Using the past to estimate sensory uncertainty

Ulrik Beierholm[1*], Tim Rohe[3,4*], Ambra Ferrari[2], Oliver Stegle[5,6,7], Uta Noppeney[2,3]

[1]Psychology Department, Durham University, Durham, United Kingdom.

[2]Centre for Computational Neuroscience and Cognitive Robotics, University of Birmingham, Birmingham, United Kingdom.

[3]Max Planck Institute for Biological Cybernetics, Tuebingen, Germany

[4]Department of Psychiatry and Psychotherapy, University of Tuebingen, Tuebingen, Germany

[5] Max Planck Institute for Intelligent Systems, Tübingen

[6] European Molecular Biology Laboratory, Genome Biology Unit, 69117 Heidelberg, Germany

[7] Division for Computational Genomics & Systems Genetics, German Cancer Research Center, 69120 Heidelberg, Germany

* These authors contributed equally to this work.

Correspondence should be addressed to ulrik.beierholm@durham.ac.uk

Abstract

To form the most reliable percept of the environment, the brain needs to represent sensory uncertainty. Current theories of perceptual inference assume that the brain computes sensory uncertainty instantaneously and independently for each stimulus.

In a series of psychophysics experiments human observers localized auditory signals that were presented in synchrony with spatially disparate visual signals. Critically, the visual noise changed dynamically over time with or without intermittent jumps. Our results show that observers integrate audiovisual inputs weighted by sensory reliability estimates that combine information from past and current signals as predicted by an optimal Bayesian learner or approximate strategies of exponential discounting

Our results challenge classical models of perceptual inference where sensory uncertainty estimates depend only on the current stimulus. They demonstrate that the brain capitalizes on the temporal dynamics of the external world and estimates sensory uncertainty by combining past experiences with new incoming sensory signals.

Perception has been described as a process of statistical inference based on noisy sensory inputs[1]. Key to this perceptual inference is the representation of sensory uncertainty. Most prominently, in multisensory perception the most reliable or 'Bayes-optimal' percept is obtained by integrating sensory signals weighted by their reliability (i.e., precision or inverse of variance) with less weight assigned to less reliable signals. Indeed, accumulating evidence suggests that human observers are close to optimal in many perceptual tasks (though see [1-3]) and weight signals according to their sensory reliabilities[4-8].

An unresolved question is how human observers compute and represent their sensory uncertainty. Current theories and experimental approaches generally assume that observers access sensory uncertainty near-instantaneously and independently across briefly (≤200 ms) presented stimuli [9,10]. At the neural level, theories of probabilistic population coding have suggested that sensory uncertainty may be represented instantaneously in the gain of the neuronal population response[11]. Yet, in our natural environment sensory noise often evolves at slow timescales. For instance, visual noise slowly varies when walking through a snow storm. Observers may capitalize on the temporal dynamics of the external world and use the past to inform current estimates of sensory uncertainty. In this alternative account, more reliable estimates of sensory uncertainty would be obtained by combining past estimates with current sensory inputs as predicted by Bayesian learning.

To arbitrate between these two critical hypotheses, we presented observers with audiovisual signals in synchrony but with a small spatial disparity in a sound localization task. Critically, the spatial standard deviation (STD) of the visual signal changed dynamically over time continuously (experiment 1-3) or discontinuously (i.e. with intermittent jumps; experiment 4). First, we investigated whether the influence of the visual signal location on observer's perceived sound location depended on the noise only of the current visual signal or also of past visual signals. Second, using computational modeling and Bayesian model

comparison, we formally assessed whether observers update their visual uncertainty estimates consistent with an i. instantaneous learner, ii. optimal Bayesian learner or iii. exponential learner.

Results

- Figure 1 to appear about here -

In a series of four experiments, we presented participants with audiovisual signals in a spatial localization task where visual noise changed dynamically over time at 5Hz (Fig. 1). On each trial, participants located the sound that was presented in synchrony with a change in color of the visual cloud of dots. The spatial disparity between the auditory signal and the mean of the visual cloud was set to ± 5° visual angle. This small audiovisual disparity enabled an influence of the visual signal location on the perceived sound location as a function of visual noise[4,12]. As a result, observers' visual uncertainty estimate could be quantified in terms of the relative weight of the auditory signal on the perceived sound location with a greater auditory weight indicating that observers estimated a greater visual uncertainty. Importantly, while the visual cloud's mean (together with the corresponding sound position) was independently resampled on each trial from five possible locations, the visual cloud was re-displayed at a higher temporal rate of 5 Hz with the cloud's horizontal standard deviation (STD) slowly varying over time. In three experiments, we used continuous sequences, where the visual cloud's STD changed periodically according to a sinusoid (n = 25; period = 30 s), a random walk (RW1; n = 33; period = 120 s) or a smoothed random walk (RW2; n = 19; period = 30 s; Fig. 2). In an additional fourth experiment, we inserted abrupt increases or decreases into a sinusoidal evolution of the visual cloud's STD (n =18, period = 30 s, Fig. 4).

- Figure 2 to appear about here -

We assigned physical visual noise (i.e., the cloud's STD) and sound localization responses to 20 (resp. 15 for experiment 4) temporally adjacent bins in each period of the three sequences. Next, we quantified the auditory and visual influence on the perceived auditory location for each time bin based on the regression model: $R_A = L_A * ß_A + L_V * ß_V + const * ß + e$ with R_A= Localization response; L_A or L_V = 'true' auditory or visual location; $ß_A$ or $ß_V$ =

auditory or visual weight. The relative auditory weight was computed as $w_A = ß_A / (ß_A + ß_V)$. As expected, the auditory weight approximately tracked the visual cloud's STD (Fig. 2). Yet, while the temporal evolution of the physical visual noise was designed to be symmetrical for each time period, we observed a temporal asymmetry for w_A in all three experiments. W_A was smaller for the 1st half of each period, when visual noise increased, than the 2nd half, when visual noise decreased over time (Fig. 3A). This impression was confirmed statistically in 2 (1st vs. flipped 2nd half) x 9 (time bins) repeated measures ANOVAs (Tab. 1) showing a significant main effect of 1st versus flipped 2nd period's half for the sinusoid ($F(1, 24) = 12.162$, $p = 0.002$, partial $\eta^2 = 0.336$) and the RW1 sequence ($F(1, 32) = 14.129$, $p < 0.001$, partial $\eta^2 = 0.306$). For the RW2 sequence, we observed a significant interaction ($F(4.6, 82.9) = 3.385$, $p = 0.010$, partial $\eta^2 = 0.158$), because the visual noise did not change monotonously within each period half. Instead, monotonic increases and decreases in visual noise alternated at nearly the double frequency as compared to RW1. The asymmetry in the auditory weights' time course suggested that the visual noise in the past influenced observers' current visual uncertainty estimate.

We further tested this in a regression model where the relative auditory weights w_A for each of the 20 time bins were predicted by the visual STD in the current time bin and the difference in STD between the current and the previous bin. Indeed, both the current visual STD ($p < 0.001$ for all three sequences; Sinusoid: $t(24)=15.767$, Cohen's $d=3.153$; RW1: $t(32)= 15.907$, Cohen's $d=2.769$; RW2: $t(18)=12.978$, Cohen's $d=2.9773$, two sided one-sample t test against zero) and the difference in STD between the current and the previous time bin (i.e. sinusoid $t(24)= -3.687$, $p = 0.001$, Cohen's $d=-0.738$; RW1 $t(32)= -2.593$, $p = 0.014$, Cohen's $d=-0.451$; RW2 $t(18)=-2.394$, $p = 0.028$, Cohen's $d=-0.549$) significantly predicted observers' relative auditory weights. Collectively, these results suggest that observers' visual uncertainty

estimates (as indexed by the relative auditory weights w_A) depend not only on the current sensory signal, but also on the recent history of the sensory noise.

Next, we aimed to characterize how human observers use information from the past to estimate current sensory uncertainty by comparing three computational models that differed in how visual uncertainty is learnt over time: Model 1, the instantaneous learner, estimates visual uncertainty independently for each trial as assumed by current standard models. Model 2, the optimal Bayesian learner, estimates visual uncertainty by updating the prior uncertainty estimate obtained from past visual signals with the uncertainty estimate from the current signal. Model 3, the exponential learner, estimates the visual uncertainty by exponentially discounting past uncertainty estimates. All three models account for the observer's uncertainty about whether auditory and visual signals were generated by common or independent sources by explicitly modeling the two potential causal structures[13] underlying the audiovisual signals (n.b. only the model component pertaining to the 'common cause' case is shown in Fig. 1B, for the full model see Supplementary Fig. S1). Models were fit individually to observers' data by sampling from the posterior over parameters for each observer (Tab. 2). We compared the three models using the Watanabe-Akaike information criterion (WAIC) as appropriate for evaluating model samples[14] (i.e., a low WAIC indicates a better model, with a difference greater than 10 considered very strong evidence for a model): across observers, the Bayesian learner ($WAIC_{Bayes} = 254158$) outperformed the exponential learner ($\Delta WAIC = +61$) and was substantially better than the instantaneous learner ($\Delta WAIC = +400$). Further, two-sided non-parametric permutation tests at the group level showed that the WAIC-values for the Bayesian learner were significantly greater across observers in comparison to the instantaneous learner ($p<0.001$) with a small non-significant trend relative to the exponential learner ($p\sim0.091$). As shown in Figure 3 both, the Bayesian and the exponential learner, accurately reproduced the temporal asymmetry for the auditory weights across all three experiments.

