Cognitive Processing (2019) 20:447–457
https://doi.org/10.1007/s10339-019-00928-9

RESEARCH ARTICLE

The Bayesian causal inference model benefits from an informed prior to predict proprioceptive drift in the rubber foot illusion

Tim Schürmann[1] · Joachim Vogt[1] · Oliver Christ[2] · Philipp Beckerle[3,4]

Received: 9 November 2018 / Accepted: 12 August 2019 / Published online: 21 August 2019

Abstract
Bayesian cognitive modeling has become a prominent tool for the cognitive sciences aiming at a deeper understanding of the human mind and applications in cognitive systems, e.g., humanoid or wearable robotics. Such approaches can capture human behavior adequately with a focus on the crossmodal processing of sensory information. The rubber foot illusion is a paradigm in which such integration is relevant. After experimental stimulation, many participants perceive their real limb closer to an artificial replicate than it actually is. A measurable effect of this recalibration on localization is called the proprioceptive drift. We investigate whether the Bayesian causal inference model can estimate the proprioceptive drift observed in empirical studies. Moreover, we juxtapose two models employing informed prior distributions on limb location against an existing model assuming uniform prior distribution. The model involving empirically informed prior information yields better predictions of the proprioceptive drift regarding the rubber foot illusion when evaluated with separate experimental data. Contrary, the uniform model produces implausibly narrow position estimates that seem due to the precision ratio between the contributing sensory channels. We conclude that an informed prior on limb localization is a plausible and necessary modification to the Bayesian causal inference model when applied to limb illusions. Future research could overcome the remaining discrepancy between model predictions and empirical observation by investigating the changes in sensory precision as a function of distance between the eyes and respective limbs.

Keywords Rubber foot illusion · Bayesian inference · Crossmodal integration · Cognitive modeling

Introduction

In the last decade, Bayesian modeling of cognitive processes has gained momentum in the cognitive sciences (Hahn 2014). It provides a computational method to handle uncertainty when researching human decision-making and crossmodal integration, leading some researchers to the conclusion that one can even consider the brain to be a Bayesian inference machine (Dayan et al. 1995; Doya 2011).

Several psychological processes like vision (Weiss et al. 2002), language acquisition (Xu and Tenenbaum 2007), and crossmodal sensory integration (Deneve and Pouget 2004; Körding et al. 2007; Orbán and Wolpert 2011) have been hypothesized to follow Bayesian principles. In crossmodal integration, it is promising to consider the different noise levels of sensory modalities as uncertain input of a Bayesian cognitive model. Consider studying the precision of the sensorimotor system with regards to a person moving their arm to pick up a glass of water: they might include uncertain sensory information of the exact position of the

Handling editor: Anna Belardinelli (University of Tübingen).

Reviewers: Karl Friston (University College London) and a reviewer who prefers to remain anonymous.

✉ Tim Schürmann
schuermann@psychologie.tu-darmstadt.de

1 Work and Engineering Psychology, Institut für Psychologie, Technische Universität Darmstadt, Alexanderstr. 10, 64283 Darmstadt, Germany

2 Institute Humans in Complex Systems, School of Psychology, University of Applied Arts and Sciences Northwestern Switzerland, Olten, Switzerland

3 Elastic Lightweight Robotics Group, Robotics Research Institute, Technische Universität Dortmund, Dortmund, Germany

4 Institute for Mechatronic Systems in Mechanical Engineering, Technische Universität Darmstadt, Darmstadt, Germany

glass, gathered by the visual system, and the position of their arm conveyed by proprioception. They also have uncertain knowledge about the size and weight of the glass. The force the person chooses to pick up the glass might now be the best estimate resulting from an inference process in which they use their prior knowledge about the usual weight of a glass full of water with the likelihood of the glass being that heavy (given the amount of water in it, the size of the glass, etc). Bayesian inference describes the way to optimally combine prior knowledge and current sensory information in order to reduce the uncertainty involved in the decision-making process and arrive at a posterior estimate. There is a lively debate in the cognitive sciences on whether such ideal observer models (Daunizeau et al. 2010) adequately capture human behavior in a given task and how much explanatory value is provided if they do. Yet, the paradigm provides some face validity regarding why we might have a harder time estimating the appropriate force to use on a non-transparent box of milk than we do on a transparent glass of water: we might be more uncertain about the amount of milk inside the box since we cannot estimate it visually. Moreover, such formalized models of human cognitive behavior are of high interest in robotics research to endow robots with more human-like behaviors and capabilities (Schürmann et al. 2019).

The rubber hand illusion (RHI) describes the illusion of owning and describing an artificial limb as part of one's own body schema. It has been extensively described and studied in different sensorimotor, contextual and stimulation settings and is also considered promising for assistive and rehabilitation robotics (Beckerle et al. 2017; Botvinick and Cohen 1998; Christ and Reiner 2014; Moseley et al. 2012). Ehrsson et al. (2008) were able to show that the RHI may enable artificial prostheses being integrated into the body schema of patients suffering from upper limb amputation. Moreover, human-in-the-loop experiments to explore rubber limb illusions for the upper and lower limbs have been suggested (Beckerle et al. 2016; Schürmann et al. 2015). Various factors such as anatomically plausible positioning or the color of the artificial limb influence bodily illusions (Tsakiris and Haggard 2005) and might be computationally predictable by Bayesian inference. If an artificial hand in front of the participant showed a different skin texture than their usual skin texture, they may consider it less likely for said hand to belong to them. Combining information about the artificial hand and the experimental stimulation, one can apply the Bayesian causal inference paradigm (Berniker and Kording 2011) to infer the probability of a common cause for the sensations experienced during bodily illusions. For example, even if the tactile stimulation of the real limb matches the sensations, the participant sees being applied to the artificial one, deviating color of the artificial limb might interfere with the illusion. Similar inference processes appear plausible for other parameters of the RHI like distance between artificial and real hand or delay between visual and tactile stimulation during the experiment (Shimada et al. 2009). In fact, Samad et al. (2015) presented a Bayesian causal inference model for the RHI. It computes the likelihood of an artificial body part belonging to oneself, given the distance from the real body part and visuotactile delay. The goal of Samad et al. (2015) was to compute participants' self-reported limb positions posterior to stimulation, which is known as proprioceptive drift in the RHI literature and a parameter indicating a successful illusion. While their results are qualitatively consistent with empirical findings, proprioceptive drift is overestimated. The reported mean proprioceptive drift of roughly 17 cm is rather far from what is usually reported in the RHI literature (Christ and Reiner 2014), which is also mentioned by Samad et al. (2015).

