $t = 0.15$  s. The observer settings were  $\sigma_s = 1$  cm,  $\sigma_v = 10$  cm/s. **(B)** The inference process in **(A)** reformulated as a series of one-dimensional inferences regarding  $x_1$  and  $x_2$  individually. *Top left:* the tap 1 likelihood function (red),  $p(x_{1m} | x_1)$ , is centered on  $x_{1m}$ . Because of its low-speed expectation, the observer predicts (red arrow) that the most probable position for a future tap 2 will also be 3 cm. *Middle right:* the observer's *predicted prior* over tap 2 (light red) represents its belief concerning the position of tap 2, projected 150 ms forward in time from the occurrence of tap 1. *Top right:* the observer's tap 2 likelihood function (blue),  $p(x_{2m} | x_2)$ , is centered on  $x_{2m}$ . Because of its low-speed expectation, the observer postdicts (blue arrow) that the most probable position for the preceding tap 1 was also 7 cm. *Middle left:* the observer's *postdicted prior* over tap 1 (light blue) represents its belief concerning the position of tap 1, projected 150 ms backward in time from the occurrence of tap 2. *Left column:* using Bayes' theorem, the observer multiplies the tap 1 likelihood function (red) by the tap 1 postdicted prior (light blue) to obtain the

tap 1 posterior (purple). *Right column:* similarly, the observer multiplies the tap 2 likelihood function (blue) by the tap 2 predicted prior (light red) to obtain the tap 2 posterior (purple). **(C)** Individual tap likelihoods, priors, and posteriors graphed with the same color scheme as in **(B)**, for three trajectories of progressively increasing ISI. At  $t = 0.05$  s, pre- and postdiction both result in relatively sharp priors that exert a strong influence over the percept (mode of the posterior). As  $t$  is increased, the pre- and postdicted priors become lower and broader: pre- and postdiction become increasingly uncertain with the passage of time. The priors thus exert diminishing influence, and the percept approaches the measurement (compare to **Figure 3A**). For all panels in **(C)**,  $\sigma_s = 1$  cm,  $\sigma_v = 10$  cm/s. **(D)** Effect of directed spatial attention, as in **Figure 8**. *Top:* a reduction in  $\sigma_{s1}$  sharpens the tap 1 likelihood function, increasing the strength of prediction (note sharp predicted prior over tap 2), while an increase in  $\sigma_{s2}$  broadens the tap 2 likelihood function, decreasing the strength of postdiction (note broad postdicted prior over tap 1). *Middle:* when  $\sigma_{s1} = \sigma_{s2}$ , pre- and postdiction have equal strength. *Bottom:* reduction in  $\sigma_{s2}$  relative to  $\sigma_{s1}$  results in effects opposite those seen in the top panel. For all panels in **(D)**,  $t = 0.06$  s,  $\sigma_v = 10$  cm/s.

Equations 6 show that the prior density over each tap's position is centered on the measurement of the other tap, reflecting the observer's low-speed expectation (the most probable speed being zero). The variance of each prior density reflects the observer's uncertainty regarding the other tap's measurement ( $\sigma_{s1}$  or  $\sigma_{s2}$ ) and the observer's prior uncertainty regarding trajectory speed ( $\sigma_v$ ), which translates into an increasing uncertainty regarding the distance traversed as the elapsed time,  $t$ , increases ( $\sigma_v t$ ). Thus, perceptual length contraction diminishes with increasing  $t$  (Figure 9C), as shown previously (Figures 3 and 5A).

Figure 9D shows that the predictive-postdictive formulation accurately reproduces the effects of directed spatial attention, previously explored in Figure 8. When attention is directed around the location of the first tap ( $\sigma_{s1} < \sigma_{s2}$ ), the predicted prior is sharper than the postdicted prior ( $\sigma_{\text{pre}}^2 < \sigma_{\text{post}}^2$ ). Consequently, prediction exerts a dominant influence, perceptually displacing the second tap asymmetrically toward the first (Figure 9D, top). When attention is directed around the location of the second tap ( $\sigma_{s2} < \sigma_{s1}$ ), the postdicted prior is sharper ( $\sigma_{\text{post}}^2 < \sigma_{\text{pre}}^2$ ). In this case, postdiction dominates, perceptually displacing the first tap asymmetrically toward the second (Figure 9D, bottom).

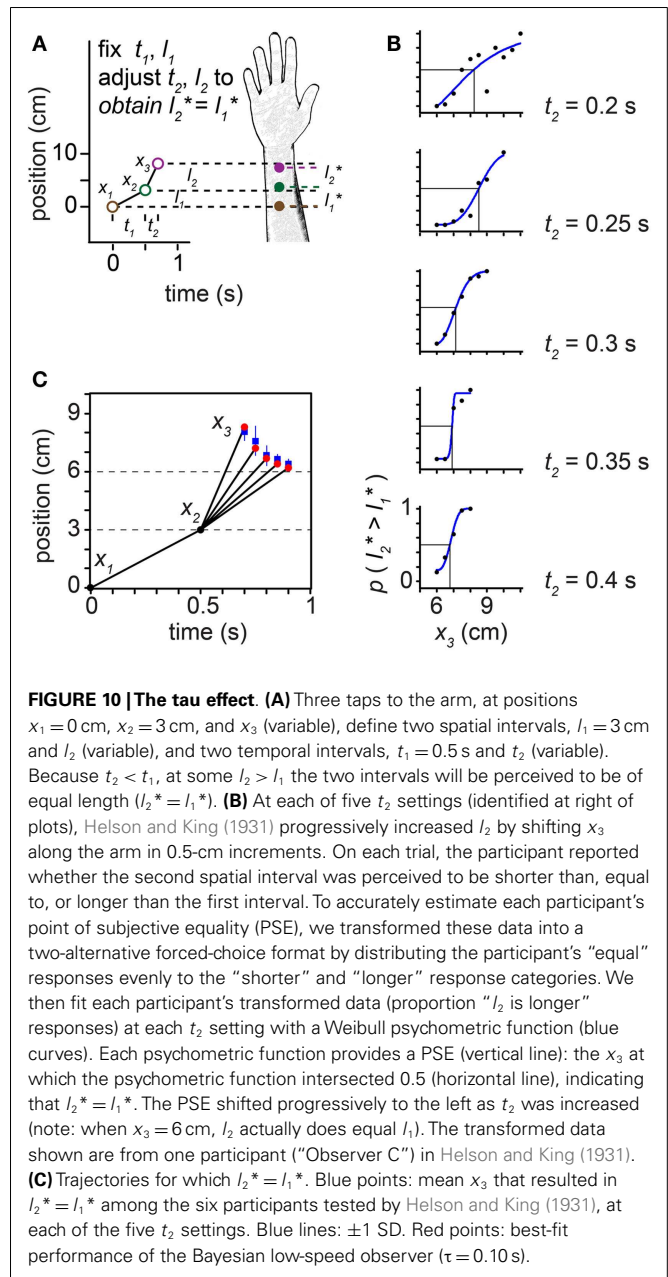
### THE PERCEPTION OF MULTI-TAP SEQUENCES

Up to this point, we have modeled the perception of two-tap trajectories<sup>5</sup>. How might a Bayesian observer handle multi-tap sequences, delivered conceivably to any number of skin sites? An observer could apply a low-speed prior independently to the movement between each tap and the next one. Alternatively, an observer might apply a low-speed prior to the first tap pair of the sequence, but thereafter incorporate an expectation that the velocity of each pair be similar to that of the preceding pair: a low-acceleration prior (See "Multi-tap perception" in Appendix).

Here, we test each of these Bayesian observers with multi-tap sequences that produce illusions in humans. We consider two well-known illusions. The first is the tau effect, so-named by Helson (1930) and subsequently described in elegant detail by Helson and King (1931). The second is a multi-tap rabbit, characterized in a delightful paper by Geldard (1982). In Figures 10 and 11, we show that the observer with a low-speed prior produces good fits to the human perceptual data; in Figure 12, we show that the observer with a low-acceleration prior does not.

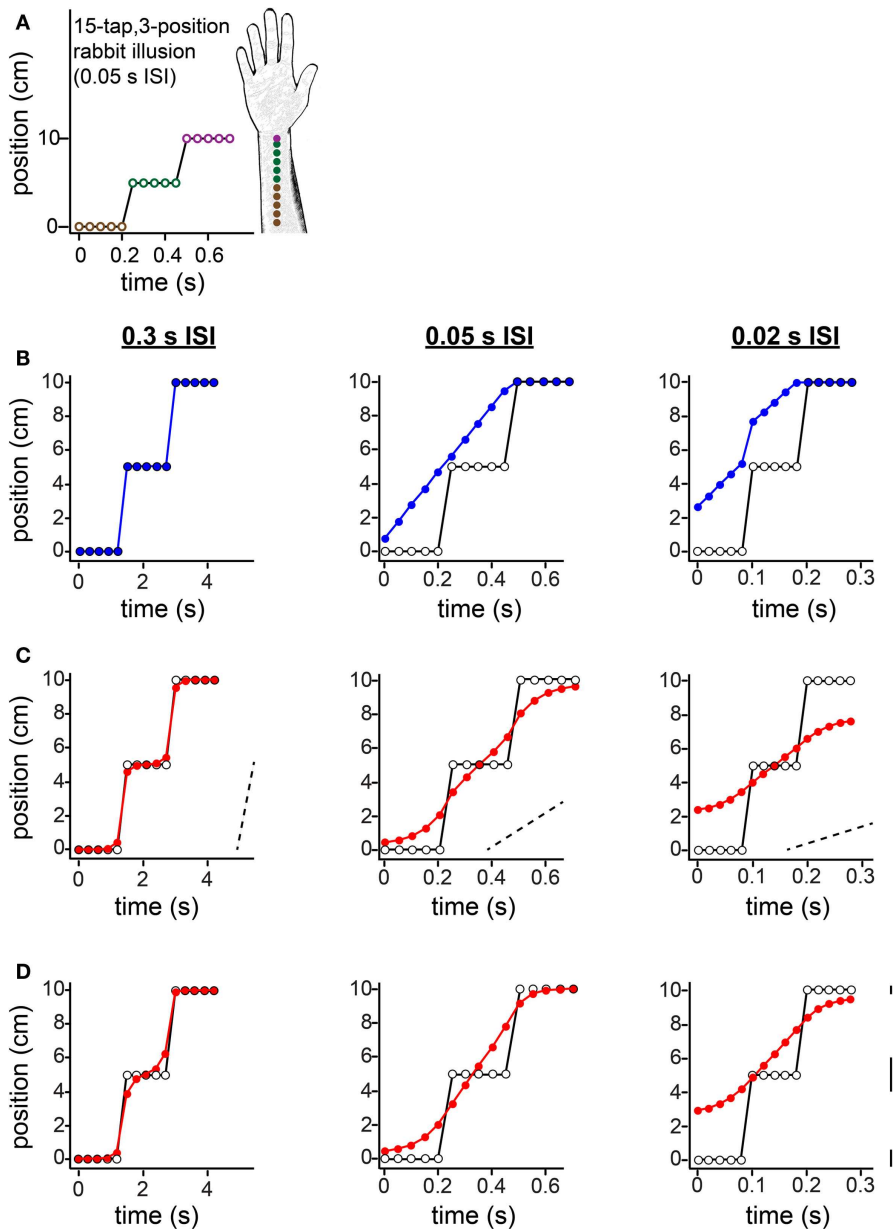
In the tau effect experiment, taps at three skin positions define two spatial and two temporal intervals (Figure 10). Helson and King (1931) reported that, when  $t_2 = t_1$  and  $l_2 = l_1$ , the participants perceived the two lengths as equal:  $l_2^* = l_1^*$ . As  $t_2$  was progressively reduced, however, tap 3 had to be located progressively farther down the arm (i.e.,  $l_2$  had to be progressively increased) in order to make  $l_2^*$  equal  $l_1^*$  (Figures 10B,C). The best-fit of our low-speed-prior observer to the average of the human data occurred at  $\tau = 0.10$  s. The Bayesian observer closely replicated the space-time curve characterizing human perception (Figure 10C).