From the optimal Bayesian learner we inferred the observers' estimated rate of change in visual reliability (i.e. parameter κ. The Sinusoidal sequence was estimated to change at a faster pace (median κ = 7.4 across observers) than the RW1 sequence (median κ = 8.1), but slower than the RW2 sequence (median κ =6.7 for RW2 sequence) indicating that the Bayesian learner accurately inferred that visual reliability changed at different pace across the three sequences (see legend of Fig. 2). The learning rates for the exponential learner (Sin 0.23, RW1 0.33, RW2 0.25) implies a timescale of respectively 4.1, 5.4 and 4.3 seconds (calculated as the time for 95 percent decay).

- Figure 4 to appear about here -

To further disambiguate between the Bayesian and the exponential learner we designed a fourth experimental 'jump sequence' that introduced abrupt increases or decreases in visual noise at three positions into the sinusoidal sequence (Fig. 3A). As shown in Fig 3B (blue vs. green line) we would expect the exponential learner to adapt faster to abrupt decreases in visual uncertainty than the Bayesian learner that takes into account its uncertainty about its visual uncertainty estimate. Again we observed a temporal asymmetry for the auditory weights (Fig 3C). For all three 'jump positions' w_A was significantly smaller for the 1st half of each period, when visual noise increased, than the 2nd half, when visual noise decreased over time . The 3 (jump positions) x 2 (1st vs. flipped 2nd half) x 7 (time bins) repeated measures ANOVA showed a significant main effect of 1st versus flipped 2nd period's half ($F(1,17) = 24.824$, $p < 0.001$, partial $\eta^2 = 0.594$), while this factor was not involved in any higher-order interactions (Tab. 1) Bayesian model comparison showed that both the Bayesian learner ($WAIC_{Bayes} = 83796$) and the exponential learner ($WAIC_{Bayes} = 83798$) substantially outperformed the instantaneous learner ($\Delta WAIC = +95$, permutation test on the WAIC values: $p<0.001$). However, as for the continuous sequences the Bayesian learner did not provide a better explanation for observers'

responses than the exponential learner (ΔWAIC = +2, permutation test on the WAIC values: p=0.44, see Fig. 4D and Supplementary Fig. S2).

In order to further disambiguate the Bayesian and exponential learning models we performed an extra set of simulations. One of the defining characteristics of Bayesian models is that they incorporate the uncertainty of estimates, something which simpler more heuristic models do not. Hence, based on the sinusoidal and the 'jump'sequences, we created a version of the task where instead of 20 dots, only 5 dots were presented for each visual presentation, thus increasing the uncertainty about the visual variance. To minimize uncontrolled noise we insured that the mean and variance of the 5 dots corresponded to the experimentally controlled visual mean and variance.

For the two sequences we simulated the choices of 12 subjects using the Bayesian learner model (parameters: sigmaA=6 deg, kappa=15, Pcommon=0.7 and sigma0=12 degrees) and then fitted the Bayesian learner and exponential learner models to each simulated data set. While the fitted parameters for the Bayesian model were very close to the generating parameters (seq 1 fitted parameters: sigmaA=6.11 deg, kappa=17.5, Pcommon=0.72 and sigma0=12.4 degrees, seq 4 fitted parameters: sigmaA=6.08 deg, kappa=17.3 , Pcommon=0.71 and sigma0=12.2 degrees), thus providing a simple version of parameter recovery and model verification, the fits to the data were very close for the two models (Supplementary Fig. 4 A, B), with mean log likeihood difference (Bayes - Exp) =-1.82 for sequence 1 and -2.74 for sequence 4. Figures S4C and D also show that the inferred changes in visual standard deviation are very similar for the two models. This implies that even when testing using a sequence designed to separate the behaviour of the two models, it will be very difficult to dissociate the Bayesian and exponential learning models given the added randomness of subjects' behavior.

Discussion

Collectively, our results challenge classical models of perceptual inference where a perceptual interpretation is obtained using a likelihood that depends solely on the current sensory inputs[5]. These models implicitly assume that sensory uncertainty (i.e., likelihood variance) is instantaneously and independently accessed from the sensory signals on each trial based on initial calibration of the nervous system[15]. Most prominently, in the field of cue combination it is generally (but see[16,17]) assumed that sensory signals are weighted by their uncertainties that are estimated only from the current sensory signals[4-6]. By contrast, our results demonstrate that human observers integrate inputs weighted by uncertainties that are estimated jointly from past and current sensory signals. While we focus on the role of prior sensory inputs on uncertainty estimation in a cue combination context, we expect that they reveal fundamental principles of how the human brain computes and encodes sensory uncertainty.

Critically, we also show that for continuous and discontinuous changes in sensory uncertainty that are common for perception in our natural environment this optimal Bayesian inference can be approximated by more simple strategies of exponential discounting that could be performed by slowly updating the sensory weights without explicit representation of uncertainty[3,9]. Future experiments that increase observers' uncertainty about their visual uncertainty estimate may potentially be able to dissociate Bayesian from approximate exponential discounting. From the perspective of neural coding, our findings suggest that current theories of probabilistic population coding[11,18] may need to be extended to accommodate additional influences of past experiences on neural representations of sensory uncertainties. Alternatively, the brain may compute sensory uncertainty using strategies of temporal sampling[19].

In conclusion, to our knowledge, this is the first demonstration that human observers do not access the uncertainty instantaneously from the current sensory signals, but learn

sensory uncertainty over time by combining past experiences and current sensory inputs as predicted by an optimal Bayesian learner or approximate strategies of exponential discounting.

Methods

Participants

76 healthy volunteers participated in the study after giving written informed consent (40 female, mean age 25.3 years, range 18-52 years). All participants were naïve to the purpose of the study. All participants had normal or corrected-to normal vision and reported normal hearing. The study was approved by the human research review committee of the University of Tübingen.

Stimuli

The visual spatial stimulus was a Gaussian cloud of twenty bright grey dots (0.56° diameter, vertical standard deviation 1.5°, luminance 106 cd/m^2) presented on a dark grey background (luminance 62 cd/m^2, i.e. 71% contrast). The auditory spatial cue was a burst of white noise with a 5ms on/off ramp. To create a virtual auditory spatial cue, the noise was convolved with spatially specific head-related transfer functions (HRTFs). The HRTFs were pseudo-individualized by matching participants' head width, heights, depth and circumference to the anthropometry of subjects in the CIPIC database[20]. HRTFs from the available locations in the database were interpolated to the desired locations of the auditory cue.

Experimental design and procedure

In a spatial ventriloquist paradigm, participants were presented with audiovisual spatial signals. Participants indicated the location of the sound by pressing one of 5 spatially corresponding buttons. The visual signal was a cloud of 20 dots sampled from a Gaussian. The visual clouds were re-displayed with variable horizontal standard deviations (see below) at a rate of 5 Hz (Fig. 1A). The cloud's location mean was temporally independently resampled from five possible locations (-10°, -5°, 0°, 5°, 10°) at a stimulus onset asynchrony jittered between 1.4

and 2.8 s. In synchrony with the change in the cloud's location, the dots changed their colour and a concurrent sound was presented. The location of the sound was sampled from ± 5° visual angle with respect to the mean of the visual cloud. Observers' visual uncertainty estimate was quantified in terms of the relative weight the auditory signal on the perceived sound location. The change in the dot's colour and the emission of sound occurred in synchrony to enhance audiovisual binding.

