
Noisy inference in information seeking without prospective reward trade-offs

Yinan Cao^{*1}, Clémence Alméras¹, Junseok K. Lee¹, and Valentin Wyart^{*1}

¹Laboratoire de Neurosciences Cognitives et Computationnelles – DEC, ENS, PSL University – France

Abstract

To make decisions under uncertainty, humans and other animals integrate ambiguous external information through sequential sampling, balancing the competing demands of exploration and exploitation. In laboratory settings, exploration has typically been studied in contexts where agents select goal-directed actions aimed at maximizing prospective rewards. However, not all choices yield immediate rewards—such as browsing restaurant reviews online. The specific role of information seeking when exploration aligns with learning about the structure of the environment, rather than directly maximizing rewards, remains unclear. Here we show that in such situations, humans ($N = 420$) use a two-stage sampling strategy: an initial phase of repeated sampling in temporal “chunks” to generate and test hypotheses about each novel single option, followed by a directed-sampling phase focusing on the more uncertain option. Using computational modeling, we found that suboptimal choices can be explained by imprecision (noise) in the learning process that updates beliefs about option outcomes. This inherent noise makes the initial phase of repeated sampling essential for mitigating the loss of accumulated evidence (compared to a noise-free decision-maker), ultimately improving overall performance. Our results offer a normative justification for this initial repetitive sampling—an apparently suboptimal strategy that does not directly reduce overall uncertainty—under specific biological constraints.

^{*}Speaker