While we assume that their computational approach works for both the RFI and the RHI qualitatively, we will outline adjustments to the model that we hypothesize will improve quantitative fit. This article draws inspiration from Samad et al. (2015) with regards to applying the paradigm of Bayesian causal inference from the RHI to the RFI. While Flögel et al. (2014) and others have already applied the idea of the more commonly studied RHI to the feet empirically, neither version of the illusion had been analyzed using a computational model prior to Samad et al. (2015). They point out that a RHI could reliably be reproduced by their model through the combination of prior knowledge about a limb's location and visual, tactile, and proprioceptive likelihoods, each represented with their individual levels of sensory uncertainty. The model correctly predicted that the illusion can occur without tactile stimulation present and that the probability of occurrence is enhanced through simultaneous brushing. Their model produces weighted estimates of participants' hand location, using both the combination of visual and proprioceptive likelihood, and each likelihood by itself. The sensory information was weighted by how likely it seemed that it was generated by a common cause, or two separate causes, respectively. For example, if a visuotactile stimulation was occurring asynchronously, with temporally different applications of the stimulation on the rubber hand and the participant's hand, a higher likelihood for separate causes of the sensation would be inferred. This can be explained similarly to how merely placing a rubber hand next to a hidden real hand can induce the RHI: if their distance to one another was not large enough, sensorimotor uncertainty would cause a proprioceptive drift to occur. If both rubber and real hand were placed farther apart from each other, it would seem clearer to the participant that their sensory information could not be generated by the same cause, i.e., seeing one's hand where one felt it.

While we hypothesize that the general model of Samad et al. (2015) should account for both the RHI and the RFI and is therefore transferrable, we argue that proprioceptive drift overestimation might depend on the implementation of prior knowledge. By employing a uniform distribution, Samad et al. (2015) assume that every possible limb position in peri-personal space is equally plausible from a participant's perspective. Uniform priors are common practice in Bayesian data analysis (Kruschke et al. 2012), but as a quantification of prior knowledge in cognitive modeling they need to be assessed for plausibility from the perspective of the cognitive system. It is likely that humans do not keep a uniform distribution of where their limbs are in space, but rather update their localizations continuously while moving. Thus, we hypothesize that an informed prior distribution, presuming earlier updating processes up until the starting point of an experiment, can place the most plausibility of an inferred limb location near its actual location in peri-personal space. This article investigates whether a contextually appropriate adjustment of the prior distribution leads to a more adequate prediction of proprioceptive drift through the Bayesian causal inference model.

- Likelihood of perceived visual information with respect to localization $p(\chi_v|X)$
- Likelihood of perceived proprioceptive information $p(\chi_p|X)$
- Likelihood of perceived visual information with respect to tactile feedback $p(\tau_v|T)$
- Likelihood of perceived tactile information with respect to tactile feedback $p(\tau_t|T)$.

X_v and X_p denote the physical spatial positions (in cm) of artificial and real limb, respectively, and T_v and T_t serve to represent potential temporal deviations between the sensory channels (in ms). The four corresponding likelihoods are conditioned on the actual spatial and temporal values of X and T, respectively. Following Bayes' Rule and adopting the equation and denotations of Samad et al. (2015), Eq. 1 generates an estimate of the posterior probability of a common cause with the given sensory input:

$$p(C=1|\chi_v,\chi_p,\tau_v,\tau_t)=\frac{p(\chi_v,\chi_p,\tau_v,\tau_t|C=1)p(C=1)}{p(\chi_v,\chi_p,\tau_v,\tau_t|C=1)p(C=1)+p(\chi_v,\chi_p,\tau_v,\tau_t|C=2)(1-p(C=1))} \tag{1}$$

Methods

In this section, we investigate the transferability of parameters applied by Samad et al. (2015) from the hands to the feet while considering the higher distance between the eyes and feet compared to the hands. We additionally describe how two separate empirical data sets are employed to inform and compare to our models.

Model parameters

Bayesian cognitive models, e.g., the applied causal inference model, represent sensory information through probability distributions. Standard deviations express uncertainty in the perceived information: the larger the standard deviation, the larger the associated uncertainty for a given value. While some parameter values could be adopted from Samad et al. (2015), it was necessary to account for the difference between the RFI and RHI paradigms in others. The following parameters were included in the model, and their distributions are considered to be Gaussian (parametrized by mean and standard deviation) except for the binary variable C.

- Prior probability of common cause/separate causes $p(C=1)/p(C=2)$
- Prior localization distribution $p(X)$
- Prior visuotactile distribution $p(T)$

The likelihood term for a common cause in Eq. 1, marginalized with respect to the physical values of X and T, is calculated through Eq. 2. The equivalent for separate causes is shown in Eq. 3.

$$p(\chi_v,\chi_p,\tau_v,\tau_t|C=1)=\iint p(\chi_v,\chi_p,\tau_v,\tau_t|X,T)p(X,T)\mathrm{d}X\mathrm{d}T \tag{2}$$

$$\begin{aligned} &p(\chi_v,\chi_p,\tau_v,\tau_t|C=2)\\ &=\iint p(\chi_v,\tau_v|X_v,T_v)p\left(X_v,T_v\right)\mathrm{d}X_v\mathrm{d}T_v\\ &*\iint p(\chi_p,\tau_t|X_p,T_t)p\left(X_p,T_t\right)\mathrm{d}X_p\mathrm{d}T_t \end{aligned} \tag{3}$$

The variability of the visual system has been found to translate to approximately 0.36 degrees (van Beers et al. 1998), or 1 mm for the distance between the participant's eyes and hands. Considering the larger distance between eyes and feet, a higher uncertainty can be expected in case of the RFI. Standard deviations of the temporal estimation for visual and tactile information were both assumed to be 20 ms according to Hirsh and Sherrick (1961). Studies that considered proprioceptive mapping of the hands found its standard deviation to be 15 mm (Jones et al. 2010). Samad et al. (2015) used the parameter values described above for likelihood distributions of incoming sensory information ($p(\chi_v)$; $p(\chi_p)$; $p(\tau_v)$; $p(\tau_t)$), combining them with uniform prior distributions across a bounded area of peri-personal space. This implies equal plausibility for every possible hand location and visuotactile delay.