Continuous sinusoidal and random walk sequences

Critically, to manipulate visual noise over time, the cloud's standard deviation changed according to i. a sinusoidal sequence, ii. a random walk sequence 1 or iii. a random walk sequence 2. In all sequences the horizontal standard deviation of the visual cloud spanned a range from 2-18°:

i. *Experiment1 - Sinusoidal sequence (Sinusoid):* A sinusoidal sequence was generated with a period of 30s. During the ~65 min of the experiment, each participant completed ~ 130 cycles of the sinusoidal sequence.
ii. *Experiment2 - Random walk sequence 1 (RW1):* First, we generated a random walk sequence of 60 s duration using a Markov chain with 76 discrete states and transition probabilities of stay (1/3), change to lower (1/3) or upper (1/3) adjacent states. To ensure that the random walk sequence segment starts and ends with the same value, this initial 60 s sequence segment was concatenated with its temporally reversed segment resulting in a RW sequence segment of 120 s duration. Each participant was presented with this 120s RW1 sequence approximately 32 times during the experiment.
iii. *Experiment3 - Random walk sequence 2 (RW2):* Likewise, we created a second random-walk sequence of 15 s duration using a Markov chain with only 38 possible states and transition probabilities similar to above. The 15 s sequence was concatenated with its

temporally reversed version resulting in a 30 s sequence. The smoothness of this sequence segment was increased by filtering it (without phase shift) with a moving average of 250 ms. Each participant was presented with this sequence segment ~130 times.

Generally, a session of a Sin, RW1 or RW2 sequence included 1676 trials. Because of experimental problems, four sessions included only 1128, 1143 or 1295 trials. Before the experimental trials, participants practiced the auditory localization task in 25 unimodal auditory trials, 25 audiovisual congruent trials with a single dot as visual spatial cue and 75 trials with stimuli as in the main experiment.

Experiment4 - Sinusoidal sequence with intermittent changes in visual noise

To dissociate the Bayesian learner from approximate exponential discounting, we designed a sinusoidal sequence (period = 30 s) with intermittent increases / decreases in visual variance (Fig. 4). As shown in figure 4A, we inserted increases by 8° in visual STD at three levels of visual STD: 7.2°, 8.6°, 9.6° STD. Conversely, we inserted decreases by 8° in visual STD at 15.3°, 16.7°, 17.7° STD. We factorially combined these 3 (increases) x 3 (decreases) such that each sinewave cycle included exactly one sudden increase and decrease in visual STD (i.e., 9 jump types). Otherwise, the experimental paradigm and stimuli were identical to the continuous sinusoidal sequence described above. During the ~80 min of this experiment, each participant completed ~ 154 cycles of the sinusoidal sequence including 16-18 cycles for each of the 9 jump types.

This jump sequence aimed to maximize differences in adaptation rate for the Bayesian and exponential learner by introducing rapid increases or decreases in visual noise. While the exponential learner will weight past and present uncertainties throughout the entire sequence based on the same exponential function, the Bayesian learner will take into account the uncertainty about the visual variance and therefore adapt its visual variance estimate more

slowly than the exponential learner for jumps from high to low visual uncertainty (see Figure 4B).

Subject numbers and inclusion criteria

29 of the 76 subjects participated in the sinusoidal and the RW1 sequence session. Eight additional subjects participated only in the RW1 sequence session. 18 independent subjects participated in the RW2 sequence session. One participant completed all three sequences. We excluded five participants of the Sin and RW1 sequence and two participants from the sinusoidal sequence with intermittent changes in visual uncertainty from further analysis, because their perceived sound location did not depend on the current visual reliability (i.e., inclusion criterion $p < 0.05$ in the linear regression; please note that this inclusion criterion is orthogonal to the question of whether participants' visual uncertainty estimate depends on visual signals prior to the current trial). In total, we analysed data from 25 participants for the sinusoidal, 33 participants for the RW1 and 19 participants for the RW2 sequence. Eighteen subjects participated in the sinusoidal sequence with intermittent changes in visual uncertainty.

Experimental setup

Audiovisual stimuli were presented using Psychtoolbox 3.09[21,22] (www.psychtoolbox.org) running under Matlab R2010b (MathWorks) on a Windows machine (Microsoft XP 2002 SP2). Auditory stimuli were presented at ~75 dB SPL using headphones (Sennheiser HD 555). As visual stimuli required a large field of view, they were presented on a 30" LCD display (Dell UltraSharp 3007WFP). Participants were seated at a desk in front of the screen in a darkened booth, resting their head on an adjustable chin rest. The viewing distance was 27.5 cm. This setup resulted in a visual field of approx. 100°. Participants responded via a standard QWERTY

keyboard. Participants used the buttons {i, 9, 0, -, =} with their right hand for localization responses.

Data analysis

Continuous sinusoidal and random walk sequences

For each period of the three continuous sinusoidal and random walk sequences, we sorted the physical visual uncertainty (i.e., the cloud's STD) and sound localization responses into 20 temporally adjacent bins. Next, we quantified the auditory and visual influence on the perceived auditory location for each time bin based on the regression model: $R_A = L_A * ß_A + L_V * ß_V + const * ß + e$ with R_A= Localization response; L_A = 'true' auditory location; L_V = 'true' visual location; $ß_A$ = auditory weight; $ß_V$ = visual weight. The relative auditory weight was computed as $w_A = ß_A / (ß_A + ß_V)$ (Fig. 2A-C).

By design, the temporal evolution of the physical visual uncertainty (i.e., STD of the visual cloud) is symmetric for each period in the sinusoid, RW1 and RW2 sequences. In other words, for physical visual uncertainty the 1st half and the flipped 2nd half within a period are identical (Fig. 2 and Fig. 3). Given this symmetry constraint we evaluated the influence of past visual uncertainty on participants' auditory weight w_A in terms of the difference between 1st half and flipped 2nd half. If prior visual uncertainty affects current visual uncertainty estimates, we would expect smaller w_A for the 1st half, when visual uncertainty increased, than the 2nd half, when visual uncertainty decreased over time. We evaluated this statistically in 2 (1st vs. flipped 2nd half) x 9 (time bins, i.e. removing the bins at maximal and minimal visual noise values) repeated measures ANOVAs performed independently for Sinusoid, RW1 and RW2 experiments (Tab. 1).

To further test whether the noise of past visual signals influenced observers' current visual uncertainty estimate, we employed a regression model where the relative auditory

weights were predicted by the visual STD in the current time bin and the difference in STD between the current and the previous time bin: $w_{A,t} = \sigma_{V,t} * ß_t + (\sigma_{V,t} - \sigma_{V,t-1}) * ß_{dt} + const * ß + e$ with $w_{A,t}$= relative auditory weight in time bin i; $\sigma_{V,t}$= visual STD in time bin i. To allow for generalization to the population level, the parameter estimates ($ß_t$, $ß_{dt}$) for each subject were entered into two-sided one-sample t-tests at the between-subject random effects level.

Sinusoidal sequence with intermittent changes in visual uncertainty

For each period of the sinusoidal sequence with intermittent changes, we sorted the physical visual uncertainty (i.e., the cloud's STD) and sound localization responses into 15 temporally adjacent bins which were positioned to optimally capture the jumps in visual uncertainty. For analysis of these sequences, we recombined the first and second halves of the 3 (increases at low, middle, high) x 3 (decreases at low, middle, high) sinewave cycles into three types of sinewave cycles such that both jumps were at low (= outer jump), middle (=middle jump) or high (= inner jump) visual uncertainty. This recombination makes the simplifying assumption that the jump position of the first half will have negligible effects on participants' uncertainty estimates of the second half. As for the continuous sequences, we quantified the auditory and visual influence on the perceived auditory location for each time bin based on the regression model: $R_A = L_A * ß_A + L_V * ß_V + const * ß + e$ independently for the 15 temporally adjacent bins and computed the relative auditory weight w_A. We statistically evaluated the influence of past visual noise on participants' auditory weight on the w_A in terms of the difference between 1st half and flipped 2nd half using a 2 (1st vs. flipped 2nd half) x 7 (time bins) x 3 (jump: inner, middle, outer) repeated measures ANOVAs (Tab. 1).

Computational Models (for continuous and discontinuous sequences)

To further characterize whether and how human observers use the uncertainty of previous visual signals to estimate the uncertainty of the current visual signal, we defined and compared three models where the visual uncertainty (λ_V, i.e. variance) was (1) estimated instantaneously for each trial (i.e., instantaneous learner), was updated via (2) Bayesian learning or (3) exponential discounting (i.e. exponential learner) (Supplementary Fig. S1).

In the following we will first describe the generative model that accounts for the fact that visual uncertainty usually changes slowly across trials (i.e., time-dependent uncertainty changes) and auditory and visual signals can be generated by one common or two independent sources (i.e., causal structure). Starting from this generative model that combines causal inference and uncertainty learning we will describe the models for the instantaneous learner, the Bayesian learner and the exponential learner. Finally, we will explain how we account for participants' internal noise and predict participants' responses from each model (i.e., the experimenter's uncertainty).

Generative model

On each trial t the subject is presented with an auditory signal $A_{,t}$, from a source $S_{A,t}$, (see Supplementary Figure S1) together with a visual cloud of dots at time t arising from a source, $S_{,Vt}$, drawn from a Normal distribution $S_{,Vt} \sim N(0, 1/\lambda_0)$ with the spatial reliability (i.e., inverse of the spatial variance): $\lambda_0 = 1/\sigma_0{}^2$. Critically, $S_{A,t}$ and $S_{V,t}$, can either be two independent sources (C = 2) or one common source (C=1): $S_{A,t} = S_{V,t} = S_t$ [13].