<sup>5</sup>Although we have encountered a four-tap rabbit experiment (Figures 1C and 8), our approach was to consider the first and fourth taps as mere reference points, so we modeled the perception of taps 2 and 3 only. Indeed, the first and fourth taps in that sequence do not interact perceptually with the second and third, from which they are separated by large ISIs.



In the 15-tap rabbit experiment, five taps are delivered consecutively at each of three positions along the arm (Figure 11). Geldard (1982) found that when the time between consecutive taps was 0.05 s, participants perceived the first 10 taps in the sequence as hopping at an approximately uniform rate up the arm, each tap displaced by a constant spatial increment from the preceding one (Figures 11A,B, center). At an ISI of 0.3 s, perception was reportedly veridical (Figure 11B, left). At an ISI of 0.02 s, the perceived sequence began partway up the arm and traced a non-linear, somewhat sigmoidal path (Figure 11B, right).

The low-speed-prior observer's perception with  $\tau = 0.10$  s agrees qualitatively with the perception of human participants (Figure 11C). To understand why, first note that, at an ISI of 0.05 s

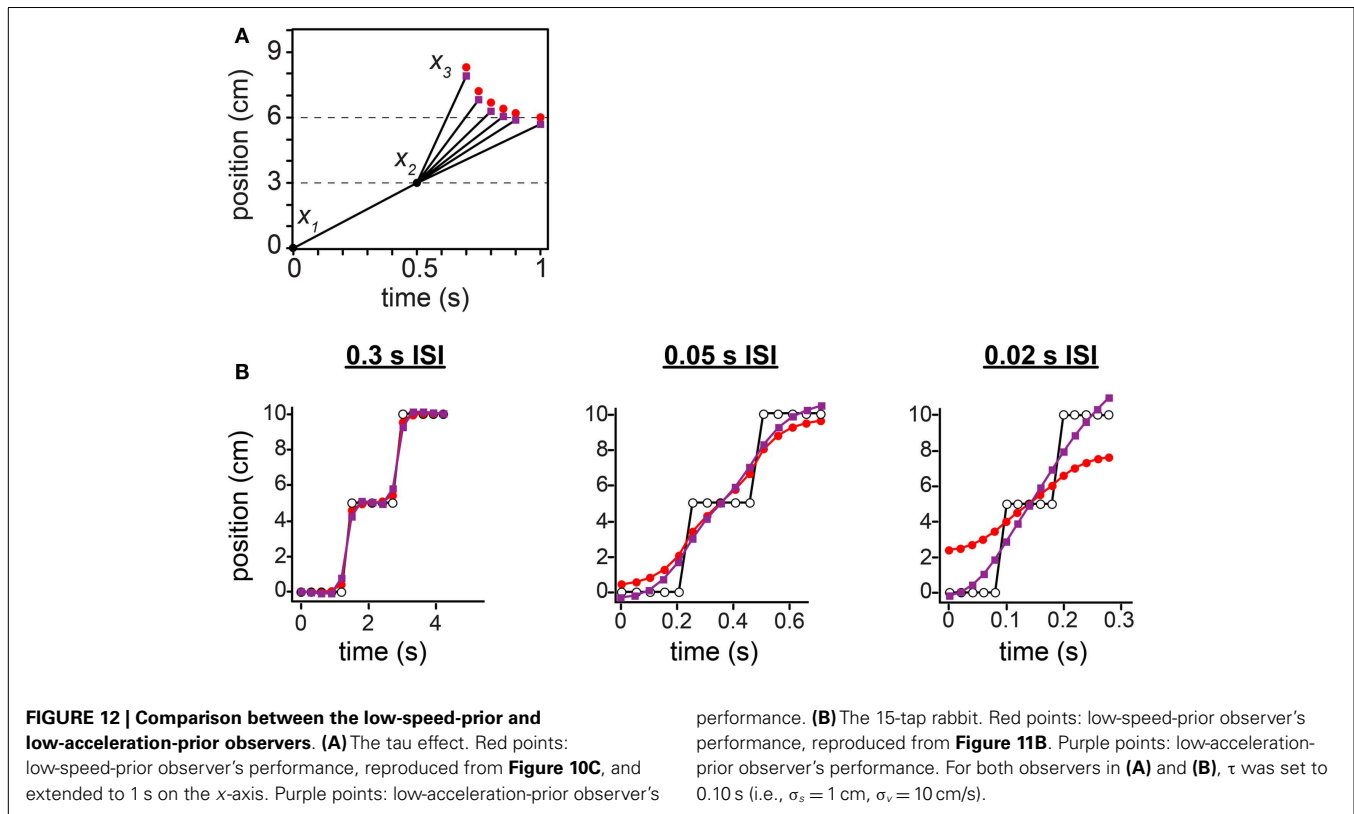


**FIGURE 11 | The 15-tap rabbit illusion.** (A) Geldard (1982) delivered five taps at each of three locations along the arm. When ISI between successive taps was 0.05 s, participants reported perceiving a linear spatial progression of taps 1 through 10 (forearm sketch). (B) The same spatial sequence shown in (A), at three different ISIs, resulted in distinct percepts (Geldard, 1982). *Left*: at 0.3 s ISI, perception was veridical. *Center*: at 0.05 s ISI, perception was as shown in (A). *Right*: at 0.02 s ISI, the taps were perceived to begin at a position between 2 and 3 cm along the arm, and to advance in a non-linear spatial progression. Open circles: true tap positions; blue points: human perceptual report. (C) The Bayesian low-speed observer's perception with a standard setting of  $\tau = 0.10$  s (e.g.,  $\sigma_s = 1$  cm,  $\sigma_v = 10$  cm/s) shows much similarity to participants' subjective reports. Open circles: true tap positions;

red points: Bayesian observer's perception (mode of the posterior). Dashed slanted lines have slope 10 cm/s (i.e.,  $1\sigma_v$ ). Note that the two rapid jumps in the true trajectory (from tap 5 to tap 6, and from tap 10 to tap 11) occur at a speed much greater than  $\sigma_v$  when the ISI is 0.05 s (*center*) or 0.02 s (*right*); thus, perceptual length contraction occurs in these cases. In contrast, at an ISI of 0.3 s (*left*), the trajectory does not strongly violate the observer's low-speed expectation; thus, perception is nearly veridical. (D) The Bayesian low-speed observer's perception can be made even closer to human reports if the value of  $\sigma_s$  varies along the arm. The observer's percept at each ISI is shown for  $\sigma_s = 1, 2$ , and 0.5 cm around the proximal, middle, and distal arm regions, respectively. Line segments at right have length equal to  $1\sigma_s$  at each location. The value of  $\sigma_v$  was fixed at 10 cm/s.

(Figure 11C, center) or 0.02 s (Figure 11C, right), the rapid jumps in the stimulus sequence are in clear violation of the observer's low-speed expectation (see diagonal dotted lines with slope  $\sigma_v$ ).

Consequently, perceptual length contraction occurs for those tap pairs: the perceived distance between taps 5 and 6, and between taps 10 and 11, is considerably smaller than the actual distance.



Now, what causes the progressive perceptual displacement of the many taps that are, in reality, at the same position? Interestingly, each jump in the actual stimulus sequence results in a chain reaction that propagates, with diminishing strength, to more distant taps. The rapid jump from tap 5 to tap 6 induces perceptual length contraction that pulls tap 5 considerably upward in the plot (and tap 6 downward). This places perceived distance between taps 4 and 5, which given the short ISI is sufficient to violate the observer's low-speed expectation as applied to that tap pair. Consequently, taps 4 and 5 are perceptually attracted, resulting in some upward perceptual displacement of tap 4, placing perceptual distance between it and tap 3, and so on.

How would perception of the 15-tap sequence change if the observer were to direct its spatial attention unequally along the arm? To explore this question, in **Figure 11D** we have plotted the low-speed-prior observer's perception under conditions of "standard" attention to the proximal arm ( $\sigma_s = 1$  cm), directed attention to the distal arm ( $\sigma_s = 0.5$  cm), and relative inattention ( $\sigma_s = 2$  cm) to the area in-between. Comparison of **Figures 11D,C** indicates that adjustment to spatial attention affects perception in ways that depend upon ISI. For the particular values of  $\sigma_s$  used in this example, perception of the 0.3 s ISI sequence remains nearly veridical (**Figure 11D**, left), whereas perception of the 0.05 s ISI sequence to some extent (center), and of the 0.02 s ISI sequence to a greater extent (right), are shifted upwards in the plots. The result is that the observer's perception even more closely resembles that of the human participants reported by Geldard (1982).

Unlike the low-speed-prior observer, the low-acceleration-prior observer distinctly fails to match human perception

(**Figure 12**). In the tau effect scenario, a discordant feature of the low-acceleration-prior observer is that, when  $t_2 = t_1$  and  $l_2 = l_1$ , the observer fails to perceive the lengths as equal, instead perceiving  $l_2^* > l_1^*$ . This perceptual asymmetry occurs because only the first segment of the trajectory is subject to a low-speed prior. Thus, when  $t_2 = t_1$ ,  $l_2$  must be made shorter than  $l_1$  in order to be perceived as equal. Consequently, in our simulation of Helson and King (1931) using the low-acceleration-prior-observer,  $x_3$  fails to converge to 6 cm as the tap 3 time approaches 1 s (**Figure 12A**, purple points). The performance of the low-speed-prior observer, in contrast, does converge as expected (red points).