The likelihood of perceived visual information with respect to localization $p(\chi_v)$ needed to be modified compared to Samad et al. (2015) because the distance between the eyes and the limb has increased. Instead of 0.2°–0.6° observed by van Beers et al. (1998), a degree of deviation of at least 0.6 degrees is to be expected. Since there is no experiment comparable to van Beers et al. (1998) regarding the lower limbs to the best of our knowledge, we assume a deviation of 0.6 degrees. Considering an average distance of 112.5 cm, this results in a standard deviation of 1.2 cm visual around a mean of 30 cm, i.e., the real position of the rubber foot.

With regards to the parameter values for the proprioceptive likelihood distribution, we generalize from the RHI values applied by Samad et al. (2015). This means that the proprioceptive likelihood distribution $p(\chi_p)$ has a mean of 50 cm and a standard deviation of 1.5 cm. In a synchronous condition, both distributions of temporal information (visual likelihood $p(\tau_v)$ and tactile likelihood $p(\tau_t)$) share an arbitrary mean depending on the timing of the stimulus. Defining an interval of one second around the stimulus application, we parametrized both distributions with a mean of 500 ms and a standard deviation of 20 ms, again following Samad et al. (2015). Participants should have no information about when tactile stimulation occurs relative to the start of the experimental stimulation phase. The prior distribution on visuotactile stimulation $p(\tau_v)$ was therefore left uninformed with a mean of 500 ms and a standard deviation of 10,000 ms. The prior probability of a common cause was left uninformed at .50.

Following our hypothesis that an informed prior would decrease overestimation of proprioceptive drift, we present two options for informing the prior distribution:

- Conceptually informed model: Assigning the real limb location (50 cm) and the proprioceptive precision (1.5 cm) to the mean and standard deviation of the Gaussian prior distribution
- Empirically informed model: Sampling a value for the mean of each participant's Gaussian prior distribution from the empirical distribution of pre-stimulation measurements of an informing dataset, while its standard deviation is equal to the proprioceptive precision (1.5 cm).

We argue that a uniform prior is not appropriate in the situation because a human participant would know quite well where their limb was before stimulation started. The simplest implementation would be to assume that each participant represents their prior limb position at its physical location. The conceptually informed model, however, ignores inter-individual variance in pre-stimulation localization that can be observed empirically by investigating the RFI (Christ et al. 2013; Flögel et al. 2015). To consider such variance, the empirically informed model relies on an informing dataset (Christ et al. 2013) that was gathered using the same RFI experimental setup as the dataset employed for model comparison. The pre-stimulation localization data in the synchronous RFI condition of that article (Christ et al. 2013) are used in the empirically informed model to establish a distribution of individual means for the prior distribution of the causal inference model. This distribution is parametrized by a mean of 50.46 cm and a standard deviation of 4.68 cm.

The information integration process following the described parameter values is outlined considering one single participant in Table 1. It describes the average model outcomes following Samad et al.'s parametrization (2015) and the changes proposed in this article. Note that the behavior of both informed models is depicted in the same column because, given a prior of 50 cm, they infer the same probability of a common cause and provide the same posterior localization if focusing on an individual participant. The difference between the informed models lies in the variation of the individual priors between different participants. While the conceptually informed model assumes a prior mean of exactly 50 cm for all participants, the empirically informed model draws that mean from a distribution governed by the second, informing dataset. Consequently, varying the prior mean as specified in the empirically informed model should result in a variation of limb localization that is closer to empirical observations. Due to focusing on one hypothetical participant and a prior of 50 cm in Table 1, the empirically informed model shows the same behavior as

Table 1 Information integration process of one exemplary participant for the uniform and informed models, focusing on an individual participant with an assumed prior limb localization of 50 cm

	Uniform model	Informed models
Common cause probability ($C=1$)		
Pre	$p=.50$	$p=.50$
Post	$p=.95$	$p=.95$
Localization of the limb		
Pre	Uniform	Mean 50 cm (Std dev 1.5 cm)
Post	Mean 38 cm (Std dev 0.8 cm)	Mean 41 cm (Std dev 0.8 cm)

Comparison of common cause and localization inference in pre- and post-conditions the conceptually informed model. Their capabilities in estimating the empirically observed inter-individual variance are analyzed in the subsequent model comparison. Starting from equal probability of a common cause or separate causes of the perceived sensory information, the posterior probability of a common cause is updated after receiving sensory information. Because all models simulate a synchronous RFI condition, temporal distributions were omitted for brevity. While the inferred probability of a common cause is identical between the models, the posterior localization distributions show the average change in the position estimate caused by the addition of an informed prior distribution under the assumption of a common cause.

After the models estimated the posterior probability of a common cause, they computed a value for the most probable location of the limb in question, presented in Eq. 4:

$$\hat{X}_p = p(C=1|\chi_v, \chi_p, \tau_v, \tau_t)\hat{X}_{p,C=1} + \left(1 - p(C=1|\chi_v, \chi_p, \tau_v, \tau_t)\hat{X}_{p,C=2}\right. \tag{4}$$

The estimates for $\hat{X}_{v,C=1} = \hat{X}_{p,C=1}$ as well as $\hat{X}_{v,C=2}$ and $\hat{X}_{p,C=2}$ are computed by Eqs. 5, 6, and 7:

$$\hat{X}_{v,C=1} = \hat{X}_{p,C=1} = \frac{\frac{\chi_v}{\sigma_v^2} + \frac{\chi_p}{\sigma_p^2} + \frac{\mu_X}{\sigma_X^2}}{\frac{1}{\sigma_v^2} + \frac{1}{\sigma_p^2} + \frac{1}{\sigma_X^2}} \tag{5}$$

$$\hat{X}_{v,C=2} = \frac{\frac{\chi_v}{\sigma_v^2} + \frac{\mu_X}{\sigma_X^2}}{\frac{1}{\sigma_v^2} + \frac{1}{\sigma_X^2}} \tag{6}$$

and

$$\hat{X}_{p,C=2} = \frac{\frac{\chi_p}{\sigma_p^2} + \frac{\mu_X}{\sigma_X^2}}{\frac{1}{\sigma_p^2} + \frac{1}{\sigma_X^2}} \tag{7}$$