We assume that the auditory signal is corrupted by noise, so that the internal signal is $A_t \sim N(S_{A,t}, 1/\lambda_A)$. By contrast, the individual visual dots (presented at high visual contrast) are assumed to be uncorrupted by noise, but presented dispersed around the location $S_{V,t}$ according to $V_{i,t} \sim N(U_t, 1/\lambda_{V,t})$, where $U_t \sim N(S_{V,t}, 1/\lambda_{V,t})$. The dispersion of the individual dots, $1/\lambda_{V,t}$,

is assumed to be identical to the uncertainty about the visual mean, allowing subjects to use the dispersion as an estimate of the uncertainty about the visual mean.

The visual reliability of the visual cloud, $\lambda_{V,t} = 1/\sigma_{V,t}^2$, varies slowly at the re-display rate of 5 Hz according to a log random walk: $\log \lambda_{V,t} \sim N(\log \lambda_{V,t-1}, 1/\kappa)$ with κ being the variability of $\lambda_{V,t}$ in log space. This assumption is clearly not valid for the fourth experimental sequence ("Jump sequence"), which was specifically designed in order to test the validity of assumptions of the two models.

The generative models of the instantaneous, Bayesian and exponential learners all account for the causal uncertainty by explicitly modeling the two potential causal structures. Yet, they differ in how they estimate the visual uncertainty on each trial, which we will describe in greater detail below.

Observer Inference

The instantaneous, Bayesian and exponential learners invert this (or slightly modified, see below) generative model during perceptual inference to compute the posterior probability of the auditory location, $S_{A,t}$, given the observed A_t and $V_{i,t}$. The observer selects a response based on the posterior using a subjective utility function which we assume to be the minimization of the squared error $(S_{A,t}-S_{true})^2$. For all models the estimate for the location of the source is given by a combination of the common source assumption (see Fig. 1) and separate sources ($S_{A,t}$ and $S_{V,t}$):

$$\widehat{S_{A,t}} = \hat{S}_{A,C=1,t}\, P(C_t = 1|A_t, V_{1:n,t}) + \hat{S}_{A,C=2,t}(1 - P(C_t = 1|A_t, V_{1:n,t})) \quad (1)$$

where $\hat{S}_{A,C=1,t}$ and $\hat{S}_{A,C=2,t}$ depend on the model (see below), and $P(C = 1|A_t, V_{1:n,t})$ is the probability that the audio and visual stimuli originated from the same source according to Bayesian causal inference[13].

$$P(C_t = 1|A_t, V_{1:n,t}) =$$

$$\frac{(P(A_t, V_{1:n,t}|C = 1)P(C_t = 1))}{\left(P(A_t, V_{1:n,t}\left|C_t = 1\right)P(C_t = 1) + \left(P(A_t, V_{1:n,t}\left|C_t = 2\right)(1 - P(C_t = 1)\right)\right.} \quad (2)$$

Finally, for all models we assume that the observer pushes the button associated with the position closest to $\widehat{S_{A,t}}$. In the following we describe the generative and inference models for the instantaneous, Bayesian and exponential learners.

Model 1: Instantaneous learner

The instantaneous learning model ignores that the visual reliability (i.e., the inverse of visual uncertainty) of the current trial depends on the reliability of the previous trial. Instead, it estimates the visual reliability independently for each trial from the spread of the cloud of visual dots:

$$P\left(S_t, U_t, \lambda_{\mathrm{V,t}}\right| A_{1:t}, V_{1:n,1:t}) = P\left(S_t, U_t, \lambda_{\mathrm{V,t}}\right| A_t, V_{1:n,t}) =$$

$$P\left(C = 1\middle|A_t, V_{1:n,t}\right)P_{C=1}\left(S_t, U_t, \lambda_{\mathrm{V,t}}\right| A_t, V_{1:n,t}) +$$

$$P\left(C = 2\middle|A_t, V_{1:n,t}\right)P_{C=2}\left(S_t, U_t, \lambda_{\mathrm{V,t}}\right| A_t, V_{1:n,t}) =$$

$$P\left(C = 1\middle|A_t, V_{1:n,t}\right)\frac{P(S_t)P(A_t|S_t)P_{C=1}\left(U_t\middle|S_t, \lambda_{\mathrm{V,t}}\right)P\left(V_{1:n,t}\middle|U_t, \lambda_{\mathrm{V,t}}\right)P\left(\lambda_{\mathrm{V,t}}\right)}{Z_1} +$$

$$\left(1 - P\left(C = 1\middle|A_t, V_{1:n,t}\right)\right) P(S_{A,t})P\left(A_t\middle|S_{A,t}\right)/Z_2.$$

Apart from $P(C = 1|A_t, V_t)$ these terms are all normal distributions, while we assume in this model that $P\left(\lambda_{\mathrm{V,t}}\right)$ is uninformative. Hence, visual reliability is computed from the sample variance: $\widehat{\lambda_{Vt}} = 1/\sigma_{Vt}^2$ where $\sigma_{Vt}^2 = 1/(\mathrm{n} - 1)\sum_{i=1}^{n}(V_{i,t} - \bar{V}_{i,t})^2$ is the sample variance (and $\bar{V}_t = 1/\mathrm{n}\sum_{i=1}^{n} V_{i,t}$ is the sample mean). The causal component estimates are given by:

$$\hat{S}_{A,C=1,t} = \frac{\widehat{\lambda}_{V,t}\bar{V}_t + \lambda_A \mathrm{A}_t + \lambda_S 0}{\lambda_{V,t} + \lambda_A + \lambda_0} \quad (4)$$

$$\hat{S}_{A,C=2,t} = \frac{\lambda_A A_t + \lambda_0 0}{\lambda_A + \lambda_0} \quad (5)$$

These two components are then combined based on their relative posterior probability (see equation 1). This model is functionally equivalent to a Bayesian causal inference model as described in Koerding et a l. (2007), but with visual reliability computed directly from the sample variance rather than a fixed unknown parameter (which the experimenter estimates during model fitting).

Model 2: Bayesian learner

The Bayesian learner capitalizes on the slow changes in visual reliability across trials and combines past and current inputs to provide a more reliable estimate of visual reliability and hence auditory location. It computes the posterior probability based on all auditory and visual signals presented until time t (here only shown for C=1, refer to supplementary material for full derivation).

According to Bayes rule the joint probability of all variables until time *t* can be written based on the generative model as:

$$P(\lambda_{\mathrm{V},1:\mathrm{t}}, A_{1:t}, U_{1:t}, V_{1:n,1:t}, S_{1:t}) = P(A_1|S_1)P(V_{1:n,1}|U_1, \lambda_{\mathrm{V},1})P(U_1|S_1, \lambda_{\mathrm{V},1})P(S_1)P(\lambda_{\mathrm{V},1})$$
$$\prod_{k=2}^{t} P(A_k|S_k)P(V_{1:n,k}|U_k, \lambda_{\mathrm{V,k}})P(U_k|S_k, \lambda_{\mathrm{V,k}})P(\lambda_{\mathrm{V,k}}|\lambda_{\mathrm{V,k-1}})P(S_k) \quad (6)$$

As above, the visual likelihood is given by the product of individual Normal distributions for each dot *i*: $P(V_{1:n,t}|U_t, \lambda_{\mathrm{V,t}}) = \prod_{i=1}^{n} N(V_{i,t}|U_t, 1/\lambda_{\mathrm{V,t}})$, and $P(U_t|S_t, \lambda_{\mathrm{V,t}}) = N(U_t|S_t, 1/\lambda_{\mathrm{V,t}})$. The prior $P(S_t)$ is a Normal distribution $N(S_t|0, 1/\lambda_{\mathrm{S}})$ and the auditory likelihood $P(A_{t,}|S_t)$ is a Normal distribution $N(A_t|S_t, 1/\lambda_A)$. As described in the generative model $P(\lambda_{\mathrm{V,k}}|\lambda_{\mathrm{V,k-1}})$ is given by $\log \lambda_{\mathrm{V,t}} \sim N(\log \lambda_{\mathrm{V,t-1}}, 1/\kappa)$.

