Optimality under fire: Dissociating learning from Bayesian integration

Luigi Acerbi[1], Bernard Marius 't Hart[2], Feryal M. P. Behbahani[3], & Megan A. K. Peters[4]
[1)] University of Edinburgh, Edinburgh, UK, [2)] Philipps-Universität Marburg, Marburg, Germany, [3)] Imperial College London, London, UK, [4)] University of California Los Angeles, Los Angeles, CA, USA

Introduction: In both unisensory and multisensory tasks, human observers have repeatedly been shown to be optimal or near-optimal in their integration of multiple cues (Ernst & Banks, 2002; Körding et al., 2007). Most of the research on cue integration has assumed that the noise in each cue follows a normal distribution, and thus that (a) the variance of the noise is a perfect indicator of the reliability of the cue, and (b) optimal integration is therefore achieved via a linear combination of the cues. However, little is known about how humans might integrate other noise distributions, e.g., those that may not be symmetric or unimodal, or which may require nonlinear cue combination. Here we ask if human observers are able to learn both lower-order and higher-order statistics (e.g., skewness) of non-normal distributions, and whether and how the acquired statistical features of such distributions affect cue integration.

Methods: In order to test human observers' capacity for learning and integrating a non-normal distribution, we had 25 observers engage in a "rifle-testing game" (Fig. 1A&B). There were two distinct "rifles", each with a different "shooting pattern": one rifle's "shots" were generated according to a normal distribution (mean: 0, std: 0.075 normalized screen units); those for the other rifle were generated according to a shifted exponential distribution (mean: 0, std: 0.05 screen units; fig. 1C). In the "distribution learning" phase of the task, observers witnessed the landing position of 20 shots fired at a visible target, and subsequently were asked to reproduce the "firing pattern" using only 10 shots (Fig. 1A). This was repeated for each gun four times. The learning phase was followed by a "cue integration" phase (~200 trials), in which observers were shown the landing position of two shots, one from each rifle, towards the same non-visible target (whose location changed randomly on every trial). Participants were asked to indicate the most likely position of the target at which the two shots had been fired (Fig. 1B). In order to test changes in the internal representation of the distributions, observers were asked to generate again the firing patterns of each of the guns intermittently during the integration phase. This whole sequence (distribution learning and cue integration) was repeated in two consecutive sessions.

Results: Observers performed well in the "distribution learning" part of the task (Fig. 2A). On average, participants' reproduced firing patterns quantitatively matched lower-order statistics (mean and variance) of the true distributions; moreover, skewness of the patterns was correctly near-zero for the Gaussian rifle and significantly positive for the exponential rifle (mean reproduced skewness: 0.6 ± 0.1; true skewness: 2.0), showing a qualitative learning of higher-order statistics as well. However, in the "cue integration" task, almost all subjects integrated the cues linearly, notwithstanding the nonlinearity of the optimal solution (Fig. 2B&C). A DIC model comparison between several Bayesian observer models with different internal representations of the rifle distributions confirmed that observers were most likely approximating both distributions as Gaussians (Fig. 2C). Interestingly, the linear weights assigned to the cues by the observers were optimal for the Gaussian approximation.

Discussion: In the "distribution learning" task, observers clearly detected and reproduced a skewed distribution for the exponential rifle's shots. However, in the main experiment only 4 out of 25 observers took into account the asymmetry of the exponential distribution when performing cue integration. Our results suggest that statistical features of the distributions and weights adopted in cue integration have independent internal representations. Whether and how it is possible to consistently elicit transfer between statistical learning of (non-Gaussian) distributions and subsequent (non-linear) integration involving the same distributions remains an exciting question, open to investigation.

References:

Ernst, M.O., & Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature, 415, 439-443.

Körding, K.P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J.B., & Shams, L. (2007). Causal inference in multisensory perception. PLoS ONE, 9, 1-10.

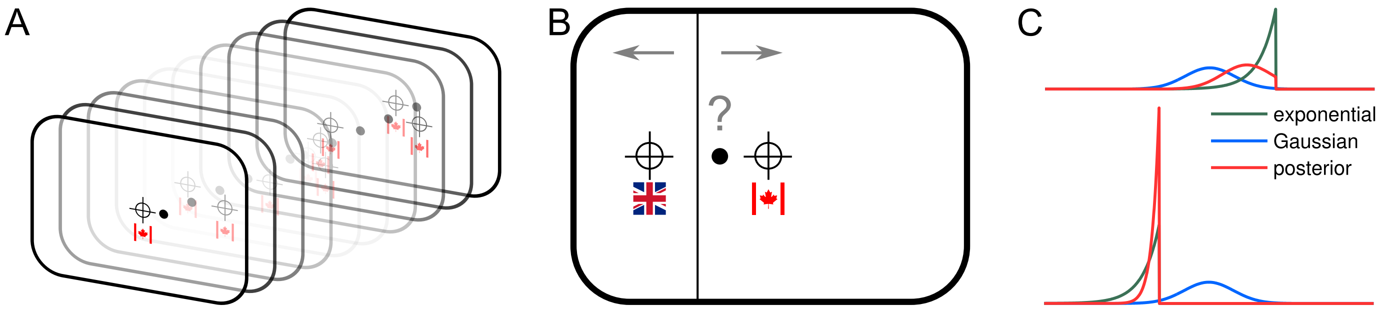

Figure 1. Task and distributions. A) Distribution learning: Participants saw the firing pattern of two "rifles" in 20 shots (represented by crosshairs and a flag identifying the rifle) relative to a target dot, and recreated the pattern using only 10 shots. **B) Cue integration:** Participants had to indicate the most likely position of the target at which the two rifles had been fired, by horizontally moving a vertical line using the mouse. After they indicated where they thought the target would be, the true location was shown. **C) Distributions:** With one of the patterns asymmetric, the shape of the posterior distribution differed when the exponential likelihood function fell on one side of the Gaussian likelihood function versus the other.