In the 15-tap rabbit experiment, at 0.05 s ISI and more markedly at 0.02 s ISI, the low-acceleration-prior observer perceives the trajectory to start below the actual tap 1 location and to end above the actual tap 15 location: the perceived trajectory is longer than the actual trajectory (**Figure 12B**, purple points). This is incompatible with human perceptual report, and opposite to the perception of the low-speed-prior observer (red points). The perceptual undershoot and overshoot occur because the rapid jumps in the actual stimulus sequence extend perceptually in both directions at nearly constant velocity, in keeping with the observer's low-acceleration expectation.

## DISCUSSION

### PERCEPTUAL LENGTH CONTRACTION AS BAYESIAN INFERENCE

Length contraction illusions have long fascinated and puzzled investigators. The tactile tau effect was first reported almost 100 years ago (Gelb, 1914). It was later named and investigated in detail in the early 1930s (Helson, 1930; Helson and King, 1931).

The best-known length contraction illusion, the cutaneous rabbit, was discovered serendipitously some 40 years later, when Geldard and colleagues, intending to study the tau effect, mistakenly produced a stimulus pattern similar to the rapid sequences shown in **Figure 11B** (Geldard and Sherrick, 1972; Geldard, 1982). The resulting perception of taps hopping up the arm led a surprised observer to exclaim “who let the rabbit loose?” (Geldard, 1982). Over the years, investigators have proposed creative explanations – geometrical, mathematical, and neural – for these and related illusions (Jones and Huang, 1982; Brigner, 1988; Wiemer et al., 2000; Grush, 2005; Flach and Haggard, 2006).

The Bayesian observer model expounded here provides a concise and coherent explanation for the tau effect, the cutaneous rabbit, and related spatiotemporal illusions. Elapsed time influences the perception of traversed space because the observer expects objects to move slowly. In its simplest form, the model contains a single free parameter, tau: a time constant for space perception (Eqs 2 and 3). While much research remains to be done, we are encouraged by the close fit of the model to human perceptual data. Because a single model replicates the tau effect (**Figure 10**), the rabbit (**Figures 1C and 11**), and other spatiotemporal illusions (**Figures 1A,B**; see also Goldreich, 2007), we suggest that these illusions are manifestations of a single perceptual assumption: a low-speed prior. Our confidence in this suggestion is strengthened by the finding that a single value of the tau parameter ( $\sim 0.1$  s) provides good fits to perception on the forearm as measured in experiments using different paradigms and carried out by multiple laboratories.

A central feature of Bayesian perceptual models is that they consider multiple hypotheses – in our case, candidate trajectories. The idea that the brain perceives by evaluating candidates is consistent with the “multiple drafts” theory of Dennett and Kinsbourne (1992). These authors propose that, confronted with stimuli such as those depicted in **Figure 11**, the brain favors a distributed sequence of taps as the most “parsimonious” interpretation. This suggestion is compatible with our model if one equates parsimony with posterior probability. However, Dennett and Kinsbourne (1992) do not explain on what grounds an observer judges a particular interpretation to be the most parsimonious, nor do they explain why the percept changes as a function of ISI.

Bayesian perceptual models make precise, quantitative predictions regarding the relationships among perceptual variables (e.g., Eq. 1). These relationships spring from Bayes’ theorem: the product of a hypothesis’ likelihood and prior probability is proportional to its posterior probability. We liken the prior distribution to the observer’s expectation derived from experience, and the likelihood function to the sensation evoked by the stimulus (**Figure 2**). In our view, then, the Bayesian perceptual framework beautifully formalizes Helmholtz’s suggestion that “previous experiences act in conjunction with present sensations to produce a perceptual image” (Helmholtz, 1925).

Bayesian observers interpret sensory data in light of an internal model – a conception of the structure and statistics of the world. Bayesian perception is optimal when the observer’s internal model accurately represents the world – that is, when the observer’s prior distribution matches the stimulus distribution, and the observer’s likelihood function accurately reflects the process by which stimuli

map to measurements (**Figure 7**). Unfortunately, the natural statistics of tactile stimuli have not been sufficiently characterized to constrain a prior distribution, nor is our knowledge of tactile sensorineural responses sufficient to specify the precise shape of a likelihood function. Accordingly, we fit a Gaussian prior and Gaussian likelihood to the human behavioral data. Subtle discrepancies between the human data and the model’s performance could result from our Gaussian assumptions. Future research is needed to determine the precise shapes of the priors and likelihoods used by individual participants. In any event, we speculate that a low-speed prior reflects the natural statistics of tactile stimuli, learned by humans through experience. If so, illusions such as the cutaneous rabbit may reveal the operation of an optimal observer who brings an expectation forged by real-world experience (the low-speed prior) into an artificial setting (the laboratory).

### THE WIDE APPLICABILITY OF THE LOW-SPEED-PRIOR OBSERVER

Our Bayesian observer model may explain a variety of perceptual phenomena beyond the tactile illusions we have considered. One such phenomenon is the out-of-body rabbit illusion. In a clever experiment, Miyazaki et al. (2010) showed that humans perceived taps as hopping progressively along an aluminum bar resting across the index fingers of the hands, when in actuality the taps were delivered only to the points on the bar directly above each finger. To apply the model to this scenario, it is necessary only to know the observer’s likelihood function evoked by a tap to the bar:  $p(\text{measurement} \mid \text{tap location along bar})$ . An interesting twist here is that both hands might detect any single tap to the bar. This does not preclude the construction of a likelihood function; it simply requires consideration of the sensory input to both hands. For instance, a more intense vibration felt with the right hand would result in a likelihood function whose peak lies to the right of the bar’s center. Once the single tap likelihood functions are determined empirically, it would be straightforward to fit the model to the behavioral data with a low-speed prior. Of interest would be to compare the value of  $\sigma_v$ , so obtained to the value ( $\sim 10$  cm/s) that fits the perception of trajectories delivered directly to the skin.

Our model provides insight into crossmodal interactions in length contraction illusions (Kawabe et al., 2008; Asai and Kanayama, 2012). In a 2-location, 3-tap rabbit paradigm, Asai and Kanayama (2012) demonstrated that the cutaneous rabbit was more consistently perceived when a visual flash occurred concurrently with, and at the typical illusory location of, the second tap. The model readily accommodates this cue-combination scenario. As shown in **Figure 6**, stochastic variability in the measurement causes trial-to-trial variability in the perceived location of either tap. Provided the Bayesian observer assumes that the concurrent visual and tactile measurements resulted independently from the same event, the observer’s likelihood function over that event’s location will be the product of the visual and tactile likelihoods. The visual measurement will therefore sharpen and shift the combined likelihood function toward the flash location, increasing the frequency with which the observer perceives the tactile stimulus to fall at that location. To test the model, one would first measure participants’ spatial uncertainty ( $\sigma_s$ ) in response to taps and flashes delivered in isolation. The model could then be used to

make testable predictions regarding the perceptual influence of the flash.

Finally, our model may account for saltation illusions in both vision (Geldard, 1976; Lockhead et al., 1980; Khuu et al., 2011) and audition (Bremer et al., 1977; Shore et al., 1998; Getzmann, 2009). Provided the brain expects visual and auditory stimuli to move slowly, the model predicts pronounced length contraction when stimulus sequences traverse areas of poor spatial acuity (high  $\sigma_s$ ). In vision, this prediction has already been confirmed: the visual rabbit illusion occurs in response to peripheral but not central stimuli (Geldard, 1976). Furthermore, a low-speed prior has been implicated in visual motion perception (Weiss et al., 2002; Stocker and Simoncelli, 2006). Future experimental studies will assess the quantitative fit of our model to visual and auditory saltation illusions.

Despite its apparently wide applicability, we do not suggest that a low-speed prior alone can account for a majority of motion illusions. Interestingly, several visual motion phenomena (Nijhawan, 2002; Hubbard, 2005) involve endpoint overestimation similar to that caused by the low-acceleration prior that did not match the tactile data considered here (**Figure 12B**). Research is needed to clarify the conditions under which perception incorporates a low-acceleration prior.

#### THE PERCEPT AS A COMBINED PRE- AND POST-DICTIVE INFERENCE

Our Bayesian observer's percept can be viewed as resulting from concomitant pre- and post-dictive inference. For instance, in two-tap trajectories, the first tap predicts the location of the second, while the second postdicts the location of the first (**Figure 9**). We suspect that Bayesian pre- and postdiction will be found to act together in many perceptual scenarios, whether or not these scenarios incorporate a low-speed prior. Indeed, it has already been reported that the two processes collaborate in the flash-lag effect (Rao et al., 2001; Soga et al., 2009), an illusion in which a brief visual flash placed alongside a moving object is perceived to lag behind the object.

By hypothesizing a link between spatial attention and  $\sigma_s$ , as suggested by point localization experiments (Moore et al., 1999; O'Boyle et al., 2001), we have shown how attention can shape the relative influence of pre- and postdiction on the percept (**Figure 9D**). When attention is directed around the location of the first tap ( $\sigma_{s1} < \sigma_{s2}$ ), prediction dominates, and the second tap is perceived as asymmetrically displaced toward the first. When attention is directed around the location of the second tap ( $\sigma_{s2} < \sigma_{s1}$ ), postdiction dominates, and the first tap is perceived as asymmetrically displaced toward the second. Under conditions of imbalanced spatial attention, the trajectory midpoint is therefore perceived as shifted toward the attended location, as specified by Eq. (5). As the spatial attention balance is adjusted from one extreme to another, the model smoothly transitions between a percept influenced predominantly by prediction to one influenced predominantly by postdiction.

Researchers have often referred to the rabbit illusion as a post-dictive phenomenon, without mentioning the involvement of prediction (Bays et al., 2006; Blankenburg et al., 2006; van Wassenhove, 2009; Miyazaki et al., 2010; Asai and Kanayama, 2012). Indeed, initial work on the rabbit described only the perceptual

displacement of the earlier tap(s) toward the later one(s) (Geldard and Sherrick, 1972), consistent with an exclusively postdictive process. However, it is clear from modern studies of the rabbit that both earlier and later taps undergo perceptual displacement – whether by equal distances or not (Kilgard and Merzenich, 1995; Flach and Haggard, 2006; Trojan et al., 2010). This supports our conclusion that the illusion involves concomitant predictive and postdictive inference.

