Target article: *Suboptimality in perceptual decision making* by Rahnev and Denison

Word count:

Abstract: 57 words (max: 60) Main text: 982 words (max: 1,000)

References: 151 words Entire text: 1,287 words

Leveraging decision consistency to decompose suboptimality in terms of its ultimate predictability

Valentin Wyart

Laboratoire de Neurosciences Cognitives et Computationnelles, Institut National de la Santé et de la Recherche Médicale, Département d'Etudes Cognitives, Ecole Normale Supérieure, PSL Research University, Paris, France

Full institutional mailing address:

Ecole Normale Supérieure, Département d'Etudes Cognitives
Laboratoire de Neurosciences Cognitives et Computationnelles
29 rue d'Ulm, 75005 Paris, France

Phone: +33 (0)1 44 32 29 54 **Fax:** +33 (0)1 44 32 26 42

E-mail: valentin.wyart@ens.fr **URL:** <http://inc2.dec.ens.fr/inference-and-decision-making>

Abstract: While the suboptimality of perceptual decision-making is indisputable in its strictest sense, characterizing the nature of suboptimalities constitutes a valuable drive for future research. I argue that decision consistency offers a rarely measured, yet important behavioral metric for decomposing suboptimality (or, more generally, deviations from any candidate model of decision-making) into ultimately predictable and inherently unpredictable components.

Main text:

The function of perceptual decision-making is to make sense of an uncertain environment whose current state is only partially observable through imperfect sensory measurements. At this 'computational' level of description (Marr 1982), the question of whether human observers process information available in their environment as accurately as possible has been the subject of a large body of work in the recent years. Rahnev & Denison (R&D) make an important case that the labelling of perceptual decisions as suboptimal does not yield much insight regarding the precise nature of suboptimalities. Rather than focusing on which aspects of decision-making are suboptimal in a particular task, R&D propose an alternative roadmap for future research which consists in developing a general 'observer model' of perceptual decision-making across many different tasks.

This proposition is particularly attractive because optimality is often undefined for certain aspects of a task. For instance, human observers carry priors which can be suboptimal for a particular laboratory experiment, but optimal when considering the overall statistics of natural environments (Girshick et al. 2011). Similarly, the cost function assigned to most perceptual decisions is unknown, such that seemingly suboptimal biases in the strictest statistical sense can be seen as optimal in terms of efficient coding (Wei & Stocker 2015). However, the endeavor proposed by R&D is likely to face challenges for which the framework outlined in Box 1 will be of little help.

Perhaps most strikingly, the long list of suboptimalities summarized by R&D in Table 1 is by definition non-exhaustive. Therefore, it remains unknown how much an observer model fails to capture unspecified suboptimalities in any given task. The approach proposed by R&D, which consists in specifying additional forms of suboptimalities and then testing whether they improve model fits, sounds a bit like fumbling in the dark. When will one know that a current ‘observer model’ captures a dominant fraction of suboptimalities? Quality-of-fit metrics are only meaningful in a relative sense - i.e., for comparing candidate models (Palminteri et al. 2017), and they are thus blind to “how wrong” a given model is in an absolute sense.

To address this difficult question, it is important to consider not which aspects of the decision process may be suboptimal, but whether suboptimalities produce random or deterministic variability in behavior - a decomposition known as the ‘bias-variance trade-off’ in statistics. These two forms of suboptimalities map onto the classical distinction between noise and bias - e.g., sensitivity and criterion in signal detection theory (Green and Swets 1966). Independently of any specific theory, the difference between random and deterministic suboptimalities is important in this context, because biases trigger suboptimal decisions which are ultimately predictable, whereas noise triggers suboptimal decisions which are inherently unpredictable. If the long-term goal of R&D’s framework is to predict decision behavior across tasks, then knowing the upper bound on the predictability of decision-making in any given task is indispensable.

While the theoretical distinction between random and deterministic suboptimalities may seem at first abstract and distant from behavior, the two produce antagonistic effects on a simple behavioral metric which can be easily measured in most perceptual tasks: the consistency of decisions across two repetitions of the exact same trial/condition (Wyart and Koechlin 2016). Indeed, deterministic biases tend to increase the consistency of decisions, whereas random noise tends to decrease the same quantity. Therefore, I propose to use decision consistency to decompose suboptimality (or, more generally, deviations from any candidate model of decision-making) into a bias (predictable) term and a variance (unpredictable) term. In practice, the only modification that needs to be made to existing tasks is that the same trial/condition has to be presented at least twice, in order to measure the fraction of repeated trial pairs for which decisions are matched - irrespectively of whether they are correct or not.

In terms of modeling, the approach consists in comparing behavior to simulations of a candidate model of decision-making in terms of decision consistency. If simulated decisions are less consistent across repeated trial pairs than human decisions are, then a fraction of the noise fitted by the model is attributable to unspecified biases - in other words, to unknown sources of suboptimalities that haven’t been captured by the model rather than to true randomness in the decision process. This discrepancy can be quantified in terms of the fraction of random variance in the model that can be pinned down to unknown biases. As an example, we obtained a value of 32% in a canonical probabilistic reasoning task when fitting an optimal model corrupted by noise to human decisions (Drugowitsch et al. 2016). This indicates that about a third of deviations from optimality are due to deterministic, predictable biases. This decomposition of suboptimality into ultimately predictable biases and unpredictable noise can serve not only to measure the effective precision of the decision process in a given task (i.e., the absolute variance of the unpredictable noise term), but also to determine how much a candidate model of decision-making lacks additional, to-be-specified biases.

Like any approach, this bias-variance decomposition of suboptimalities has its limits. First, the bias term will by definition capture only within-trial biases, not sequential biases that propagate across successive trials (Wyart and Koechlin 2016). Sequential biases should thus be specified in the model to be accounted for in the analysis of decision consistency. Second, biases which change over the course of the experiment will spill into the variance term. To control for the existence of such time-varying biases, the experimental design can be made such that the distance between repeated trials is varied across trial pairs (Drugowitsch et al. 2016).

An important corollary of R&D's roadmap is to build an observer model that provides an accurate split between suboptimal biases and true randomness in decision-making. Therefore, analyzing decision consistency should become standard practice to determine how much a candidate model approximates the decision process as good as it possibly can.

Alphabetical reference list:

Drugowitsch, J., Wyart, V., Devauchelle, A.-D., & Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, 92(6), 1398–1411.

Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14(7), 926–932.

Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics. New York: John Wiley.

Marr, D. (1982). Vision: a computational investigation into the human representation and processing of visual information. New York: Freeman.

Palminteri, S., Wyart, V., & Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences*, 21(6), 425–433.

Wei, X.-X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain “anti-Bayesian” percepts. *Nature Neuroscience*, 18(10), 1509–1517.

Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, 11, 109–115.