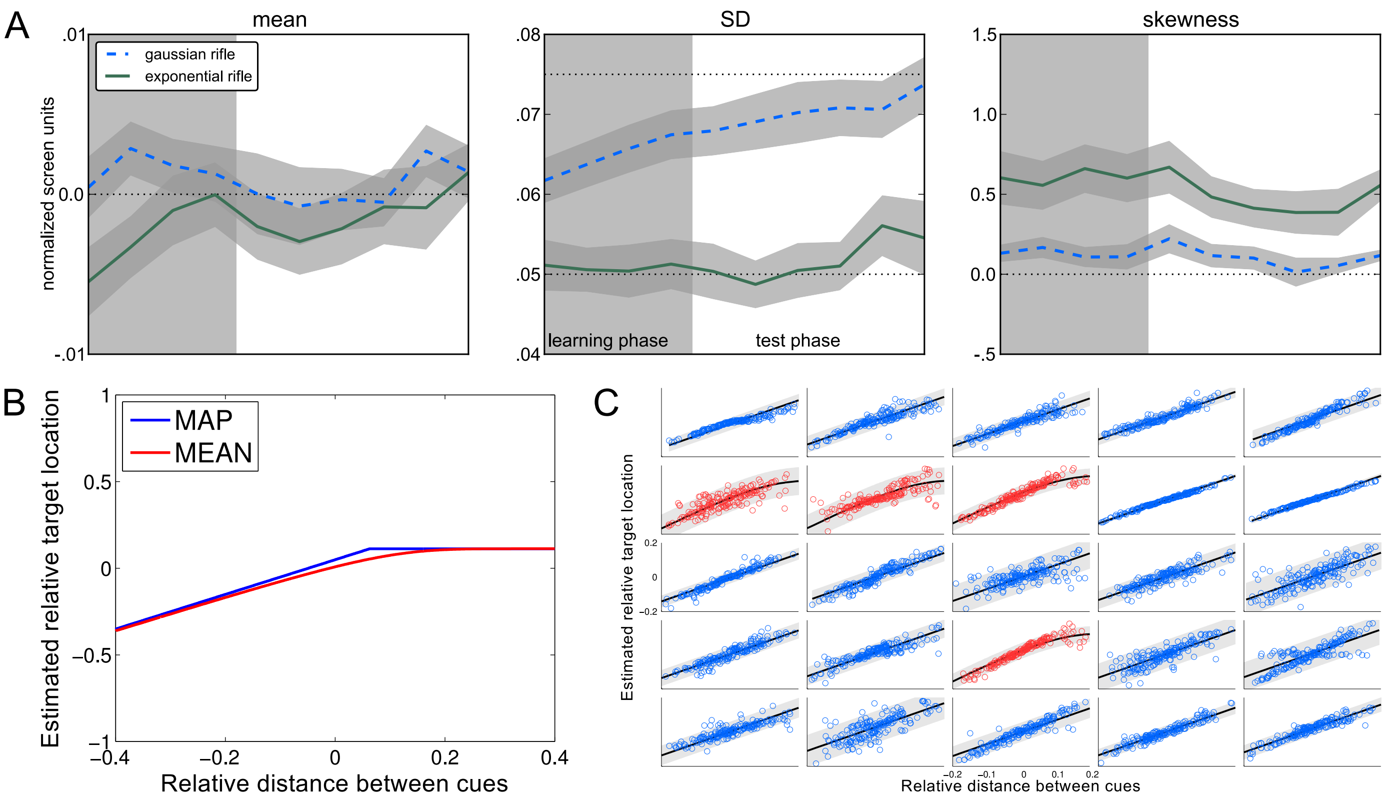

Figure 2. Results. A) Distribution learning: Reproduction data from the first session were used to assess learning of distribution moments. Lines indicate average moments of reproduced firing patterns (*continuous*: Gaussian rifle; *dashed*: exponential rifle). Shaded areas denote standard errors of the mean. Dotted lines indicate true values (true skew is 2.0 for the exponential). Estimates of the mean of the exponential and the standard deviation of the Gaussian improved over the course of the session. Skewness was approximated slightly better in the learning phase as compared to the test phase. **B) Optimal response profile:** Optimal guess for target location as a function of the displacement of the exponential rifle's shot with respect to the Gaussian rifle's shot (*blue:* MAP rule; *red*: MEAN rule). **C) Cue integration and model comparison:** Dots represent individual responses in the cue integration task, lines are the best model fit; each panel shows a participant. Despite demonstrating learning of both lower-order and higher-order statistics, participants do not integrate the distributions optimally. A Gaussian approximation model using only the mean and variance of the presented distributions, or variants thereof, fit participant responses best (in *blue*) for all but four subjects (in *red*).

Acknowledgments: The majority of the work was done as a team project during the CoSMo 2013 summer school at Queen's University, Kingston ON, Canada. We thank Dr. Gunnar Blohm and Dr. Konrad Körding for organizing the summer school, and Dr. Paul Schrater for valuable feedback given during the workshop. We also thank our fellow CoSMo students for volunteering as participants in this study.