Why did initial rabbit illusion investigations describe only the displacement of earlier taps toward later ones? In his three-tap “reduced rabbit” paradigm, Geldard (1982) stimulated with a “locator” (tap 1) followed at large ISI by an “attractee” (tap 2) at the same position, which he reported as perceptually displaced toward the subsequent “attractant” (tap 3) delivered at a different location. The participants' report that tap 2 was perceptually displaced toward tap 3, but not vice versa, may have owed to the absence of a second locator tap placed at the position of tap 3. Without a locator tap for spatial comparison, participants may have been unaware that tap 3 was perceptually displaced. This hypothesis was considered and discarded by Geldard (1982) upon preliminary investigation, but Kilgard and Merzenich (1995), using a 4-tap paradigm that included a second locator tap, did find symmetric perceptual displacement of taps 2 and 3 (**Figure 1C**).

Alternatively, as demonstrated by Kilgard and Merzenich (1995) and modeled here, asymmetric rabbit percepts could reflect an imbalance in spatial attention (**Figures 8 and 9D**; Eq. 5). An interesting possibility is that – particularly during multi-tap sequences – participants have time to redistribute their spatial attention on the fly. When investigators randomize the direction of movement (up or down along the arm), the participants cannot know where to expect the first tap, so they presumably distribute their spatial attention equally. After the first tap has occurred, however, experienced participants will know where the trajectory is heading, and might direct their attention fully toward the upcoming final location. This would cause a decrease in  $\sigma_s$  at the final location, consequently shifting the percept toward that point (e.g., **Figure 11D**).

#### SPECULATIONS REGARDING NEURAL IMPLEMENTATION

We have described two computational approaches by which our Bayesian observer could obtain its percept: either multi-dimensional inference (e.g., the two-dimensional inference shown in **Figure 9A**) or equivalent one-dimensional prediction-postdiction (**Figure 9B**). Which, if either, approach might the brain implement? The two approaches yield the same percept, but they scale very differently in difficulty as the number of taps increases. In the case of a sequence of  $n$  taps, the joint likelihood function, prior, and posterior would each require  $n$  dimensions. The neural representation of such multi-dimensional distributions would appear to pose considerable challenges. More plausibly, the brain could undertake one-dimensional predictive-postdictive inference recursively.

It is tempting to reinterpret the graphs in **Figure 9** as plots of activity (e.g., spike rates) of a series of cortical neurons that represent the corresponding skin positions ( $x$ -axes). Under this interpretation, the predicted prior is a mound of cortical neural activity evoked by tap 1 that decays and broadens over time

(Figure 9C). When the second tap initiates a second mound of cortical activity (the tap 2 likelihood function), the two mounds interact (e.g., through summation), resulting in a tap 2 percept that is shifted toward the tap 1 location. For trajectories with greater ISI, the tap 1 mound would have more time to decay, and would thus exert less influence over the tap 2 percept. This idea is similar to a model proposed by Flach and Haggard (2006). The idea is attractively simple; nevertheless, it seems able to account satisfactorily only for prediction, not postdiction. A more complex network model was proposed by Wiemer et al. (2000), but that model produces perceptual length dilation at large ISIs, a result contradicted by behavioral data.

Computationally, the perception of multi-tap sequences can be achieved with recursive predictive-postdictive Bayesian inference. The Kalman filter is an algorithm for recursive predictive inference (Haykin, 2001), for which plausible neural implementation schemes have been proposed (Deneve et al., 2007; Beck et al., 2011). Kalman smoothing combines the Kalman filter with recursive postdictive inference (Haykin, 2001). The percepts obtained by our Bayesian observer are identical to those that would result from an appropriately configured Kalman smoother (see “Multi-tap perception” in Appendix). Smoothing has already been implicated in the flash-lag effect (Rao et al., 2001) and proposed to contribute to a variety of motion illusions, including the rabbit (Grush, 2005), though to our knowledge a specific neural implementation for the Kalman smoother has not yet been proposed.

## TESTABLE PREDICTIONS

Our Bayesian observer model makes many testable predictions; we encourage other investigators to pursue these experimentally.

The model predicts that perceptual length contraction will be more pronounced on body areas with worse spatial acuity or – on a given body area – in response to stimuli that are harder to

localize (e.g., weaker taps to the skin). Because  $\sigma_s$  can be independently manipulated and measured using single taps, the length contraction formula (Eq. 1) can be used to make specific testable predictions regarding the effect of body area or stimulus strength on the perception of two-tap trajectories.

Under conditions of imbalanced spatial attention, the model predicts that perceptual length contraction will occur in accordance with Eq. 4 and that the midpoint of the perceived two-tap trajectory will vary in accordance with Eq. 5. These predictions could be tested experimentally by independently measuring an observer’s  $\sigma_{s1}$  and  $\sigma_{s2}$  under different degrees of directed spatial attention, then measuring the trajectory percepts under the same conditions.

As explained above, the model can be used to make testable predictions regarding a variety of perceptual length contraction phenomena beyond those that we have modeled in this paper. These include the out-of-body rabbit, crossmodal influences on the rabbit percept, and the visual and auditory rabbit illusions.

We encourage readers to generate their own predictions by using our freely downloadable computer program, Leaping Lagomorphs ([http://psych.mcmaster.ca/goldreich-lab/LL/Leaping\\_Lagomorphs.html](http://psych.mcmaster.ca/goldreich-lab/LL/Leaping_Lagomorphs.html)). This convenient program implements the Bayesian observer, with either balanced or imbalanced spatial attention, and outputs its perception in response to any stimulus sequence that the user cares to enter.

## ACKNOWLEDGMENTS

This research was supported by an Individual Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors thank Andy Bhattacharjee, Luxi Li, Ryan Peters, Mike Wong, and Deda Gillespie for their insightful comments.

## REFERENCES

- Anton-Erxleben, K., and Carrasco, M. (2013). Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nat. Rev. Neurosci.* 14, 188–200.
- Asai, T., and Kanayama, N. (2012). “Cutaneous rabbit” hops toward a light: unimodal and cross-modal causality on the skin. *Front. Psychol.* 3:427. doi:10.3389/fpsyg.2012.00427
- Bays, P. M., Flanagan, J. R., and Wolpert, D. M. (2006). Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biol.* 4:e28. doi:10.1371/journal.pbio.0040028
- Beck, J. M., Latham, P. E., and Pouget, A. (2011). Marginalization in neural circuits with divisive normalization. *J. Neurosci.* 31, 15310–15319.
- Blankenburg, F., Ruff, C. C., Deichmann, R., Rees, G., and Driver, J. (2006). The cutaneous rabbit illusion affects human primary sensory cortex somatotopically. *PLoS Biol.* 4:e69. doi:10.1371/journal.pbio.0040069
- Bremer, C. D., Pittenger, J. B., Warren, R., and Jenkins, J. J. (1977). An illusion of auditory saltation similar to the cutaneous “rabbit.” *Am. J. Psychol.* 90, 645–654.
- Brigner, W. L. (1988). Saltation as a rotation of space-time axes. *Percept. Mot. Skills* 66, 637–638.
- Cody, F. W., Garside, R. A., Lloyd, D., and Poliakoff, E. (2008). Tactile spatial acuity varies with site and axis in the human upper limb. *Neurosci. Lett.* 433, 103–108.
- Deneve, S., Duhamel, J. R., and Pouget, A. (2007). Optimal sensorimotor integration in recurrent cortical networks: a neural implementation of Kalman filters. *J. Neurosci.* 27, 5744–5756.
- Dennett, D. C., and Kinsbourne, M. (1992). Time and the observer: the where and when of consciousness in the brain. *Behav. Brain Sci.* 15, 183–201.
- Eagleman, D. M., and Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science* 287, 2036–2038.
- Flach, R., and Haggard, P. (2006). The cutaneous rabbit revisited. *J. Exp. Psychol. Hum. Percept. Perform.* 32, 717–732.
- Gelb, A. (1914). “Versuche auf dem Gebiete der Zeit- und Raumschauung,” in *Bericht Über Den VI. Kongress für Experimentelle Psychologie: in Göttingen April 1914*, ed. F. Schumann (Leipzig: J. A. Barth), 36–42.
- Geldard, F. A. (1976). The saltatory effect in vision. *Sens. Processes* 1, 77–86.
- Geldard, F. A. (1982). Saltation in somesthesia. *Psychol. Bull.* 92, 136–175.
- Geldard, F. A., and Sherrick, C. E. (1972). The cutaneous “rabbit”: a perceptual illusion. *Science* 178, 178–179.
- Getzmann, S. (2009). Exploring auditory saltation using the “reduced-rabbit” paradigm. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 289–304.
- Goldreich, D. (2007). A Bayesian perceptual model replicates the cutaneous rabbit and other tactile spatiotemporal illusions. *PLoS ONE* 2:e333. doi:10.1371/journal.pone.0000333
- Grush, R. (2005). Internal models and the construction of time: generalizing from state estimation to trajectory estimation to address temporal features of perception, including temporal illusions. *J. Neural Eng.* 2, S209–218.
- Haykin, S. S. (2001). “Kalman filters,” in *Kalman Filtering and Neural Networks*, ed. S. S. Haykin (New York: Wiley), 1–21.
- Helmholtz, H. V. (1925). *Treatise on Physiological Optics, III: The Perceptions of Vision (1910)*. Rochester, NY: Optical Society of America.
- Helson, H. (1930). The tau effect – an example of psychological relativity. *Science* 71, 536–537.