Model evaluation

We compare the RFI model variants described in the previous section with an informed prior distribution on localization to the original model proposed by Samad et al. (2015). Subsequently, the former will be referred to as the conceptually or empirically "informed models," while the latter will be referred to as the "uniform model." To compare the models, we use Bayes factors (Annis and Palmeri 2017), which represent the ratio of marginal likelihoods of the data given the models, as is shown in Eq. 8.

$$B_{iu} = \frac{p(D|M_i)}{p(D|M_u)} \tag{8}$$

The index of B_{iu} denotes a directed Bayes factor of an informed model relative to the uniform model, with the reciprocal describing the other direction of model comparison. An established convention describes a Bayes factor of above 3 to indicate moderate evidence in favor of a model, with a Bayes factor of above 10 indicating strong evidence toward it (Lee and Wagenmakers 2013). Typically, a model with an increasing number of free parameters is able to arbitrarily fit an increasing amount of data (Annis and Palmeri 2017). The Bayes factor accounts for this problem of overfitting by marginalizing a model's free parameters. One can consider the marginal likelihood of the data under a given model to represent its average prediction, weighted by the prior uncertainty about its parameter values. This process results in a model with broad priors placing less probability mass at any given prediction than a very constrained model. The very constrained model might in turn make more precise, yet inaccurate, predictions.

In this paper, the cognitive models under comparison provide distributions of predicted position estimates after the experimental stimulation. Using an empirical data set described in the next section, we will judge model performance by comparing the likelihood of that data set under each model's predictions and report the resulting Bayes factors.

Participants and descriptive data analysis

To compare model results with actual participant behavior, we reanalyzed participant data reported by Flögel et al. (2015). Here, participants were seated at an experimental setup showcased in Fig. 1. The experienced foot position was measured by stopping a sliding light that was moving horizontally above the participant's limb. The proprioceptive drift of each participant was calculated as the difference of the pre- and post-stimulation position measurements. Flögel et al.'s (2015) experiment was originally designed to apply the RHI procedure to the feet and verify the existence of proprioceptive drift and questionnaire effects in the RFI for synchronous and asynchronous conditions. The study showed that for the synchronous condition, illusion strength was comparable for both limbs.

Flögel et al.'s (2015) empirical dataset consisted of 31 participants (20 females, age $M_{\mathrm{ea}} = 24$, $\mathrm{SD}_{\mathrm{ea}} = 6.65$ years) who were randomly assigned to differing sequences of stimulation. The synchronous RFI stimulation condition reported a proprioceptive drift of $M_{\mathrm{ed}} = 3.22$ cm ($\mathrm{SD}_{\mathrm{ed}} = 5.0$ cm), which differed significantly from the asynchronous ($t(30) = 4.78$, $p < 0.001$) and control condition ($t(30) = 2.91$, $p = 0.007$) of the RFI. Empirical data and Gaussian summary distributions are shown in Fig. 2.

To provide prior knowledge to the empirically informed, we used pre-stimulation data from a separate experiment by

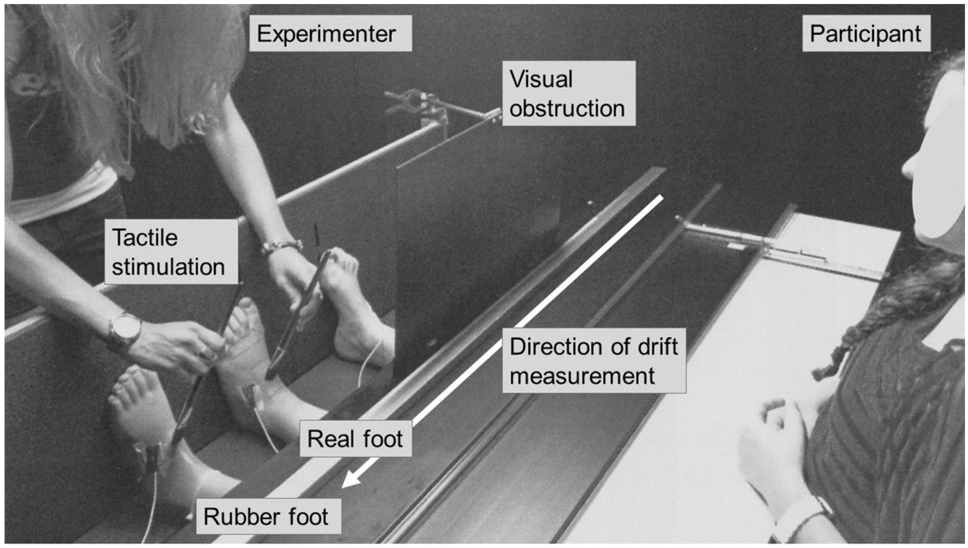

Fig. 1 Experimental setup of the rubber foot illusion as conducted by Christ et al. (2013)