Importantly, only the visual reliability, $\lambda_{V,t}$, is directly dependent on the previous trial ($P(\lambda_{V,k}, \lambda_{V,k-1}) = P(\lambda_{V,k}|\lambda_{V,k-1})P(\lambda_{V,k-1}) \neq P(\lambda_{V,k})P(\lambda_{V,k-1})$). Because of the Markov property (i.e. $\lambda_{V,t}$ depends only on $\lambda_{V,t-1}$) the joint distribution for time *t* can be written as

$$P(\lambda_{V,t}, \lambda_{V,t-1}, A_t, U_t, V_{1:n,t}, S_t) =$$

$$P(A_t|S_t)P(U_t|S_t, \lambda_{V,t})P(V_{1:n}|U_t, \lambda_{V,t})P(\lambda_{V,t}|\lambda_{V,t-1})P(\lambda_{V,t-1}|V_{1:n,t-1}, A_{t-1})P(S_t). \quad (7)$$

Hence, the joint posterior probability over location and visual reliability given a stream of auditory and visual inputs can be rewritten as:

$$P(S_t, U_t, \lambda_{V,t}|\, A_{1:t}, V_{1:n,1:t}) =$$

$$P(S_t)P(A_t|S_t)P(U_t|S_t, \lambda_{V,t})P(V_{1:n}|U_t, \lambda_{V,t}) \int P(\lambda_{V,t}|\lambda_{V,t-1})P(\lambda_{V,t-1}|V_{1:n,t-1}, A_{t-1})\, d\lambda_{V,t-1}/$$

$$Z. \quad (8)$$

As this equation cannot be solved analytically, we obtain an approximate solution by factorizing the posterior in terms of the unknown variables ($S_t, U_t, \lambda_{V,t}$) according to the method of variational Bayes[23]. For a full derivation of this approach see Supplementary Material, Methods.

In short, a derivation is performed for the case of a common cause (C=1), as well as for the separate causes (C=2), and the two are then combined based on their relative posterior likelihood $P(C = 1|A_t, V_{1:n,t})$, as described above in equation 1.

Model 3: Exponential learner

Finally, the observer may approximate the full Bayesian inference of the Bayesian learner by the more simple heuristic strategy of exponential discounting. In the exponential discounting model the observer learns the visual reliability by exponentially discounting past visual reliability estimates:

$$\hat{\lambda}_{V,t-1} = 1/\sigma_{Vt}^2\,(1 - \gamma) + \hat{\lambda}_{V,t-1}\gamma \quad (9)$$

Similar to the optimal Bayesian learner (above) this observer model uses the past to compute the current reliability, but it does so based on a fixed discounting rate γ. Computation is otherwise performed in accordance with models M1 and M2, eqs 1-2, 4-5.

Inference by the experimenter

From the observer's viewpoint this completes the inference process. However, from the experimenter's viewpoint the internal variable for the auditory stimulus, A_t, is unknown and not directly under the experimenter's control. To integrate out this unknown variable we generated 1,000 samples of the internal auditory value for each trial from the generative process $A_t \sim N(S_{A,t,true}, \sigma_A^2)$, where $S_{A,t,true}$ was the true location the auditory stimulus came from. For each value of A_t we obtained a single estimate $\hat{S}_t$ (as described above). To link these estimates with observers' button response data, we assumed that subjects push the button associated with the position closest to $\hat{S}_{A,t}$. In this way, we obtained a histogram of responses for each subject and trial which provide the likelihood of the model parameters given subject's responses, $P(resp_t|\kappa, \sigma_A, P_{common}, S_{A,t,true}, S_{V,t,true})$.

Model estimation and comparison

Parameters for each model (all models: σ_A, P_{common}=P(C=1), σ_0. Bayesian learner: □. exponential learner: γ) were fit for each individual subject by sampling using a symmetric proposal Metropolis-Hasting (MH) algorithm (with A_t integrated out via sampling, see above). The MH algorithm iteratively draws samples set_n from a probability distribution through a variant of rejection sampling: if the likelihood of the parameter set is larger than the previous set the new set is accepted, otherwise it is accepted with probability $L(model|set_n)/L(model|set_{n-1})$, where $L(model|set)= \prod_t P(resp_t|\kappa, \sigma_A, P_{common}, S_{A,t,true}, S_{V,t,true})$(for Bayesian learner). We sampled 4000 steps from 4 sampling chains with thinning (only using every 4th sampling

to avoid correlations in samples), giving a total of 4000 samples per subject data sets. Convergence was assessed through scale reduction (using criterion $R<1.1$[24]). Using sampling does not just provide a single parameter estimate for a data set (as when fitting maximum likelihood), but can instead be used to assess the uncertainty in estimation for the data set. The model code was implemented in Matlab (Mathworks, MA) and ran on two dual Xeon workstations. Each sample step, per subject data set, took 30 seconds on a single core (~42 hours per sampling chain).

Quantitative Bayesian model comparison of the three candidate models was based on the Watanabe-Akaike Information Criterion (WAIC) as an approximation to the out of sample cross validation[24]. For the comparison between models we performed a non-parametric permutation test (also known as a randomization test), comparing the observed difference in WAIC values with a null distribution where the labels of the models were randomly allocated (i.e. the sign of the WAIC difference between models for each subject were randomly permuted).

Simulated localization responses

To qualitatively compare the localization responses given by the participants and the responses predicted by the instantaneous, Bayesian and exponential learner, we computed the auditory weight w_A from the predictions of the three models exactly as in the analysis for the behavioral data. For illustration, we then show and compare the model's w_A from the 1st and the flipped 2nd half of the periods for each of the four experiments (cf Fig. 3B-D and Supplementary Fig. S2A, B).