- Helson, H., and King, S. M. (1931). The tau effect: an example of psychological relativity. *J. Exp. Psychol.* 14, 202–217.
- Higashiyama, A., and Hayashi, M. (1993). Localization of electrocutaneous stimuli on the fingers and forearm: effects of electrode configuration and body axis. *Percept. Psychophys.* 54, 108–120.
- Hubbard, T. L. (2005). Representational momentum and related displacements in spatial memory: a review of the findings. *Psychon. Bull. Rev.* 12, 822–851.
- Jones, B., and Huang, Y. L. (1982). Space-time dependencies in psychological judgment of extent and duration: algebraic models of the tau and kappa effects. *Psychol. Bull.* 91, 128–142.
- Kawabe, T., Miura, K., and Yamada, Y. (2008). Audiovisual tau effect. *Acta Psychol. (Amst.)* 128, 249–254.
- Khuu, S. K., Kidd, J. C., and Badcock, D. R. (2011). The influence of spatial orientation on the perceived path of visual saltatory motion. *J. Vis.* 11, ii:5.
- Kilgard, M. P., and Merzenich, M. M. (1995). Anticipated stimuli across skin. *Nature* 373, 663.
- Knill, D. C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719.
- Lechelt, E. C., and Borchert, R. (1977). The interdependence of time and space in somesthesia: the Tau effect reexamined. *Bull. Psychon. Soc.* 10, 191–193.
- Lockhead, G. R., Johnson, R. C., and Gold, F. M. (1980). Saltation through the blind spot. *Percept. Psychophys.* 27, 545–549.
- Marks, L. E., Girvin, J. P., Quest, D. O., Antunes, J. L., Ning, P., O'Keefe, M. D., et al. (1982). Electrocutaneous stimulation II. The estimation of distance between two points. *Percept. Psychophys.* 32, 529–536.
- Martikainen, I. K., and Pertovaara, A. (2002). Spatial discrimination of one versus two test stimuli in the human skin: dissociation of mechanisms depending on the task and the modality of stimulation. *Neurosci. Lett.* 328, 322–324.
- Miyazaki, M., Hirashima, M., and Nozaki, D. (2010). The “cutaneous rabbit” hopping out of the body. *J. Neurosci.* 30, 1856–1860.
- Moore, C. E., Partner, A., and Sedgwick, E. M. (1999). Cortical focusing is an alternative explanation for improved sensory acuity on an amputation stump. *Neurosci. Lett.* 270, 185–187.
- Nijhawan, R. (2002). Neural delays, visual motion and the flash-lag effect. *Trends Cogn. Sci. (Regul. Ed.)* 6, 387.
- O'Boyle, D. J., Moore, C. E., Poliakoff, E., Butterworth, R., Sutton, A., and Cody, F. W. (2001). Human locognosic acuity on the arm varies with explicit and implicit manipulations of attention: implications for interpreting elevated tactile acuity on an amputation stump. *Neurosci. Lett.* 305, 37–40.
- Rao, R. P., Eagleman, D. M., and Sejnowski, T. J. (2001). Optimal smoothing in visual motion perception. *Neural Comput.* 13, 1243–1253.
- Shore, D. I., Hall, S. E., and Klein, R. M. (1998). Auditory saltation: a new measure for an old illusion. *J. Acoust. Soc. Am.* 103, 3730–3733.
- Soga, R., Akaishi, R., and Sakai, K. (2009). Predictive and postdictive mechanisms jointly contribute to visual awareness. *Conscious. Cogn.* 18, 578–592.
- Sripati, A. P., Yoshioka, T., Denchev, P., Hsiao, S. S., and Johnson, K. O. (2006). Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system. *J. Neurosci.* 26, 2101–2114.
- Stocker, A. A., and Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci.* 9, 578–585.
- Trojan, J., Stolle, A. M., Carl, A. M., Kleinbohl, D., Tan, H. Z., and Holzl, R. (2010). Spatiotemporal integration in somatosensory perception: effects of sensory saltation on pointing at perceived positions on the body surface. *Front. Psychol.* 1:206. doi:10.3389/fpsyg.2010.00206
- van Wassenhove, V. (2009). Minding time in an amodal representational space. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 364, 1815–1830.
- Vilares, I., and Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Ann. N. Y. Acad. Sci.* 1224, 22–39.
- Weinstein, S. (1968). “Intensive and extensive aspects of tactile sensitivity as a function of body part, sex, and laterality,” in *The Skin Senses: Proceedings*, ed. D. R. Kenshalo (Springfield, IL: Thomas), 195–222.
- Weiss, Y., Simoncelli, E. P., and Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nat. Neurosci.* 5, 598–604.
- Wiemer, J., Spengler, F., Joublin, F., Stagge, P., and Wacquant, S. (2000). Learning cortical topography from spatiotemporal stimuli. *Biol. Cybern.* 82, 173–187.

**Conflict of Interest Statement:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 20 March 2013; accepted: 11 April 2013; published online: 10 May 2013.

Citation: Goldreich D and Tong J (2013) Prediction, postdiction, and perceptual length contraction: a Bayesian low-speed prior captures the cutaneous rabbit and related illusions. *Front. Psychol.* 4:221. doi: 10.3389/fpsyg.2013.00221

This article was submitted to *Frontiers in Consciousness Research, a specialty of Frontiers in Psychology*.

Copyright © 2013 Goldreich and Tong. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in other forums, provided the original authors and source are credited and subject to any copyright notices concerning any third-party graphics etc.

## APPENDIX

Here, we further develop mathematically, and offer new conceptual insights into, the *basic Bayesian observer* model put forth by Goldreich (2007). In the following seven sections, we: 1) specify the observer's generative model, and derive the posterior probability density over tap trajectories and the perceptual length contraction formula; 2) generalize the derivation to include inhomogeneous spatial acuity caused by selective spatial attention; 3) consider useful one-dimensional reductions of the two-dimensional posterior density; 4) reformulate the observer's percept as a combined predictive-postdictive inference; 5) model the perception of multi-tap sequences; 6) consider extensions of the model that incorporate additional sources of uncertainty; and 7) describe how we fit the model to human perceptual data.

### THE BAYESIAN MODEL

We consider here an observer whose goal is to perceive the locations of two-taps delivered to the skin in rapid succession. We assume that the observer has an internal generative model – a conception of the statistics of moving tactile stimuli – and that it interprets the stimulus sequence optimally within the context of its generative model. Briefly, the observer considers two taps that occur in rapid succession to result from a single moving object, and it considers that tactile objects tend to move slowly. Specifically, according to the generative model: (1) An object briefly touches the skin at a location,  $x_1$ , drawn from a uniform density. (2) The object moves away from  $x_1$  with velocity  $v$ , drawn from a Gaussian density with mean zero and standard deviation  $\sigma_v$ ; at some elapsed time  $t$  (independent of  $x_1$ ), the object again briefly touches the skin, at location  $x_2$ . (3) Noisy sensorineural activity evoked by each tap results in measured values for the tap positions,  $x_{1m}$  and  $x_{2m}$ , drawn from Gaussian densities centered on the actual tap positions,  $x_1$  and  $x_2$ , with standard deviations  $\sigma_s$ .

#### Bayes' formula

The observer's goal is to infer the positions of the taps ( $x_1, x_2$ ), which we refer to as the movement trajectory. We assume in this basic model that the observer perceives the time between taps,  $t$ , veridically. Thus, the observer knows  $x_{1m}$ ,  $x_{2m}$ , and  $t$ , and wishes to infer  $x_1$  and  $x_2$ . According to Bayes' formula, the posterior over trajectories is proportional to the product of likelihood and prior:

$$p(x_1, x_2 | x_{1m}, x_{2m}, t) \propto p(x_{1m}, x_{2m} | x_1, x_2, t) p(x_1, x_2 | t) \quad (\text{A1})$$

We now work out the observer's prior and likelihood.

#### Prior probability density

The observer's prior probability density over trajectories is:

$$p(x_1, x_2 | t) = p(x_2 | x_1, t) p(x_1 | t) \quad (\text{A2})$$

Because  $t$  and  $x_1$  are independent,  $p(x_1 | t) = p(x_1)$ , and this is a constant ( $x_1$  being drawn from a uniform distribution). Therefore, we can write more concisely:

$$p(x_1, x_2 | t) \propto p(x_2 | x_1, t) \quad (\text{A3})$$

We note that, given  $x_1$  and  $t$ ,  $x_2$  is a function of the velocity,  $v$ :

$$x_2 = x_1 + vt \quad (\text{A4})$$

Thus, the probability that  $v$  resides in the infinitesimal region ( $v \pm \frac{dv}{2}$ ) is equal to the probability that  $x_2$  resides in the corresponding infinitesimal region ( $x_2 \pm \frac{dx_2}{2}$ ):

$$p(x_2 | x_1, t) dx_2 = p(v) dv \quad (\text{A5})$$

It follows that:

$$p(x_2 | x_1, t) = p(v) \left| \frac{dv}{dx_2} \right| = \frac{p(v)}{t} \quad (\text{A6})$$

Now recall that the observer has a low-velocity prior expectation:

$$p(v) = \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{v^2}{2\sigma_v^2}\right) = \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{((x_2 - x_1)/t)^2}{2\sigma_v^2}\right) \quad (\text{A7})$$

Referring to Eqs A3, A6, and A7, we therefore have:

$$p(x_1, x_2|t) \propto p(x_2|x_1, t) = \frac{1}{\sqrt{2\pi\sigma_v t}} \exp\left(-\frac{(x_2 - x_1)^2}{2(\sigma_v t)^2}\right) \quad (\text{A8})$$

The observer's prior probability density over trajectories is proportional to a Gaussian distribution over the distance between taps, with mean zero and standard deviation  $\sigma_v t$ . Reflecting the low-speed prior, when the elapsed time,  $t$ , is large, a wide range of displacements is permissible; when  $t$  is shorter, the observer expects the two taps to more closely coincide spatially.

For future reference, we note that  $x_2$ , like  $x_1$ , is independent of  $t$ . We see this by integrating Eq. A8 with respect to  $x_1$ :

$$p(x_2|t) = \int_{x_1} p(x_1, x_2|t) dx_1 \propto \int_{x_1} \frac{1}{\sqrt{2\pi\sigma_v t}} \exp\left(-\frac{(x_2 - x_1)^2}{2(\sigma_v t)^2}\right) dx_1 = 1 \quad (\text{A9})$$

Thus,  $x_2$  is independent of  $t$ , and, like  $p(x_1)$ ,  $p(x_2)$  is a constant. Eq. A8 shows that  $x_2$  is conditionally dependent on  $t$ , given  $x_1$ .