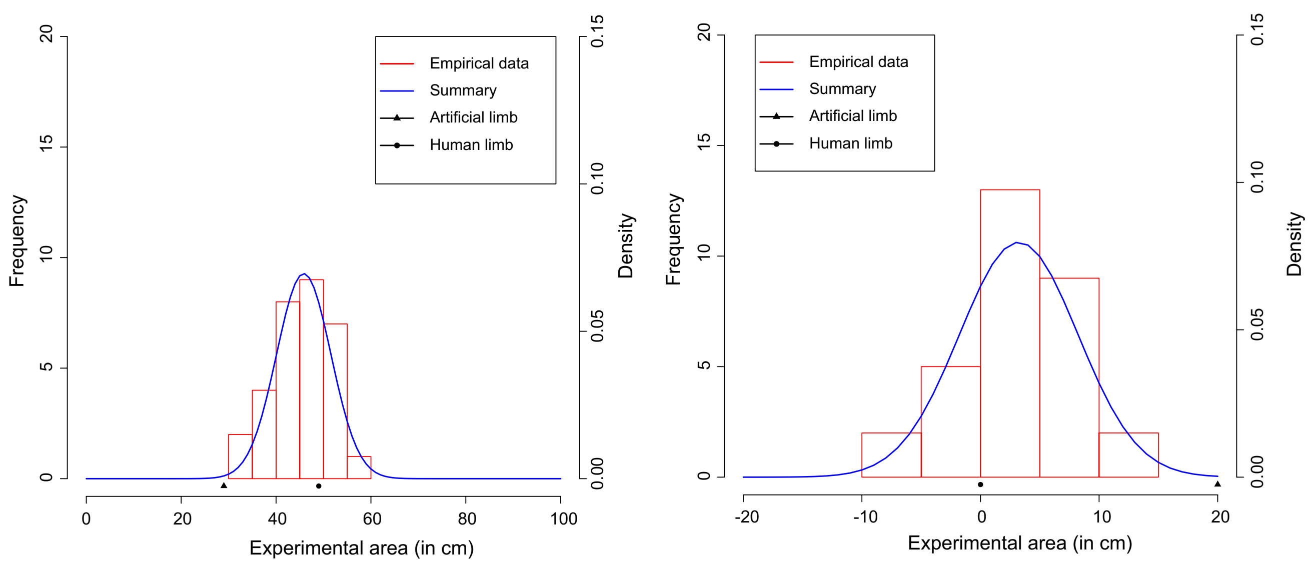

Fig. 2 Empirical data and Gaussian summary distributions for position estimates post-stimulation (left) and proprioceptive drift (right). For position estimates, the abscissa represents the horizontal location in front of a participant (position estimates closer to the artificial limb are plotted farther left). For proprioceptive drift, values closer to the artificial limb, which indicate a stronger drift, are plotted farther right

Christ et al. (2013), which includes data from 19 participants (11 females, age $M_{ia} = 24$, $SD_{ia} = 4.5$) with the above-mentioned pre-stimulation positions estimates ($M_{ip} = 50.46$ cm, $SD_{ip} = 4.68$ cm).

Results

This section describes the results of the model comparisons using visual presentation and Bayes factors. Figure 3 shows the position estimates provided by the uniform and both informed models and the empirically observed distribution (blue). The predicted position estimates of the empirically informed model (solid green) approximate the empirically observed amount and variance of post-stimulation localization. The position estimates of the uniform model (red), however, deviate distinctly regarding their mean and are noticeably narrow, leading to very low probability mass at other position estimates. This yields a very large Bayes factor between the empirically informed model and the uniform model. The conceptually informed model approximates the mean of the empirically observed post-stimulation localization similarly well as the empirically informed model, but does not describe the empirically observed variance. Comparing the empirically informed model to the conceptually informed one, a positively infinite Bayes factor is found,

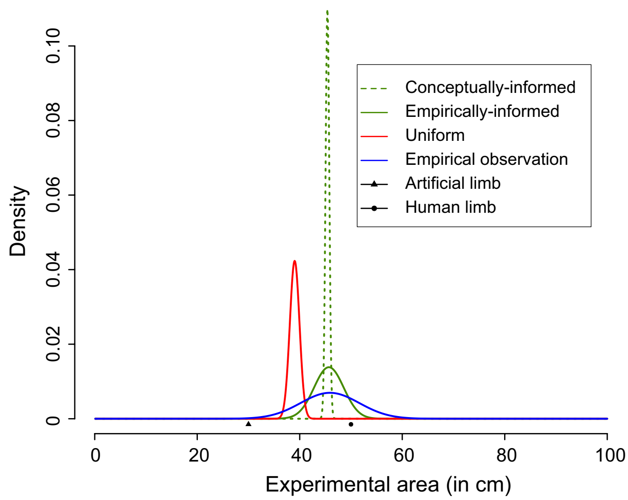

Fig. 3 Model predictions (uniform in red, conceptually informed in dashed green, empirically informed in solid green) and empirical distribution (blue) of position estimates after the experimental stimulation

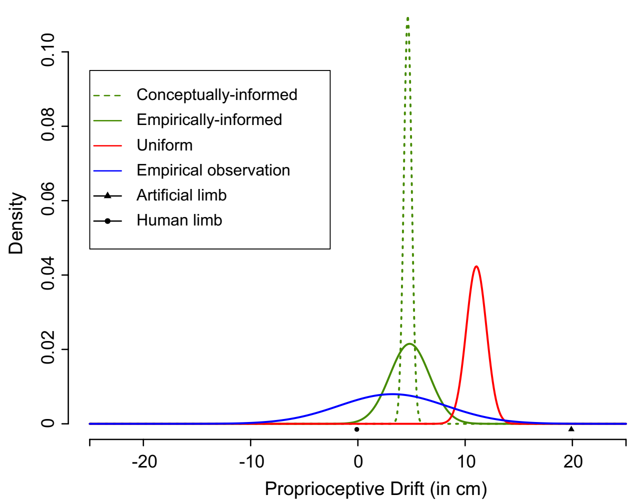

Fig. 4 Model predictions and empirical distribution of proprioceptive drift. Positive values indicate a drift from the position of the human limb toward the artificial limb

which underlines that the empirically informed model provides a better prediction of the empirically observed localization. The comparison between the conceptually informed model and the uniform model produces no Bayes factor because the product of each marginal likelihood is zero and the Bayes factor fraction is therefore undefined. Hence, the data do not provide support for either of those two models.

In addition to investigating position estimates after stimulation, Fig. 4 shows the proprioceptive drift estimates of the models as well as the empirically observed (blue) distribution of drifts. Note that the direction of the x-axis has been inverted here. While Fig. 3 represented a horizontal axis in front of a participant, with position estimates closer to the artificial limb plotted further left (away from the participant's body center), Fig. 4 represents proprioceptive drift values as absolute values increasing to the right. The Bayes factor from the empirically informed model (solid green) to the uniform model (red) is positively infinite, substantiating that the empirically informed model approximates the empirical values better. The comparison between the empirically informed and conceptually informed models is also positively infinite. As with the predicted position estimates before, the Bayes factor between the conceptually informed model and the uniform model is not defined because both receive approximately zero support from the data.

Discussion

The aim of the investigated models is to develop a better understanding of human body experience and to generate a quantitative estimation of the integration of multisensory percepts at the computational level (Marr 1982). This is a clear distinction to models that are applied in robotics for planning, control, and navigation (Siciliano and Khatib 2008). Yet, methods like Kalman filters (Roncone et al. 2016) and Bayesian filters (Lanillos et al. 2017) have also been applied to the problem of robotic self-perception, which might be combined with the proposed Bayesian models, e.g., for robot control (Roncone et al. 2016) or scene understanding (Lanillos et al. 2017). Considering Marr's levels, those methods from robotics could also serve as algorithmic realizations of our computational model for applications in robotics.