Figures

Figure 1

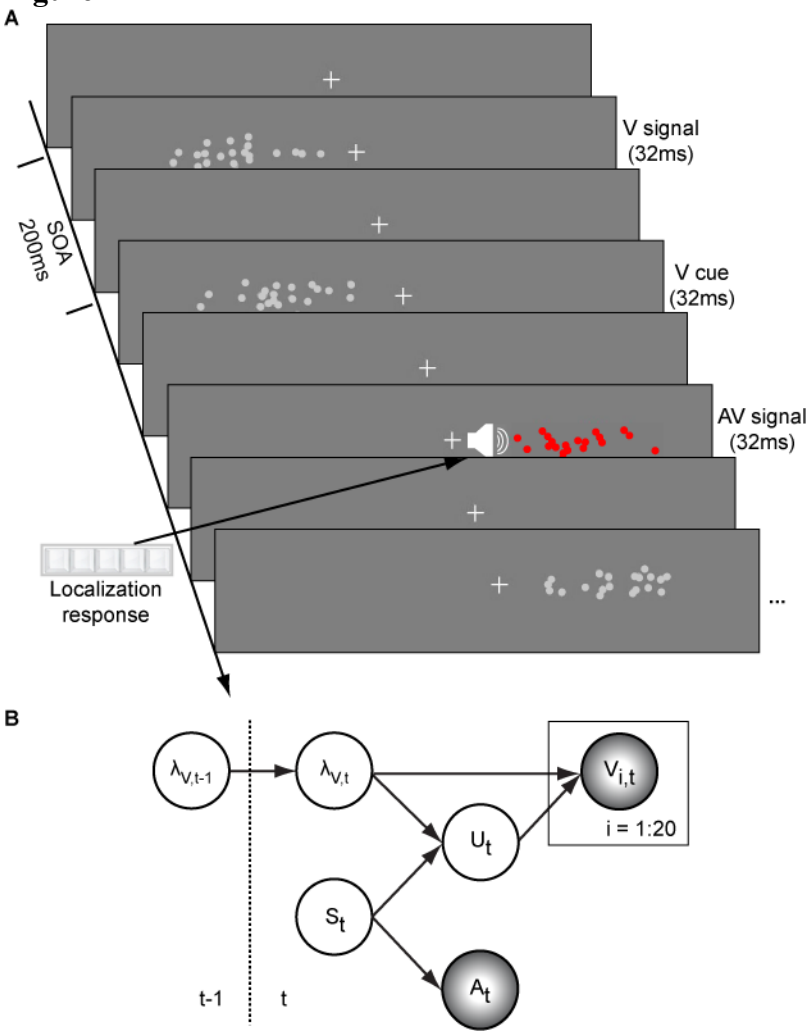

Figure 2

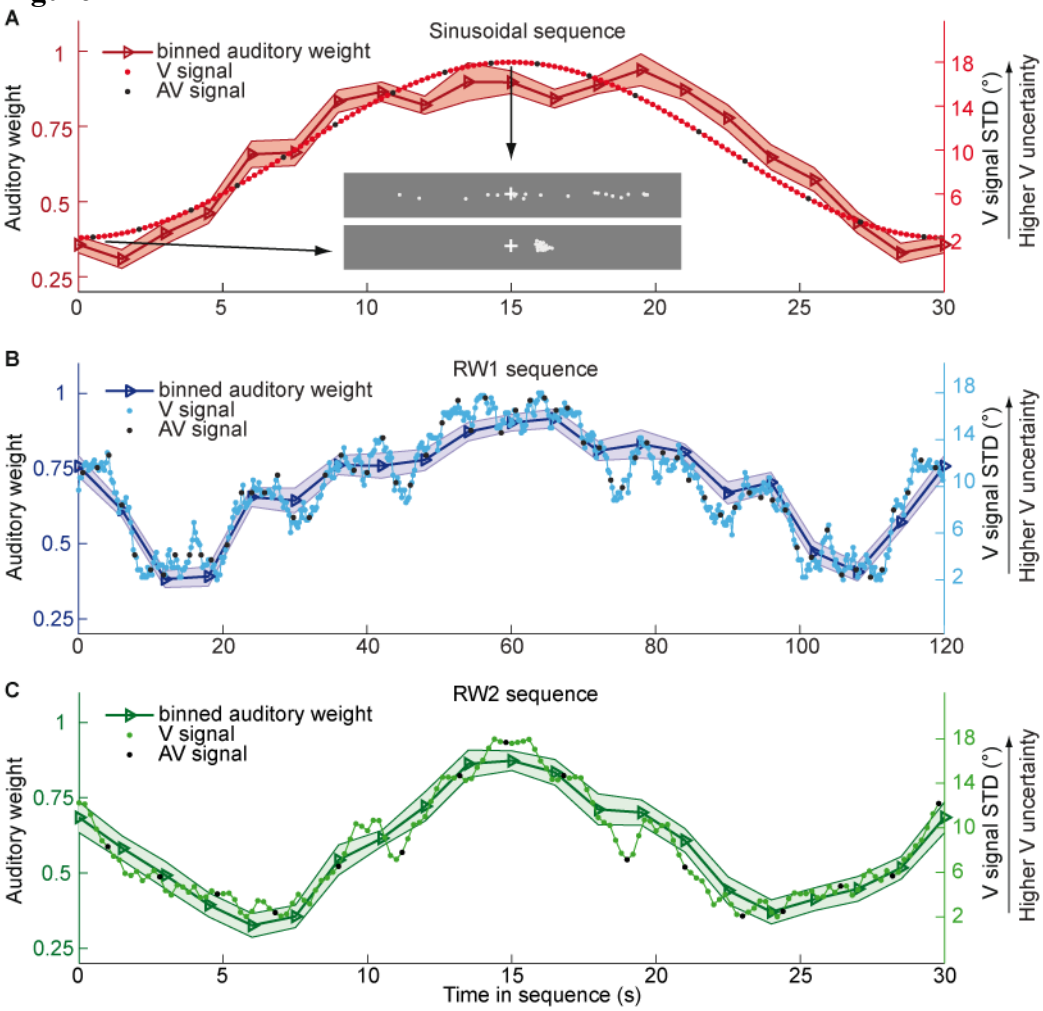
A
Sinusoidal sequence
binned auditory weight
V signal
AV signal
Auditory weight
V signal STD (°)
Higher V uncertainty
B
RW1 sequence
binned auditory weight
V signal
AV signal
Auditory weight
V signal STD (°)
Higher V uncertainty
C
RW2 sequence
binned auditory weight
V signal
AV signal
Auditory weight
V signal STD (°)
Higher V uncertainty
Time in sequence (s)

Figure 3

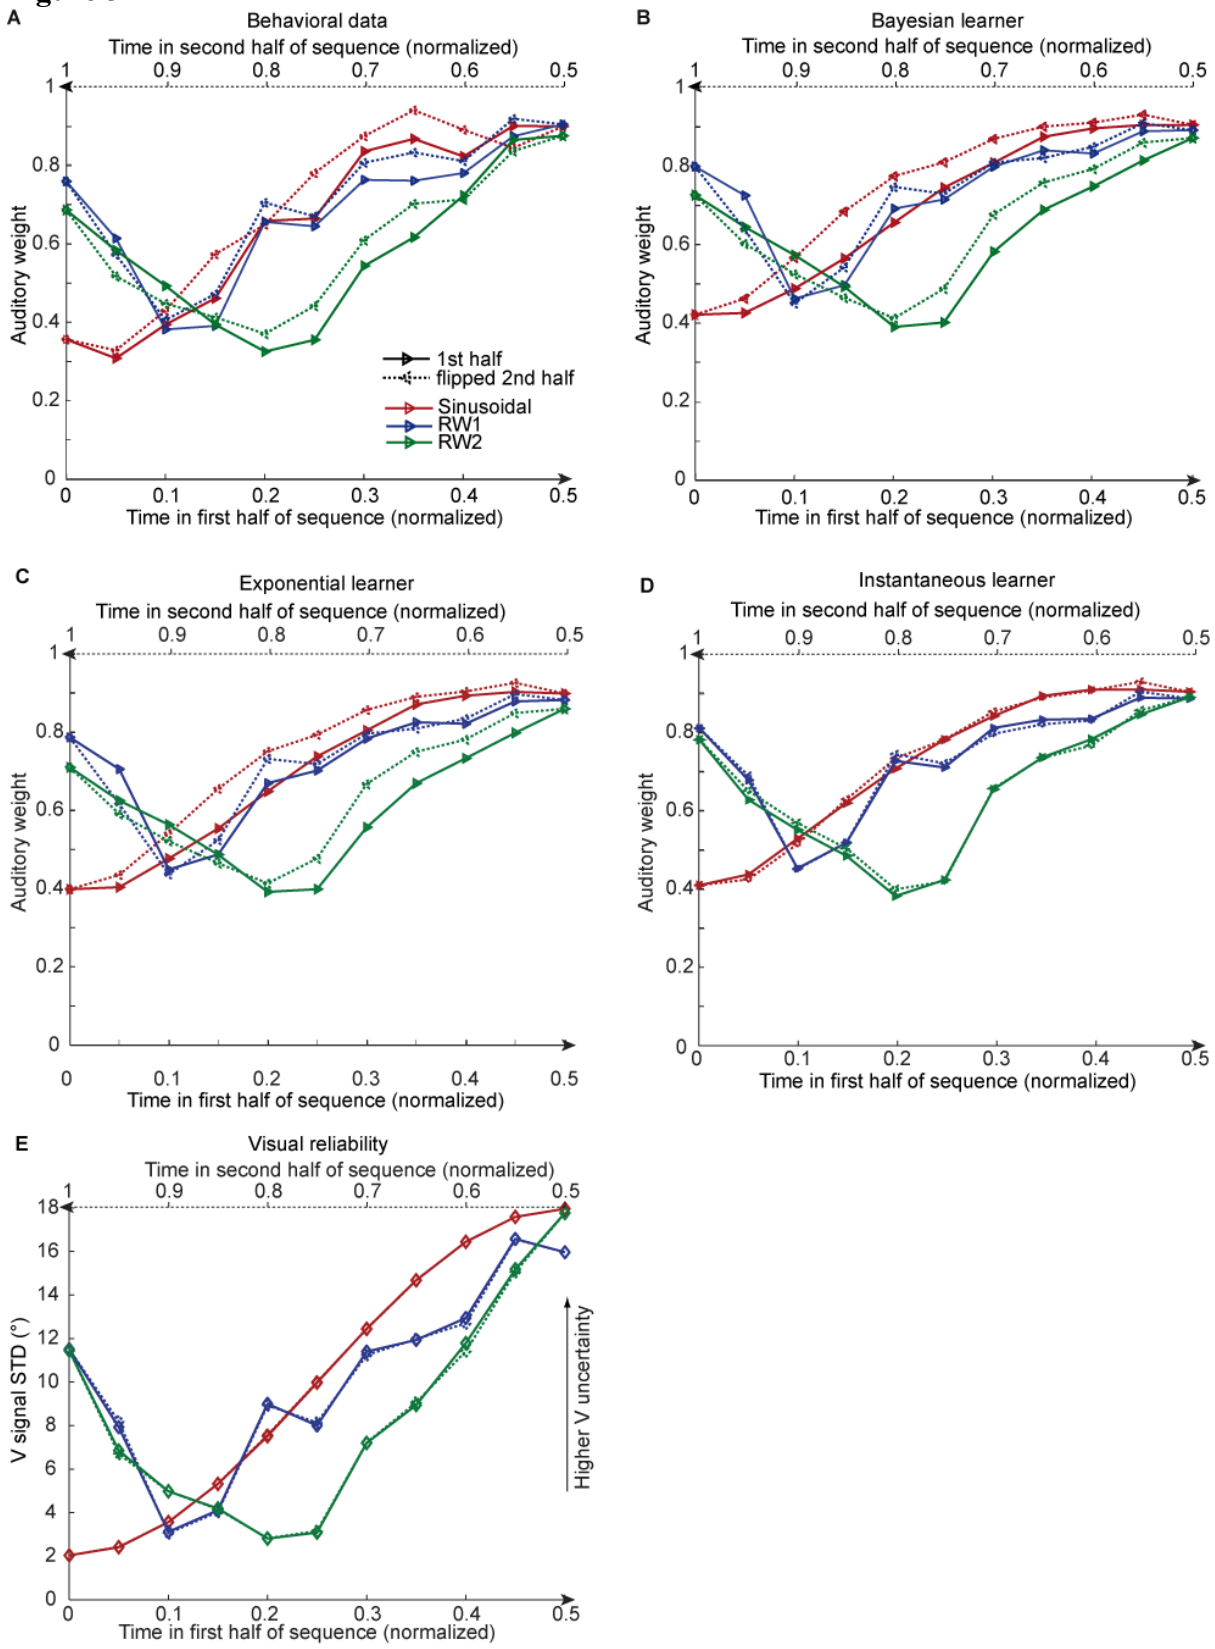
A
Behavioral data
Time in second half of sequence (normalized)
Auditory weight
Time in first half of sequence (normalized)
1st half
flipped 2nd half
Sinusoidal
RW1
RW2
B
Bayesian learner
C
Exponential learner
D
Instantaneous learner
E
Visual reliability
V signal STD (°)
Higher V uncertainty

