Bayesian Experimental Design with Mutual Information and Learned Errors for Human-Computer Interaction

Hugo Miquel[*,†], Julien Gori[‡], and Olivier Rioul[*]

[*] LTCI Télécom Paris, Institut Polytechnique de Paris, France
[†] ENS Paris-Saclay, Université Paris-Saclay, France
[‡] Sorbonne Université, CNRS, Inserm, Institut des Systèmes Intelligents et de Robotique, ISIR, France

Abstract

This work provides a Bayesian framework for handling user errors in interactive systems, with applications in human-computer interaction (HCI) and user modeling.

The Bayesian Information Gain (BIG) algorithm [1, 2, 3, 4] is an iterative variant of Bayesian experimental design with mutual information as a cost function, used in HCI. It is a principled approach that maximizes expected information gained from each interaction. More precisely, let Θ be the potential target in the user's mind with prior distribution $p(\theta)$, X be the system feedback, and Y be the corresponding user's input. In each interaction loop, BIG selects feedback x that maximizes mutual information $I(\Theta; Y|X = x)$, assuming a known user model (likelihood) $p(y|x, \theta)$, and then updates the posterior distribution $p(\theta|x, y)$.

This work extends the BIG algorithm to learn from user errors while preserving its mathematical foundations. We incorporate an error rate parameter ϵ into the likelihood function $p(y|x, \theta, \epsilon)$ and develop an adaptive algorithm that jointly infers both θ and ϵ by updating the posterior $p(\theta, \epsilon|x, y)$ at each interaction step. We also discuss three simplifying hypotheses for the prior expression $p(\theta, \epsilon)$ and three user models: (i) zero error; (ii) fixed error rate; (iii) arbitrary random error rate. We prove mathematical continuity between these three models, showing that our adaptive approach naturally extends BIG.

We also investigate model mismatch on the overall performance and degradation properties with respect to the standard BIG algorithm. While standard BIG converges quickly with perfect responses, it degrades with even small error rates. The fixed-error model depends critically on correctly estimating the error parameter, while our adaptive model achieves the highest accuracy under varying error conditions, at the expense of additional interactions.

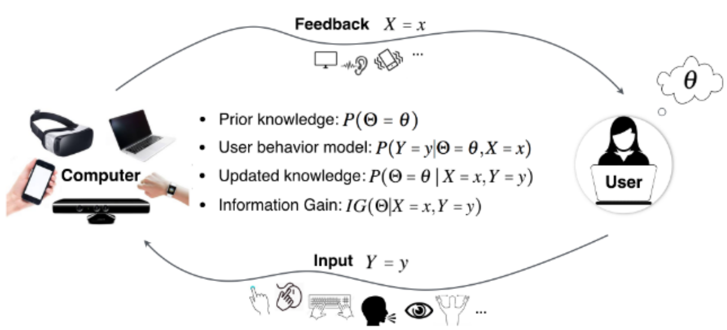

Figure 1: The BIG framework

References

[1] Wanyu Liu, Rafael Lucas d'Oliveira, Michel Beaudouin-Lafon, and Olivier Rioul. Bignav: Bayesian information gain for guiding multiscale navigation. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 5869–5880, 2017.

[2] Wanyu Liu, Olivier Rioul, Joanna Mcgrenere, Wendy E Mackay, and Michel Beaudouin-Lafon. Bigfile: Bayesian information gain for fast file retrieval. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018.

[3] Wanyu Liu, Olivier Rioul, and Michel Beaudouin-Lafon. Bayesian information gain to design interaction. *Bayesian Methods for Interaction Design*, 2022.

[4] Seung Won Lee, Hwan Kim, Taeha Yi, and Kyung Hoon Hyun. Bigaze: an eye-gaze action-guided bayesian information gain framework for information exploration. *Advanced Engineering Informatics*, 58:102159, 2023.