The presented models focus on the computational problem and the adequacy of their respective prior distributions. They consider the integration of sensory information to be a continuous process of Bayesian causal inference (Körding et al. 2007), making Bayesian computational modeling an appropriate choice. However, other approaches to implement Bayesian inference, e.g., predictive processing (Clark 2013) or active inference (Friston and Stephan 2007) can be assumed to apply as well.

Model performance

We have proposed a modification of the Bayesian causal inference model of bodily illusions used by Samad et al. (2015) aimed at serving two purposes. First, changes to the parameter value of visual precision were made for all models under comparison to reflect the different physical setup of the RFI compared to the RHI. Secondly, we argued for the inclusion of an informed prior distribution on localization before stimulation. We implemented this informed prior in two separate ways: (1) the conceptually informed model assumes the mean of each participant's prior distribution to correspond to the physical location of their limb on a horizontal plane in front of them, (2) we informed prior distributions of participants in the empirically informed model using a second dataset that was gathered with the same experimental setup as the evaluation dataset. Here, empirical variation in pre-stimulation localization was taken into account by sampling an individual participant's prior distribution mean from a distribution with parametrization equal to the second dataset. The resulting empirically informed model is therefore capable to account for inter-individual differences. These changes were made to reflect the intuition that participants should not be assumed to treat any position in front of them as equally plausible to be the position of their limb, but to have a fairly accurate estimate before the experimental stimulation interferes. An alternative and widely used approach in computational modeling (Sun 2008) would be to estimate participant-specific prior beliefs that rendered their empirical reports the most likely, i.e., optimizing prior beliefs instead of trying to compare sets of fixed priors. Such computational phenotyping is applied, for example, to the estimation of individual motor priors (Wolpe et al. 2014) or to the identification of prior belief structures in clinically relevant cohorts, e.g., different psychiatric cohorts (Schwartenbeck and Friston 2016). Our analysis can be considered a hypothesis-driven version of this search for empirically informed priors. Therefore, we tested the performance of our two suggested informed models against the uniform model suggested by Samad et al. (2015) with adjusted visual precision. Analysis via Bayes factors revealed that the empirically informed model was strongly favored compared to the other two, both concerning absolute position estimates and proprioceptive drift values.

Figure 4 shows that all models estimate proprioceptive drift values larger than the empirical observation, with the highest likelihood of the data under the empirically informed model. Only considering this measure might not seem unordinary to the reader, but the more important result comes from considering position estimates and proprioceptive drifts simultaneously. As can be seen in Fig. 3, both the uniform model and the conceptually informed model yielded position estimates that were all narrowly focused around specific positions on the horizontal plane in front of the participant. For the uniform model, predictions focused around approximately 39 cm, while the conceptually informed model showed similar behavior for a mean of approximately 45 cm. The small variation in these values is in accordance with several figures in Samad et al.'s work (Samad et al. 2015, Figs. 2a, 4).

Uniform model

The assumption that each possible limb position before stimulation is equally probable dominates any individual a priori position estimate of the uniform model. This result returns a position estimate that is solely impacted by the relative precision of visual and proprioceptive information channels. As we treated both sensory precision parameters as fixed and equal across the participants, the uniform model predicted approximately identical position estimates after stimulation for all participants, disregarding the location of its prior distribution mean completely. The proprioceptive drift results shown in Fig. 4 alone would not point out this characteristic of the uniform model, as it was outperformed by the empirically informed model but still produced plausible results at first glance. As Samad et al. (2015) have shown, the model behavior also holds up to qualitative investigation, for example showing a decrease in common cause probability at increasing distances. We showed, however, that the use of a uniform prior distribution on localization not only implements some unrealistic assumptions on plausible position estimates. It also results in the causal inference model returning the signal ratio between visual and proprioceptive channels as a prediction. This implausible model behavior did not receive support from the empirical data.

Conceptually informed model

The conceptually informed model shares the uniform model's issue of producing posterior position estimates that are narrowly focused around a specific value. While the mean of these estimates lies close to the empirically observed mean, the missing inter-individual variance leads to an even stronger compression of its posterior estimates compared to the uniform model's predictions. The model lacks the ability to display an adequate amount of inter-individual variance through its assumption of a fixed prior across individuals, but it improves upon the uniform model by allowing changes to the prior mean to be reflected in its posterior position estimates. If one were to fit the prior mean to a dataset instead of setting it to the physical location of a participant's limb, the conceptually informed model's predictions would approach the empirical data, while the uniform model could not account for such parameter changes. This property, however, does not detract from the fact that both models are equally unsupported by the data due to their narrow predictive distributions, leading to an incalculable Bayes factor.

Empirically informed model

The empirically informed model incorporates the possibility to reflect prior parameter changes in its posterior predictions, which turns them into more than a mere signal ratio between sensory channel precisions. Moreover, it abandons the conceptually informed model's assumption of a prior distribution fixed to a limb's physical location. Participants in RHI and RLI studies showcase a higher amount of variation, both pre- and post-stimulation (Christ et al. 2013; Flögel et al. 2015), than a prior fixed to a limb's physical location can account for. The empirically informed model creates more empirically plausible predictions by implementing a hierarchical dependency where the mean of a participant's prior is drawn from an inter-individual distribution governing the spread of the individual means, i.e., a hyperparameter or hyperprior (Farrell and Lewandowsky 2018; Gelman 2006; Kruschke 2015). While the implementation helps in achieving posterior position estimates that result in greater evidence provided by the data for the empirically informed model than the others, a graphical examination of Figs. 3 and 4 shows that its prediction mean is nearly identical to the conceptually informed model. The empirically informed model is more complex than its competitors and produces posterior localization estimates better supported by the empirical data, as shown through model comparison via Bayes factors.