Figure 4

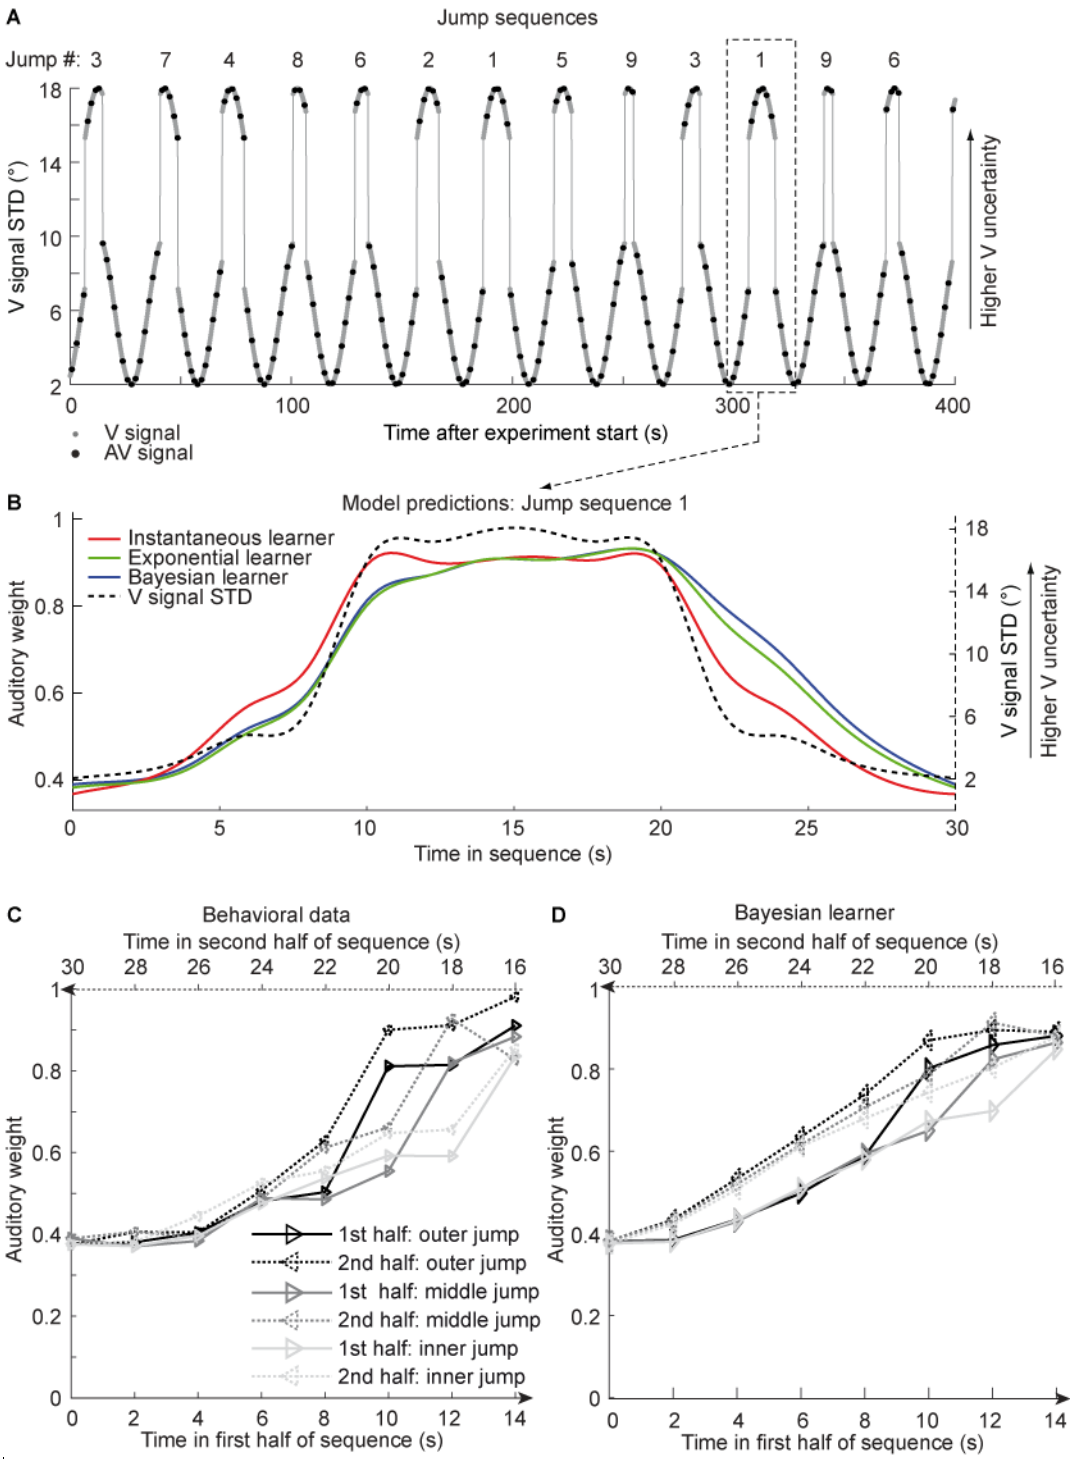
A
Jump sequences
Jump #: 3 7 4 8 6 2 1 5 9 3 1 9 6
V signal STD (°)
Higher V uncertainty
V signal
AV signal
Time after experiment start (s)
B
Model predictions: Jump sequence 1
Instantaneous learner
Exponential learner
Bayesian learner
V signal STD
Auditory weight
V signal STD (°)
Higher V uncertainty
Time in sequence (s)
C
Behavioral data
Time in second half of sequence (s)
Auditory weight
1st half: outer jump
2nd half: outer jump
1st half: middle jump
2nd half: middle jump
1st half: inner jump
2nd half: inner jump
Time in first half of sequence (s)
D
Bayesian learner
Time in second half of sequence (s)
Auditory weight
Time in first half of sequence (s)

References

1 Acerbi, L., Vijayakumar, S. & Wolpert, D. M. On the origins of suboptimality in human probabilistic inference. *PLoS computational biology* **10**, e1003661 (2014).

2 Drugowitsch, J., Wyart, V., Devauchelle, A.-D. & Koechlin, E. Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* **92**, 1398-1411 (2016).

3 Shen, S. & Ma, W. J. A detailed comparison of optimality and simplicity in perceptual decision making. *Psychological review* **123**, 452 (2016).

4 Alais, D. & Burr, D. The ventriloquist effect results from near-optimal bimodal integration. *Current biology : CB* **14**, 257-262 (2004).

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

6 Jacobs, R. A. Optimal integration of texture and motion cues to depth. *Vision Res* **39**, 3621-3629 (1999).

7 Knill, D. C. & Pouget, A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* **27**, 712-719 (2004).

8 van Beers, R. J., Sittig, A. C. & Gon, J. J. D. v. d. Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of neurophysiology* **81**, 1355-1364 (1999).

9 Ma, W. J. & Jazayeri, M. Neural coding of uncertainty and probability. *Annual review of neuroscience* **37**, 205-220 (2014).

10 Zemel, R. S., Dayan, P. & Pouget, A. Probabilistic interpretation of population codes. *Neural computation* **10**, 403-430 (1998).

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

12 Battaglia, P. W., Jacobs, R. A. & Aslin, R. N. Bayesian integration of visual and auditory signals for spatial localization. *J Opt Soc Am A Opt Image Sci Vis* **20**, 1391-1397 (2003).

13 Kording, K. P. *et al.* Causal inference in multisensory perception. *PloS one* **2**, e943 (2007).

14 Gelman, A., Hwang, J. & Vehtari, A. Understanding predictive information criteria for Bayesian models. *Statistics and computing* **24**, 997-1016 (2014).

15 Jacobs, R. A. & Fine, I. Experience-dependent integration of texture and motion cues to depth. *Vision Research* **39**, 4062-4075 (1999).