### Likelihood function

The tap positions measured by the observer,  $x_{1m}$  and  $x_{2m}$ , are drawn independently from Gaussian densities centered on the actual tap positions, with standard deviations  $\sigma_s$ . Therefore, the observer's likelihood function is:

$$p(x_{1m}, x_{2m}|x_1, x_2, t) = p(x_{1m}|x_1) p(x_{2m}|x_2) \quad (\text{A10})$$

where

$$p(x_{1m}|x_1) = \frac{1}{\sqrt{2\pi\sigma_s}} \exp\left(-\frac{(x_{1m} - x_1)^2}{2\sigma_s^2}\right) \quad p(x_{2m}|x_2) = \frac{1}{\sqrt{2\pi\sigma_s}} \exp\left(-\frac{(x_{2m} - x_2)^2}{2\sigma_s^2}\right) \quad (\text{A11})$$

### Posterior probability density

The observer uses Bayes' formula (Eq. A1) to calculate the posterior density over trajectories. It is useful to express the posterior density in several ways. First, referring to Eqs A3 and A10, we see that Bayes' formula can be rewritten:

$$p(x_1, x_2|x_{1m}, x_{2m}, t) \propto p(x_{1m}|x_1) p(x_{2m}|x_2) p(x_2|x_1, t) \quad (\text{A12})$$

Next, from Eqs A8 and A11, we have

$$p(x_1, x_2|x_{1m}, x_{2m}, t) \propto \exp\left(-\left(\frac{(x_{1m} - x_1)^2}{2\sigma_s^2} + \frac{(x_2 - x_1)^2}{2(\sigma_v t)^2}\right)\right) \quad (\text{A13})$$

Finally, following some rearrangement, Eq. A13 can be written as a two-dimensional (2D) Gaussian distribution

$$p(x_1, x_2|x_{1m}, x_{2m}, t) \propto \exp\left(-\frac{1}{2(1-\rho^2)} \left(\frac{(x_1 - x_1^*)^2}{\sigma^2} + (x_2 - x_2^*)^2 - 2\rho(x_1 - x_1^*)(x_2 - x_2^*)\right)\right) \quad (\text{A14})$$

where the posterior mode  $(x_1^*, x_2^*)$  is given by

$$x_1^* = x_{1m} \left(\frac{(\sigma_v t)^2 + \sigma_s^2}{(\sigma_v t)^2 + 2\sigma_s^2}\right) + x_{2m} \left(\frac{\sigma_s^2}{(\sigma_v t)^2 + 2\sigma_s^2}\right)$$

$$x_2^* = x_{1m} \left(\frac{\sigma_s^2}{(\sigma_v t)^2 + 2\sigma_s^2}\right) + x_{2m} \left(\frac{(\sigma_v t)^2 + \sigma_s^2}{(\sigma_v t)^2 + 2\sigma_s^2}\right)$$

and the variance ( $\sigma^2$ ) and correlation coefficient ( $\rho$ ) are given by:

$$\sigma^2 = \sigma_s^2 \frac{\sigma_s^2 + (\sigma_v t)^2}{2\sigma_s^2 + (\sigma_v t)^2} \quad \rho = \frac{\sigma_s^2}{\sigma_s^2 + (\sigma_v t)^2}$$

We assume that the observer reads out the posterior mode as the percept. Note that the perceived positions,  $x_1^*$  and  $x_2^*$ , are weighted averages of the measurements,  $x_{1m}$  and  $x_{2m}$ . The perceived positions are drawn toward one another as the time between taps shortens,

converging toward the measurement midpoint,  $(x_{1m} + x_{2m})/2$ , in the limit that  $t$  approaches zero. As  $t$  approaches infinity, by contrast,  $x_{1*}$  and  $x_{2*}$  approach the measured values,  $x_{1m}$  and  $x_{2m}$ .

Subtracting  $x_{1*}$  from  $x_{2*}$ , we find that the perceived distance between taps,  $l^* = x_{2*} - x_{1*}$ , relates to the measured distance,  $l_m = x_{2m} - x_{1m}$ , according to the formula:

$$l^* = x_{2*} - x_{1*} = \frac{x_{2m} - x_{1m}}{1 + 2\left(\frac{\sigma_s}{\sigma_v t}\right)^2} = \frac{l_m}{1 + 2\left(\frac{\tau}{t}\right)^2} \quad (\text{A15})$$

where we have defined the parameter tau as the ratio of the observer's spatial uncertainty to the width of the low-speed prior:  $\tau = \frac{\sigma_s}{\sigma_v}$ .

Although the measured tap positions will vary stochastically from trial to trial, on average they will equal the actual tap positions. Thus, on average the perceived distance is related to the true distance,  $l$ , as:

$$l^* = \frac{l}{1 + 2\left(\frac{\tau}{t}\right)^2} \quad (\text{A16})$$

This is the perceptual length contraction formula, previously derived – using a different approach and expressed in a slightly different form – by Goldreich (2007).

### GENERALIZATION TO INHOMOGENEOUS SPATIAL UNCERTAINTY

So far we have assumed equal spatial uncertainty,  $\sigma_s$ , at each point on the skin. Here, we consider the more general situation in which each tap may be associated with a different spatial uncertainty,  $\sigma_{s1}$  and  $\sigma_{s2}$ , as might occur if the participant were to focus spatial attention on one skin region. In this case, the likelihood functions, Eq. A11, become:

$$p(x_{1m}|x_1) = \frac{1}{\sqrt{2\pi}\sigma_{s1}} \exp\left(-\frac{(x_{1m} - x_1)^2}{2\sigma_{s1}^2}\right) \quad p(x_{2m}|x_2) = \frac{1}{\sqrt{2\pi}\sigma_{s2}} \exp\left(-\frac{(x_{2m} - x_2)^2}{2\sigma_{s2}^2}\right) \quad (\text{A17})$$

Consequently, the posterior density over tap positions (Eq. A13) becomes

$$p(x_1, x_2|x_{1m}, x_{2m}, t) \propto \exp\left(-\left(\frac{(x_{1m} - x_1)^2}{2\sigma_{s1}^2} + \frac{(x_{2m} - x_2)^2}{2\sigma_{s2}^2} + \frac{(x_2 - x_1)^2}{2(\sigma_v t)^2}\right)\right) \quad (\text{A18})$$

Following rearrangement, Eq. A18 can be re-written as a 2D Gaussian distribution,

$$p(x_1, x_2|x_{1m}, x_{2m}, t) \propto \exp\left(-\frac{1}{2(1-\rho^2)}\left(\frac{(x_1 - x_{1*})^2}{\sigma_1^2} + \frac{(x_2 - x_{2*})^2}{\sigma_2^2} - \frac{2\rho(x_1 - x_{1*})(x_2 - x_{2*})}{\sigma_1\sigma_2}\right)\right) \quad (\text{A19})$$

where the posterior mode  $(x_{1*}, x_{2*})$  is given by

$$x_{1*} = x_{1m} \left(\frac{(\sigma_v t)^2 + \sigma_{s2}^2}{(\sigma_v t)^2 + \sigma_{s1}^2 + \sigma_{s2}^2}\right) + x_{2m} \left(\frac{\sigma_{s1}^2}{(\sigma_v t)^2 + \sigma_{s1}^2 + \sigma_{s2}^2}\right)$$

$$x_{2*} = x_{1m} \left(\frac{\sigma_{s2}^2}{(\sigma_v t)^2 + \sigma_{s1}^2 + \sigma_{s2}^2}\right) + x_{2m} \left(\frac{(\sigma_v t)^2 + \sigma_{s1}^2}{(\sigma_v t)^2 + \sigma_{s1}^2 + \sigma_{s2}^2}\right)$$

and the variances  $(\sigma_1^2, \sigma_2^2)$  and correlation coefficient ( $\rho$ ) are given by:

$$\sigma_1^2 = \sigma_{s1}^2 \frac{\sigma_{s2}^2 + (\sigma_v t)^2}{\sigma_{s1}^2 + \sigma_{s2}^2 + (\sigma_v t)^2} \quad \sigma_2^2 = \sigma_{s2}^2 \frac{\sigma_{s1}^2 + (\sigma_v t)^2}{\sigma_{s1}^2 + \sigma_{s2}^2 + (\sigma_v t)^2} \quad \rho = \frac{\sigma_{s1}\sigma_{s2}}{\sqrt{(\sigma_{s1}^2 + (\sigma_v t)^2)(\sigma_{s2}^2 + (\sigma_v t)^2)}}$$

It follows that

$$l^* = x_{2*} - x_{1*} = \frac{l_m}{1 + \frac{\sigma_{s1}^2 + \sigma_{s2}^2}{(\sigma_v t)^2}} = \frac{l_m}{1 + 2\left(\frac{\sigma_{s(\text{rms})}}{\sigma_v t}\right)^2} \quad (\text{A20})$$

Thus, the uniform spatial uncertainty,  $\sigma_s$ , of Eq. A15 is replaced by the root-mean-square of the uncertainty at the two locations:

$$\sigma_{s(\text{rms})} = \sqrt{\frac{\sigma_{s1}^2 + \sigma_{s2}^2}{2}}.$$

Interestingly, when  $\sigma_{s1} \neq \sigma_{s2}$ , the midpoint of the perceived trajectory no longer coincides with the midpoint of the measured trajectory. From the expressions (Eq. A19) for  $x_{1*}$  and  $x_{2*}$  it is easily shown that the shift,  $\Delta_{\text{midpt}}$ , in the perceived trajectory midpoint away from the measured trajectory midpoint is:

$$\Delta_{\text{midpt}} = \frac{x_{1*} + x_{2*}}{2} - \frac{x_{1m} + x_{2m}}{2} = \frac{l_m}{2} \left( \frac{\sigma_{s1}^2 - \sigma_{s2}^2}{(\sigma_v t)^2 + \sigma_{s1}^2 + \sigma_{s2}^2} \right) \quad (\text{A21})$$

### ONE-DIMENSIONAL REDUCTIONS

The two-dimensional joint  $(x_1, x_2)$  posterior density (Eq. A19) fully represents the observer's belief distribution over stimulus trajectories, and it captures dependencies between the variables. Nevertheless, it can be useful to express the observer's belief about a single parameter of interest, although this entails a loss of information about dependencies. One such parameter of interest is the length,  $l$ , between taps. Other parameters of interest are the tap positions,  $x_1$  and  $x_2$ , considered individually. Here we derive the observer's one-dimensional posterior densities over each of these parameters.

#### Posterior density over trajectory length

The posterior over trajectory length,  $l = x_2 - x_1$ , can be found by integrating across the joint posterior:

$$p(l|x_{1m}, x_{2m}, t) = \int_{x_1} p(x_1, x_2 = l + x_1 | x_{1m}, x_{2m}, t) dx_1 \quad (\text{A22})$$

The posterior over  $l$  can also be found by noting that, from Eq. A8, the observer's prior over  $l$  is:

$$p(l|t) = \frac{1}{\sqrt{2\pi}\sigma_v t} \exp\left(-\frac{l^2}{2(\sigma_v t)^2}\right) \quad (\text{A23})$$

Further, from Eq. A17, we see that the observer's displacement measurement,  $l_m = x_{2m} - x_{1m}$ , is normally distributed with mean  $l$  and variance  $\sigma_{s1}^2 + \sigma_{s2}^2$ :

$$p(l_m|l) = \frac{1}{\sqrt{2\pi(\sigma_{s1}^2 + \sigma_{s2}^2)}} \exp\left(-\frac{(l_m - l)^2}{2(\sigma_{s1}^2 + \sigma_{s2}^2)}\right) \quad (\text{A24})$$

Thus, by Bayes' rule, the posterior over  $l$  is proportional to the product of these two Gaussian densities:

$$p(l|l_m, t) \propto p(l_m|l, t) p(l|t) \quad (\text{A25})$$

The result is a Gaussian posterior density with mean and variance given by:

$$\mu_{l \text{ posterior}} = \frac{l_m}{1 + \frac{\sigma_{s1}^2 + \sigma_{s2}^2}{(\sigma_v t)^2}}, \quad \sigma_{l \text{ posterior}}^2 = \frac{1}{\frac{1}{\sigma_{s1}^2 + \sigma_{s2}^2} + \frac{1}{(\sigma_v t)^2}} \quad (\text{A26})$$

The mean of the posterior over  $l$  is again the length contraction formula, Eq. A20. The variance of the posterior over  $l$  is smaller than the variance of  $l_m$ , given  $l$ . For this reason, the observer's length percept is more accurate than the length measurement (see **Figure 7**).