Limitations

Using Bayes factors to determine the relative evidence provided by competing models only points toward the model that performs better in the set of models under consideration. This method of model comparison, however, does not reveal whether even the best model at hand should be considered a good model of the data. As can be seen by graphical examination of Figs. 3 and 4, the proposed empirically informed model does make predictions that are in better accordance with the empirical observation, but it still leaves room for improvement. An obvious starting point to improve model performance might be the adjustment of precision parameters from the hands to the feet. To our knowledge, there is no literature available on how visual and proprioceptive precision might change between the extremities, aside from assuming a linear increase. The previous literature on proprioceptive precision of the feet does not appear to be focused on azimuthal positioning relevant in the RFI setup (Robbins et al. 1997; Robbins et al. 1995).

Intuitively, one might assume a decreased proprioceptive precision in the feet compared to the hands due to the requirements of everyday tasks and the associated neural representation of both body parts. However, decreasing proprioceptive precision in the proposed informed model would increase proprioceptive drift and move position estimates toward the artificial limb, decreasing the fit to the empirical data. The only change related to precision one could implement to improve fit is to decrease visual precision more than we have done to compensate for increased distance between the eyes and feet. We are, however, unaware of prior work justifying more than a linear increase in visual precision.

Finally, the present article does not consider different experimental conditions of the RFI besides synchronous stimulation. While previous work has focused on testing and extending the conditions necessary for bodily illusions to occur (Christ and Reiner 2014; Crea et al. 2015; Lenggenhager et al. 2015), our scope was to investigate the assumptions of the causal inference model necessary to approximate empirical data. However, we expect our main finding of improved fit through an adequately informed prior on localization to translate to different experimental conditions.

Conclusion

Understanding bodily illusions is not only a very challenging topic of psychological research (Christ and Reiner 2014; Giummarra et al. 2008), but also has paramount engineering potential (Beckerle et al. 2017; Caspar et al. 2015; Schürmann et al. 2019). For instance, reliable models of human perception and cognition could enable new human–machine interaction strategies, e.g., interfaces based on online user modeling and adaptation. To improve the quantitative predictions of a Bayesian cognitive model of the RFI compared to observed data, we proposed an informed prior on localization. The comparison via Bayes factors revealed clear relative support for informing models with empirical prior data. This is confirmed by graphical examination, which outlines that an empirically informed model outperforms a uniform model due to the implausibly precise position estimates of the latter. A third model that is only conceptually informed with geometrical information about the setup showed similar predictions for the mean value, but is not able to explain the empirically observed variance. These results suggest that cognitive models of bodily illusions should be informed based on empirical data from prior experiments if available.

Future research ought to investigate the visual precision on targets at eye to foot distance. Further investigation is also needed when it comes to proprioceptive precision between hands and feet. Lastly, the presented models assume a 50% prior on the probability of a common cause. This probability would certainly be dependent on knowledge about the RHI/RFI paradigm, something the average participant might have acquired earlier since many bodily illusion experiments employ psychology students as participants for course credit. We have shown that including an adequately informed prior distribution on localization can improve the predictive accuracy of the Bayesian causal inference model in bodily illusions. Applying the above-mentioned model comparison technique of parameter estimation and fit maximization might enable more detailed individual phenotyping. Given further development, utilization of such models for online user adaptation in assistive robotic systems may increase their user acceptance and body scheme integration.

Acknowledgements This work received support from the German Research Foundation (DFG) through the project "Users' Body Experience and Human–Machine Interfaces in (Assistive) Robotics" (No. BE 5729/3&11). In addition, we would like to thank Mareike Flögel for providing their dataset as well as an average distance estimate between eyes and feet in an RFI experiment. Further, we would like to thank Frank Jäkel for advice on the manuscript.

Compliance with ethical standards

Conflict of interest The authors declare no conflict of interest.

References

Annis J, Palmeri TJ (2017) Bayesian statistical approaches to evaluating cognitive models: Bayesian statistical approaches. Cognit Sci. https://doi.org/10.1002/wcs.1458

Beckerle P, De Beir A, Schurmann T, Caspar EA (2016) Human body schema exploration: analyzing design requirements of robotic hand and leg illusions, pp 763–768. https://doi.org/10.1109/ROMAN.2016.7745205

Beckerle P, Salvietti G, Unal R, Prattichizzo D, Rossi S, Castellini C, Bianchi M (2017) A human–robot interaction perspective on assistive and rehabilitation robotics. Front Neurorobotics. https://doi.org/10.3389/fnbot.2017.00024

Berniker M, Kording K (2011) Bayesian approaches to sensory integration for motor control. Cognit Sci 2(4):419–428. https://doi.org/10.1002/wcs.125

Botvinick M, Cohen J (1998) Rubber hands "feel" touch that eyes see. Nature 391:756

Caspar EA, De Beir A, De Saldanha M, Da Gama PA, Yernaux F, Cleeremans A, Vanderborght B (2015) New frontiers in the rubber hand experiment: when a robotic hand becomes one's own. Behav Res Methods 47(3):744–755. https://doi.org/10.3758/s13428-014-0498-3

Christ O, Reiner M (2014) Perspectives and possible applications of the rubber hand and virtual hand illusion in non-invasive rehabilitation: technological improvements and their consequences. Neurosci Biobehav Rev 44:33–44. https://doi.org/10.1016/j.neubiorev.2014.02.013

Christ O, Elger A, Schneider K, Rapp A, Beckerle P (2013) Identification of haptic paths with different resolution and their effect on body scheme illusion in lower limbs. Presented at the European conference on technically assisted rehabilitation (TAR-2013), Berlin, Germany

Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behav Brain Sci 36(3):181–204

Crea S, D'Alonzo M, Vitiello N, Cipriani C (2015) The rubber foot illusion. J NeuroEng Rehabil. https://doi.org/10.1186/s12984-015-0069-6

Daunizeau J, den Ouden HEM, Pessiglione M, Kiebel SJ, Stephan KE, Friston KJ (2010) Observing the observer (I): meta-Bayesian models of learning and decision-making. PLoS ONE 5(12):e15554. https://doi.org/10.1371/journal.pone.0015554

Dayan P, Hinton GE, Neal RM, Zemel RS (1995) The helmholtz machine. Neural Comput 7(5):889–904

Deneve S, Pouget A (2004) Bayesian multisensory integration and cross-modal spatial links. J Physiol Paris 98(1–3):249–258. https://doi.org/10.1016/j.jphysparis.2004.03.011

Doya K (ed) (2011) Bayesian brain: probabilistic approaches to neural coding. MIT Press, Cambridge