16 Mikula, L., Gaveau, V., Pisella, L., Khan, A. Z. & Blohm, G. Learned rather than online relative weighting of visual-proprioceptive sensory cues. *Journal of neurophysiology* **119**, 1981-1992 (2018).

17 Triesch, J., Ballard, D. H. & Jacobs, R. A. Fast temporal dynamics of visual cue integration. *Perception* **31**, 421-434 (2002).

18 Beck, J. M. *et al.* Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142-1152 (2008).

19 Fiser, J., Berkes, P., Orban, G. & Lengyel, M. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn Sci* **14**, 119-130 (2010).

20 Algazi, V. R., Duda, R. O., Thompson, D. M. & Avendano, C. in *Applications of Signal Processing to Audio and Acoustics, 2001 IEEE Workshop on the.* 99-102 (IEEE).

21 Brainard, D. H. The psychophysics toolbox. *Spatial vision* **10**, 433-436 (1997).

22 Kleiner, M. *et al.* What's new in Psychtoolbox-3. *Perception* **36**, 1.1-16 (2007).

23 Bishop, C. M. *Pattern recognition and machine learning*. Vol. 4 (Springer, 2006).
24 Gelman, A. *et al. Bayesian data analysis*. (Chapman and Hall/CRC, 2013).

Acknowledgements

This study was funded by the ERC (ERC-multsens, 309349) and the Deutsche Forschungsgemeinschaft (DFG; grant number RO 5587/1-1).

Author contributions

TR, OS and UN conceived the experiment. TR and AF collected the data. UB, TR and UN analyzed the data. UB developed the computational models. UB, TR and UN wrote the manuscript.

Competing interests

The authors report no conflict of interest.

Figure legends

Figure 1. Audiovisual localization paradigm and Bayesian causal inference model for learning visual reliability. (**A**) Visual (V) signals (cloud of 20 bright dots) were presented at 5 Hz (i.e. every 200 ms). The cloud's location mean was temporally independently resampled from five possible locations (-10°, -5°, 0°, 5°, 10°) at a SOA jittered between 1.4 and 2.8 s. In synchrony with the change in the cloud's location, the dots changed their colour and a sound was presented (AV signal) which the participants localized using five response buttons. The location of the sound was sampled from the two possible locations adjacent to the visual cloud's mean location (i.e. ± 5° AV spatial discrepancy). (**B**) The generative model for the Bayesian learner explicitly modelled the potential causal structures, i.e. whether visual (V_i) signals and an auditory (A) signal are generated by one common audiovisual source S_t , i.e. C = 1, or by two independent sources S_{Vt} and S_{At} , i.e. C = 2 (n.b. only the model component for the common source case is shown to illustrate the temporal updating, for complete generative model, see Supplementary Fig. S1). Importantly, the reliability (i.e., 1/variance) of the visual signal at time t (λ_t) depends on the reliability of the previous visual signal (λ_{t-1}) for both model components (i.e. common source and independent source).

Figure 2. Time course of visual noise and relative auditory weights for continuous sequences of visual noise. Relative auditory weights (mean across participants ± SEM, left ordinate) and visual noise (i.e., STD of the cloud of dots, right ordinate) are displayed as a function of time. The relative auditory weight varies between one (i.e. pure auditory influence on the localization responses) and zero (i.e. pure visual influence). The STD of the visual cloud was manipulated as (**A**) a sinusoid (period 30s, N = 25), (**B**) a random walk (RW1, period 120s, N = 33) and (**C**) a smoothed random walk (RW2, period 30s, N = 19). Please note that the

period for RW1 sequence is 120 s, while the periods of the sinusoid and RW2 is only 30 s. As a result the overall dynamics as quantified by the power spectrum is faster for RW2 than RW1 (peak in frequency range [0 0.2] Hz: Sin: 0.033 Hz, RW1: 0.025 Hz, RW2: 0.066 Hz). The sequence of visual clouds was presented at 5 Hz, while trials, i.e. audiovisual (AV) signals (colour change with sound presentation, black dots), were interspersed with a SOA jittered between 1.4 and 2.8 s. For illustration purposes, the cloud of dots for the lowest (i.e., V signal STD = 2°) and the highest (i.e., V signal STD = 18°) visual variance are shown in (**A**).

Figure 3. Observed and predicted relative auditory weights for continuous sequences of visual noise. Relative auditory weights w_A of the 1st (solid) and the flipped 2nd half (dashed) of a period (binned into 20 time bins) plotted as a function of the normalized time in the sinusoidal (red), the RW1 (blue) and the RW2 (green) sequences. Relative auditory weights were computed from auditory localization responses of human observers **(A)**, Bayesian **(B)**, exponential **(C)** or instantaneous **(D)** learning models. For comparison, the standard deviation of the visual signal is shown in **(E)**. Please note that all models were fitted to observers' auditory localization responses (i.e. not the auditory weight w_A).

Figure 4. Time course of visual noise and relative auditory weights for sinusoidal sequence with intermittent jumps in visual noise. (A) The visual noise (i.e., STD of the cloud of dots, right ordinate) are displayed as a function of time. Each cycle included one abrupt increase and decrease in visual noise. The sequence of visual clouds was presented at 5 Hz while audiovisual (AV) signals (black dots) were interspersed with a SOA jittered between 1.4 and 2.8 s. **(B)** Visual noise (i.e STD of the cloud of dots; right y axis) and the relative auditory weights (left y axis) computed from auditory localization responses of the Bayesian,

exponential, or instantaneous learning models for jump sequence 1. Critically, while the Bayesian and the exponential learners adapt similarly to abrupt increases in visual variance, the Bayesian learner adapts more slowly to abrupt decreases in visual variance. **(C, D)** Relative auditory weights w_A of the 1st (solid) and the flipped 2nd half (dashed) of a period (binned into 15 time bins) plotted as a function of the time in the sinusoidal sequence with intermitted inner (light gray), middle (gray) and outer (dark gray) jumps. Relative auditory weights were computed from auditory localization responses of human observers **(C)** and Bayesian learning model **(D).** Please note that all models were fitted to observers' auditory localization responses (i.e. not the auditory weight w_A).

Tables

Table 1. Analyses of the temporal asymmetry of the relative auditory weights across the four sequences of visual noise using repeated measures ANOVAs with the factors sequence part (1st vs. flipped 2nd half), time bin and jump position (only for the sinusoid sequences with intermittant changes).

	Effect	F	df1	df2	p	Partial η^2
Sinusoid	Part	12.162	1	24	0.002	0.336
	Bin	92.007	3.108	74.584	<0.001	0.793
	PartXBin	2.167	2.942	70.617	0.101	0.083
RW1	Part	14.129	1	32	0.001	0.306
	Bin	76.055	4.911	157.151	<0.001	0.704
	PartXBin	1.225	4.874	155.971	0.300	0.037
RW2	Part	2.884	1	18	0.107	0.138
	Bin	60.142	3.304	59.467	<0.001	0.770
	PartXBin	3.385	4.603	82.849	0.010	0.158
Sinusoid with intermittent changes	Jump	28.306	2	34	<0.001	0.625
	Part	24.824	1	17	<0.001	0.594
	Bin	76.476	1.873	31.839	<0.001	0.818
	JumpXPart	0.300	2	34	0.743	0.017
	JumpXBin	8.383	3.309	56.247	<0.001	0.330

PartXBin	1.641	3.248	55.222	0.187	0.088
JumpXPartXBin	0.640	5.716	97.175	0.690	0.036

Note: The factor time bin comprised 9 levels in the first three and 7 levels in the fourth sequence. In this sequence, the factor Jump comprised three levels. If Mauchly tests indicated significant deviations from sphericity ($p < 0.05$), we report Greenhouse-Geisser corrected degrees of freedom and p values.

Table 2. Model parameters (median) and relative WAIC values for the three candidate models in the four sequences of visual noise.

Sequence	Model	σ_A	P_{common}	σ_0	κ or γ	ΔWAIC
Sinusoid	Instantaneous learner	5.56	0.63	8.95	-	109.9
	Bayesian learner	5.64	0.65	9.03	κ: 7.37	0
	Exponential discounting	5.62	0.64	9.02	γ: 0.23	45.6
RW1	Instantaneous learner	6.30	0.69	8.46	-	89.0
	Bayesian learner	6.29	0.72	8.68	κ: 8.06	0
	Exponential discounting	6.26	0.70	8.75	γ : 0.33	32.3
RW2	Instantaneous learner	6.36	0.72	10.79	-	201.3
	Bayesian learner	6.49	0.78	10.9	κ: 6.7	0
	Exponential discounting	6.46	0.73	11.0	γ : 0.25	46.3
Sinusoid with intermittent changes	Instantaneous learner	6.38	0.65	8.19	-	94.9
	Bayesian learner	6.45	0.68	8.26	κ: 6.13	0
	Exponential discounting	6.43	0.67	8.20	γ: 0.24	1.64