#### Marginal posterior densities over $x_1$ and $x_2$

To express the observer's belief about each tap's position individually, we can integrate the joint posterior along  $x_2$  to find the marginal posterior over  $x_1$ , and integrate the joint posterior along  $x_1$  to find the marginal posterior over  $x_2$ :

$$p(x_1|x_{1m}, x_{2m}, t) = \int_{x_2} p(x_1, x_2|x_{1m}, x_{2m}, t) dx_2 \quad (\text{A27})$$

$$p(x_2|x_{1m}, x_{2m}, t) = \int_{x_1} p(x_1, x_2|x_{1m}, x_{2m}, t) dx_1$$

Because the joint posterior density is a 2D Gaussian (Eq. A19), the marginalization integrals (Eq. A27) have simple solutions:

$$\begin{aligned} p(x_1|x_{1m}, x_{2m}, t) &= \frac{1}{\sqrt{2\pi}\sigma_1} \exp\left(-\frac{(x_1 - x_{1*})^2}{2\sigma_1^2}\right) \\ p(x_2|x_{1m}, x_{2m}, t) &= \frac{1}{\sqrt{2\pi}\sigma_2} \exp\left(-\frac{(x_2 - x_{2*})^2}{2\sigma_2^2}\right) \end{aligned} \quad (\text{A28})$$

### THE PREDICTION-POSTDICTION FORMULATION

Here, we show that the observer's marginal posterior over  $x_2$  can be equivalently derived from predictive inference: upon observing tap 1, the observer predicts (infers forward in time) a prior over tap 2; the observer then combines this *predicted prior* with the tap 2 likelihood to obtain the posterior over  $x_2$ . Conversely, the marginal posterior over  $x_1$  can be derived from postdictive inference: upon observing tap 2, the observer postdicts (infers backward in time) a prior over tap 1; the observer then combines this *postdicted prior* with the tap 1 likelihood to obtain the posterior over  $x_1$ .

#### Predicting tap 2 upon observing tap 1

Replacing the integrand in lower Eq. A27 with the expression from Eq. A1, we have:

$$p(x_2|x_{1m}, x_{2m}, t) \propto \int_{x_1} p(x_{1m}, x_{2m}|x_1, x_2, t) p(x_1, x_2|t) dx_1 \quad (\text{A29})$$

Further expanding the integrand, we have:

$$p(x_2|x_{1m}, x_{2m}, t) \propto \int_{x_1} p(x_{1m}|x_1) p(x_{2m}|x_2) p(x_2|x_1, t) p(x_1) dx_1 \quad (\text{A30})$$

Because  $p(x_{2m}|x_2)$  does not depend on  $x_1$ , we move it outside the integral. Thus, we have:

$$p(x_2|x_{1m}, x_{2m}, t) \propto p(x_{2m}|x_2) \int_{x_1} p(x_{1m}|x_1) p(x_2|x_1, t) p(x_1) dx_1 \quad (\text{A31})$$

Now we note that, according to Bayes' formula:

$$p(x_{1m}|x_1) p(x_1) \propto p(x_1|x_{1m}) \quad (\text{A32})$$

Substituting Eq. A32 into Eq. A31 yields:

$$p(x_2|x_{1m}, x_{2m}, t) \propto p(x_{2m}|x_2) \int_{x_1} p(x_2|x_1, t) p(x_1|x_{1m}) dx_1 \quad (\text{A33})$$

Equation A33 is Bayes' formula for the tap 2 position,  $x_2$ . It states that the marginal posterior density over  $x_2$  is proportional to the product of the tap 2 likelihood,  $p(x_{2m}|x_2)$ , and the tap 2 *predicted prior density*,

$$p(x_2|x_{1m}, t) = \int_{x_1} p(x_2|x_1, t) p(x_1|x_{1m}) dx_1 \quad (\text{A34})$$

The predicted prior projects belief forwards in time. It reflects the observer's beliefs about tap 2, given the tap 1 measurement and the elapsed time. Based on  $x_{1m}$ , the observer can generate a posterior over tap 1,  $p(x_1|x_{1m})$ . The predicted prior over a particular tap 2 position is then calculated by integrating across every possible tap 1 the product of this tap 1 posterior with the probability that the particular tap 2 will follow.

**Postdicting tap 1 upon observing tap 2**

Replacing the integrand in upper Eq. A27 with the expression from Eq. A1, we have:

$$p(x_1|x_{1m}, x_{2m}, t) \propto \int_{x_2} p(x_{1m}, x_{2m}|x_1, x_2, t) p(x_1, x_2|t) dx_2 \quad (\text{A35})$$

Further expanding the integrand, we have:

$$p(x_1|x_{1m}, x_{2m}, t) \propto \int_{x_2} p(x_{1m}|x_1) p(x_{2m}|x_2) p(x_1|x_2, t) p(x_2) dx_2 \quad (\text{A36})$$

Because  $p(x_{1m}|x_1)$  does not depend on  $x_2$ , we move it outside the integral. Thus, we have:

$$p(x_1|x_{1m}, x_{2m}, t) \propto p(x_{1m}|x_1) \int_{x_2} p(x_{2m}|x_2) p(x_1|x_2, t) p(x_2) dx_2 \quad (\text{A37})$$

Now we note that, according to Bayes' formula:

$$p(x_{2m}|x_2) p(x_2) \propto p(x_2|x_{2m}) \quad (\text{A38})$$

Substituting Eq. A38 into Eq. A37 yields:

$$p(x_1|x_{1m}, x_{2m}, t) \propto p(x_{1m}|x_1) \int_{x_2} p(x_1|x_2, t) p(x_2|x_{2m}) dx_2 \quad (\text{A39})$$

Equation A39 is Bayes' formula for the tap 1 position,  $x_1$ . It states that the marginal posterior density over  $x_1$  is proportional to the product of the tap 1 likelihood,  $p(x_{1m}|x_1)$ , and the tap 1 *postdicted prior density*,

$$p(x_1|x_{2m}, t) = \int_{x_2} p(x_1|x_2, t) p(x_2|x_{2m}) dx_2 \quad (\text{A40})$$

The postdicted prior projects belief backwards in time. It reflects the observer's beliefs about tap 1, given the tap 2 measurement and the elapsed time. Based on  $x_{2m}$ , the observer can generate a posterior over tap 2,  $p(x_2|x_{2m})$ . The postdicted prior over a particular tap 1 position is then calculated by integrating across every possible tap 2 the product of this tap 2 posterior with the probability that the particular tap 1 preceded.

**Formulas for the predicted and postdicted prior densities**

We now solve the predicted and postdicted prior integrals (Eqs A34 and A40). To find the predicted prior, we substitute from Eqs A8 and A17 left, into Eq. A34:

$$\begin{aligned} p(x_2|x_{1m}, t) &= \int_{x_1} \frac{1}{\sqrt{2\pi}\sigma_v t} \exp\left(-\frac{(x_2 - x_1)^2}{2(\sigma_v t)^2}\right) \frac{1}{\sqrt{2\pi}\sigma_{s1}} \exp\left(-\frac{(x_{1m} - x_1)^2}{2\sigma_{s1}^2}\right) dx_1 \\ &= \frac{1}{2\pi\sigma_v t \sigma_{s1}} \int_{x_1} \exp\left[-\left(\frac{(x_2 - x_1)^2}{2(\sigma_v t)^2} + \frac{(x_{1m} - x_1)^2}{2\sigma_{s1}^2}\right)\right] dx_1 \end{aligned} \quad (\text{A41})$$

We note that, upon much rearrangement:

$$\frac{(x_2 - x_1)^2}{2(\sigma_v t)^2} + \frac{(x_{1m} - x_1)^2}{2\sigma_{s1}^2} = \frac{(\sigma_v t)^2 + \sigma_{s1}^2}{2\sigma_{s1}^2(\sigma_v t)^2} \left(x_1 - \frac{x_2\sigma_{s1}^2 + x_{1m}(\sigma_v t)^2}{(\sigma_v t)^2 + \sigma_{s1}^2}\right)^2 + \frac{1}{2} \left(\frac{(x_2 - x_{1m})^2}{(\sigma_v t)^2 + \sigma_{s1}^2}\right) \quad (\text{A42})$$

Thus, Eq. A41 becomes,

$$p(x_2|x_{1m}, t) = \frac{1}{2\pi\sigma_v t \sigma_{s1}} \exp\left(-\frac{(x_2 - x_{1m})^2}{2((\sigma_v t)^2 + \sigma_{s1}^2)}\right) \int_{x_1} \exp\left(-\frac{\left(x_1 - \frac{x_2\sigma_{s1}^2 + x_{1m}(\sigma_v t)^2}{(\sigma_v t)^2 + \sigma_{s1}^2}\right)^2}{\frac{2\sigma_{s1}^2(\sigma_v t)^2}{(\sigma_v t)^2 + \sigma_{s1}^2}}\right) dx_1 \quad (\text{A43})$$

The integrand is a Gaussian function with standard deviation

$$\frac{\sigma_{s1}\sigma_v t}{\sqrt{(\sigma_v t)^2 + \sigma_{s1}^2}}.$$

Because the integral of an un-normalized Gaussian function of standard deviation  $\sigma$  is  $\sqrt{2\pi}\sigma$ , Eq. A43 simplifies to:

$$p(x_2|x_{1m}, t) = \frac{1}{2\pi\sigma_v t\sigma_{s1}} \exp\left(-\frac{(x_2 - x_{1m})^2}{2((\sigma_v t)^2 + \sigma_{s1}^2)}\right) \frac{\sqrt{2\pi}\sigma_{s1}\sigma_v t}{\sqrt{(\sigma_v t)^2 + \sigma_{s1}^2}} \quad (\text{A44})$$