Ehrsson HH, Rosen B, Stockselius A, Ragno C, Kohler P, Lundborg G (2008) Upper limb amputees can be induced to experience a rubber hand as their own. Brain 131(12):3443–3452. https://doi.org/10.1093/brain/awn297

Farrell S, Lewandowsky S (2018) Computational modeling of cognition and behavior, 1st edn. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9781316272503

Flögel M, Beckerle P, Christ O (2014) Rubber hand and rubber foot illusion: a comparison and perspective in rehabilitation. Clin Neurophysiol 125:S113. https://doi.org/10.1016/S1388-2457(14)50371-9

Flögel M, Kalveram K, Christ O, Vogt J (2015) Application of the rubber hand illusion paradigm: comparison between upper and lower limbs. Psychol Res. https://doi.org/10.1007/s00426-015-0650-4

Friston KJ, Stephan KE (2007) Free-energy and the brain. Synthese 159(3):417–458

Gelman A (2006) Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). Bayesian Anal 1(3):515–534

Giummarra MJ, Gibson SJ, Georgiou-Karistianis N, Bradshaw JL (2008) Mechanisms underlying embodiment, disembodiment and loss of embodiment. Neurosci Biobehav Rev 32(1):143–160. https://doi.org/10.1016/j.neubiorev.2007.07.001

Hahn U (2014) The Bayesian boom: Good thing or bad? Front Psychol. https://doi.org/10.3389/fpsyg.2014.00765

Hirsh IJ, Sherrick CE Jr (1961) Perceived order in different sense modalities. J Exp Psychol 62(5):423–432. https://doi.org/10.1037/h0045283

Jones S, Cressman EK, Henriques DYP (2010) Proprioceptive localization of the left and right hands. Exp Brain Res 204(3):373–383. https://doi.org/10.1007/s00221-009-2079-8

Körding KP, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, Shams L (2007) Causal inference in multisensory perception. PLoS ONE 2(9):e943. https://doi.org/10.1371/journal.pone.0000943

Kruschke JK (2015) Doing Bayesian data analysis: a tutorial with R, JAGS, and stan, 2nd edn. Academic Press, Boston

Kruschke JK, Aguinis H, Joo H (2012) The time has come: Bayesian methods for data analysis in the organizational sciences. Organ Res Methods 15(4):722–752. https://doi.org/10.1177/1094428112457829

Lanillos P, Dean-Leon E, Cheng G (2017) Yielding self-perception in robots through sensorimotor contingencies. IEEE Trans Cognit Dev Syst 9(2):100–112. https://doi.org/10.1109/TCDS.2016.2627820

Lee MD, Wagenmakers E-J (2013) Bayesian cognitive modeling: a practical course. Cambridge University Press, Cambridge

Lenggenhager B, Hilti L, Brugger P (2015) Disturbed body integrity and the "rubber foot illusion". Neuropsychology 29(2):205–211. https://doi.org/10.1037/neu0000143

Marr D (1982) Vision: a computational investigation into the human representation and processing of visual information. MIT Press, Cambridge

Moseley GL, Gallace A, Spence C (2012) Bodily illusions in health and disease: physiological and clinical perspectives and the concept of a cortical 'body matrix'. Neurosci Biobehav Rev 36(1):34–46. https://doi.org/10.1016/j.neubiorev.2011.03.013

Orbán G, Wolpert DM (2011) Representations of uncertainty in sensorimotor control. Curr Opin Neurobiol 21(4):629–635. https://doi.org/10.1016/j.conb.2011.05.026

Robbins S, Waked E, Mcclaran J (1995) Proprioception and stability: foot position awareness as a function of age and footware. Age Ageing 24(1):67–72. https://doi.org/10.1093/ageing/24.1.67

Robbins S, Waked E, Allard P, McClaran J, Krouglicof N (1997) Foot position awareness in younger and older men: the influence of footwear sole properties. J Am Geriatr Soc 45(1):61–66. https://doi.org/10.1111/j.1532-5415.1997.tb00979.x

Roncone A, Hoffmann M, Pattacini U, Fadiga L, Metta G (2016) Peripersonal space and margin of safety around the body: learning visuo-tactile associations in a humanoid robot with artificial skin. PLoS ONE 11(10):e0163713. https://doi.org/10.1371/journal.pone.0163713

Samad M, Chung AJ, Shams L (2015) Perception of body ownership is driven by Bayesian sensory inference. PLoS ONE 10(2):e0117178. https://doi.org/10.1371/journal.pone.0117178

Schürmann T, Overath P, Christ O, Vogt J, Beckerle P (2015) Exploration of lower limb body schema integration with respect to body-proximal robotics, pp 61–65. https://doi.org/10.1109/RTSI.2015.7325072

Schürmann T, Mohler BJ, Peters J, Beckerle P (2019) How cognitive models of human body experience might push robotics. Front Neurorobot 13:14

Schwartenbeck P, Friston K (2016) Computational phenotyping in psychiatry: a worked example. ENeuro. https://doi.org/10.1523/ENEURO.0049-16.2016

Shimada S, Fukuda K, Hiraki K (2009) Rubber hand illusion under delayed visual feedback. PLoS ONE 4(7):e6185. https://doi.org/10.1371/journal.pone.0006185

Siciliano B, Khatib O (eds) (2008) Springer handbook of robotics: with… 84 tables. Springer, Berlin

Sun R (ed) (2008) The Cambridge handbook of computational psychology. Cambridge University Press, New York

Tsakiris M, Haggard P (2005) The rubber hand illusion revisited: visuotactile integration and self-attribution. J Exp Psychol Hum Percept Perform 31(1):80–91. https://doi.org/10.1037/0096-1523.31.1.80

van Beers RJ, Sittig AC, Denier van der Gon JJ (1998) The precision of proprioceptive position sense. Exp Brain Res 122(4):367–377. https://doi.org/10.1007/s002210050525

Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. Nat Neurosci 5(6):598–604. https://doi.org/10.1038/nn858

Wolpe N, Wolpert DM, Rowe JB (2014) Seeing what you want to see: priors for one's own actions represent exaggerated expectations of success. Front Behav Neurosci 8:232

Xu F, Tenenbaum JB (2007) Word learning as Bayesian inference. Psychol Rev 114(2):245–272. https://doi.org/10.1037/0033-295X.114.2.245