Therefore, the predicted prior density over  $x_2$  is

$$p(x_2|x_{1m}, t) = \frac{1}{\sqrt{2\pi((\sigma_v t)^2 + \sigma_{s1}^2)}} \exp\left(-\frac{(x_2 - x_{1m})^2}{2((\sigma_v t)^2 + \sigma_{s1}^2)}\right) \quad (\text{A45})$$

That is, the predicted prior is a Gaussian with mean and variance

$$\mu_{\text{pre}} = x_{1m} \quad \sigma_{\text{pre}}^2 = (\sigma_v t)^2 + \sigma_{s1}^2 \quad (\text{A46})$$

A similar derivation reveals that the postdicted prior density over  $x_1$  is

$$p(x_1|x_{2m}, t) = \frac{1}{\sqrt{2\pi((\sigma_v t)^2 + \sigma_{s2}^2)}} \exp\left(-\frac{(x_1 - x_{2m})^2}{2((\sigma_v t)^2 + \sigma_{s2}^2)}\right) \quad (\text{A47})$$

That is, the postdicted prior is a Gaussian with mean and variance

$$\mu_{\text{post}} = x_{2m} \quad \sigma_{\text{post}}^2 = (\sigma_v t)^2 + \sigma_{s2}^2 \quad (\text{A48})$$

### MULTI-TAP PERCEPTION

So far, we have considered trajectories composed of just two taps. An interesting question arises in modeling the perception of multi-tap stimuli: is the observer's generative model (a) a direct extension of the one we have considered here, such that a zero-mean low-speed prior applies independently to each pair of consecutive taps, or (b) does the observer expect velocity to be consistent across the multi-tap trajectory, such that the prior applied to each tap pair might be a Gaussian centered on the velocity of the preceding pair (a zero-mean low-acceleration prior)?

Considering trajectories with an arbitrary number of taps,  $n$ , and permitting inhomogeneous spatial acuity, possibilities (a) and (b) result in the following generalizations of Eq. A18:

(a)

$$p(\{x_i\} | \{x_{im}\}, \{t_i\}) \propto \exp\left(-\left(\sum_{i=1}^n \frac{(x_{im} - x_i)^2}{2\sigma_{si}^2} + \sum_{i=1}^{n-1} \frac{(x_{i+1} - x_i)^2}{2(\sigma_v t_i)^2}\right)\right) \quad (\text{A49})$$

(b)

$$p(\{x_i\} | \{x_{im}\}, \{t_i\}) \propto \exp\left(-\left(\sum_{i=1}^n \frac{(x_{im} - x_i)^2}{2\sigma_{si}^2} + \frac{(x_2 - x_1)^2}{2(\sigma_v t_1)^2} + \sum_{i=2}^{n-1} \frac{\left(\frac{x_{i+1} - x_i}{t_i} - \frac{x_i - x_{i-1}}{t_{i-1}}\right)^2}{2\sigma_v^2}\right)\right) \quad (\text{A50})$$

Here  $\{x_i\}$  refers to the set of tap positions,  $x_1, x_2, \dots, x_n$ ;  $\{x_{im}\}$  to the corresponding set of measurements;  $\{t_i\}$  to the set of times elapsed between each tap  $i$  and tap  $i + 1$ ; and  $\sigma_{si}$  to the spatial uncertainty associated with tap  $i$ .

The observer's percept  $\{x_i^*\}$  in case (a) or (b) can be found by taking partial derivatives of Eq. A49 or Eq. A50 with respect to each of the  $\{x_i\}$ , setting these to zero, and solving the simultaneous equations. We used this method to find the percepts depicted in **Figures 10** and **11** [case (a)] and **Figure 12** [case (b)].



Alternatively, the identical percept can be found through Kalman smoothing (Haykin, 2001), a recursive extension of the predictive-postdictive formulation described above. The Kalman smoother consists of an iterative forward (predictive) pass through the stimulus sequence, followed by a backward (postdictive) pass. For model (a), the algorithm for the forward pass (the Kalman *filter*) is:

$$\begin{aligned} K_i &= \frac{\sigma_{i-1|i-1}^2 + (\sigma_v t)^2}{\sigma_{i-1|i-1}^2 + (\sigma_v t)^2 + \sigma_s^2} \\ \hat{x}_{i|i} &= \hat{x}_{i-1|i-1} + K_i (x_{im} - \hat{x}_{i-1|i-1}) \\ \sigma_{i|i}^2 &= (1 - K_i) (\sigma_{i-1|i-1}^2 + (\sigma_v t)^2) \end{aligned} \quad (\text{A51})$$

Here,  $K_i$  is the *Kalman gain* at time  $i$ ; the notation  $\hat{x}_{i|j}$  refers to the estimated position of tap  $i$  based on all taps up to and including tap  $j$ ; and  $\sigma_{i|j}^2$  is the variance of that estimate. The filter is initialized at the first tap, with  $\hat{x}_{1|1} = x_{1m}$ ,  $\sigma_{1|1}^2 = \sigma_s^2$ , and runs forward until tap  $n$  is reached. The Rauch-Tung-Striebel algorithm for the subsequent backward pass is:

$$\begin{aligned} C_i &= \frac{\sigma_{i|i}^2}{\sigma_{i|i}^2 + (\sigma_v t)^2} \\ \hat{x}_{i|n} &= \hat{x}_{i|i} + C_i (\hat{x}_{i+1|n} - \hat{x}_{i|i}) \\ \sigma_{i|n}^2 &= \sigma_{i|i}^2 + C_i^2 (\sigma_{i+1|n}^2 - \sigma_{i|i}^2 - (\sigma_v t)^2) \end{aligned} \quad (\text{A52})$$

We verified that Eqs A51 and A52 yielded the same percepts plotted in **Figures 10** and **11**.

## EXTENSIONS

Although skin is a two-dimensional surface, we have so far considered only a single position axis,  $x$ , along which stimuli occur. In essence, we have assumed that the orthogonal,  $y$  coordinate, of the taps is a known constant. We have also assumed that the time,  $t$ , is known. Each of these restrictions can be removed.

### Two-dimensional movement

A more realistic generative model would allow stimuli to move in any direction along a two-dimensional skin surface. To accomplish this, we can adopt an  $(x, y)$  Cartesian coordinate system in which the orthogonal components of the velocity vector are independently specified by low-speed priors:

$$\begin{aligned} p(v_x) &= \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{v_x^2}{2\sigma_v^2}\right) = \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{((x_2 - x_1)/t)^2}{2\sigma_v^2}\right) \\ p(v_y) &= \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{v_y^2}{2\sigma_v^2}\right) = \frac{1}{\sqrt{2\pi}\sigma_v} \exp\left(-\frac{((y_2 - y_1)/t)^2}{2\sigma_v^2}\right) \end{aligned} \quad (\text{A53})$$

The tap 1 and 2 likelihood functions generalize to:

$$\begin{aligned} p(x_{1m}, y_{1m} | x_1, y_1) &= \frac{1}{\sqrt{2\pi}\sigma_{s1}} \exp\left(-\frac{(x_{1m} - x_1)^2 + (y_{1m} - y_1)^2}{2\sigma_{s1}^2}\right) \\ p(x_{2m}, y_{2m} | x_2, y_2) &= \frac{1}{\sqrt{2\pi}\sigma_{s2}} \exp\left(-\frac{(x_{2m} - x_2)^2 + (y_{2m} - y_2)^2}{2\sigma_{s2}^2}\right) \end{aligned} \quad (\text{A54})$$

The posterior over trajectories then takes the form:

$$\begin{aligned} p(x_1, y_1, x_2, y_2 | x_{1m}, y_{1m}, x_{2m}, y_{2m}, t) \\ \propto \exp\left(-\left(\frac{(x_{1m} - x_1)^2 + (y_{1m} - y_1)^2}{2\sigma_{s1}^2} + \frac{(x_{2m} - x_2)^2 + (y_{2m} - y_2)^2}{2\sigma_{s2}^2} + \frac{((x_2 - x_1)^2 + (y_2 - y_1)^2)^2}{2(\sigma_v t)^2}\right)\right) \end{aligned} \quad (\text{A55})$$

It is straightforward to show that the length contraction formula resulting from Eq. A55 is identical to Eq. A20. Indeed, if we define the  $x$ -axis as the axis along which the tap measurements lie, then marginalization of Eq. A55 over  $y_1$  and  $y_2$  recovers the posterior density Eq. A18.

### Temporal uncertainty

Our model has assumed that the time between stimuli,  $t$ , is perceived veridically. This assumption can be removed. Goldreich (2007) showed that the Bayesian observer with temporal uncertainty tends to overestimate  $t$  in addition to underestimating  $l$ . Thus, the Bayesian observer can model time dilation as well as length contraction illusions.

### FITTING TO HUMAN PERCEPTUAL DATA

We found the value of tau that minimized the mean-squared error (MSE) between human and model performance. This was done separately for the perceptual data from Marks et al. (1982), Lechelt and Borchert (1977), and Kilgard and Merzenich (1995), shown in Figures 1A–C, and for the data from Helson and King (1931), shown in Figure 10.

The data of Helson and King (1931) required some processing prior to the fitting procedure. We fit the data reported in Tables 2–6 of Helson and King (1931). In those experiments, on each trial the participant reported whether the second spatial interval was perceived to be shorter than, equal to, or longer than the first interval (which was fixed at 3 cm). To fit these data, we first transformed them into an equivalent two-alternative forced-choice format by distributing each participant's "equal" responses evenly to the "shorter" and "longer" response categories. We then fit each participant's transformed data (proportion " $l_2$  is longer" responses) at each  $t_2$  setting with a Weibull psychometric function:

$$\Psi_{a,b,\gamma,\delta}(l_2) = (1 - \delta) \left[ \gamma + (1 - \gamma) \left( 1 - 2^{-\left(\frac{l_2 - 3\text{cm}}{a}\right)^b} \right) \right] + \frac{\delta}{2}$$

Here  $\delta$  is a lapse rate,  $\gamma$  is the probability that the concentrating participant would answer " $l_2$  is longer" when in fact  $l_2 = l_1$  (i.e., 3 cm),  $a$  is a position parameter, and  $b$  is a slope parameter. We found the maximum likelihood parameter settings, and from them read off the point of subjective equality (PSE:  $l_2$  that the participant judged longer than  $l_1$  with 50% probability). We fit the Bayesian observer's tau to minimize the MSE between its performance and the average PSE of the six human participants across the five  $t_2$  values tested by Helson and King (1931). Before doing these fits, we discarded the data from one of the six participants on one of the five  $t_2$  points: "Observer B" of Helson and King (1931) did not have a valid PSE at  $t_2 = 0.25$  s because that participant's transformed " $l_2$  is longer" response proportion was greater than 50% at all  $l_2$  values.