Interface Tutoring

Context of the project Software companies invest a lot of effort into designing new features. However, most users, professionals included, exploit an underwhelming amount of these features.For example, the AutoCAD software offers more than 2000 commands, yet the typical user functions with about 30-40 commands [13]. As a result, the exploitation of software on the whole is suboptimal: few software users reach the level of expertise that could be expected, even after an extended use [6], with the majority of users essentially failing to learn the interfaces they interact with. Making users transition towards expert behavior is thus crucial.

This transition has been studied in the field of Human Computer Interaction (HCI) primarily from a design and empirical perspective [7], and remains largely an open problem. In this work, we aim at facilitating this transition using tools from AI.

Objectives Most users operate in a goal-driven context which requires short-term efficiency. This leaves little space open for exploration and/or learning of the interface. The main idea of this work is to teach users about interfaces in the least obtrusive manner possible, so that users learn about the interface *incidentally*. In contrast with most existing works in HCI, we adopt a computational and modelling perspective, by using an Intelligent Tutoring System (ITS) that leverages tools and methods of AI. The goals of this thesis are:

- understanding and quantifying how users perceive teaching interventions, as well as perceived utilities to transitioning to expert modes of interaction;
- designing an ITS to teach interfaces, via model-based planning that trades off a model of retention (how users learn and forget) with estimates of utilities;
- transferring results from extrinsic and intrinsic motivation theories for incidental learning of interfaces.

Positioning in relation to the state of the art Intelligent Tutoring Systems have become quite popular — for example a famous foreign language ITS totals more than 100M+ downloads on a major application store. ITS consist of three modules: a domain model (which amounts to expert knowledge), a student model, maintained by the system about the knowledge, abilities and/or particularities of the student, and a tutoring model, which based on the domain and student model and other extrinsic information (such as the date of an upcoming evaluation), plans the tutoring strategy [2]. Our work does not tackle the domain model; in the context of interfaces that model is usually trivial (the command map is documented). Instead, our work focuses on the *user and tutoring models* for the specific problem of *teaching interfaces.*

Student Model Student's cognitive and/or affective states can be inferred by maintaining parametrized user models, where the parameters are identified from observed behavior. For example, exponential memory models have been leveraged and their parameters updated with bayesian belief updating [14] to predict the effect of scheduled learning with repetitions. Computational approaches to (intrinsic) motivation [15] of students can also be leveraged [5] to increase the engagement of students with the learning material. We will transpose these methods to the context of teaching interfaces, where specific user models will have to be developed.

Tutoring models While many strategies exist Several decision-making models, such as Multi-armed bandits [5] and Partially Observable Markov Decision Processes (POMDP) [18, 14] have been recently leveraged to plan optimal recommendations for students. These models are particularly relevant because they express utilities in sequential decision-making problems, which are exactly at stake in this thesis. Approximate solutions to the POMDP can be found via Deep Reinforcement Learning [12, 19]).

Impact A large and increasing amount of work is performed today via computers. Even a slight increase in efficiency will generate large savings, and therefore has the potential for a large economic impact. From a scientific perspective, this context of "learning while working" constitutes an interesting edge case to validate theories of motivation. We also believe that intended work on the perceived utility of the interventions can transfer to other AI-infused systems, such as recommender systems.

Originality Perceived cost of interruptions, and the design and admissability of various interventions have been studied in the HCI literature, but mostly from a qualitative point of view[8], and various heuristics exist to select interventions. Here we propose a single framework to tradeoff the objectives of the system (teaching) with its acceptability from a computational perspective. Computational modeling in HCI is an emerging topic [16], with very few representatives in France, among which two of the thesis supervisors Julien Gori and Gilles Bailly.

Challenges The first challenge of this thesis is to apply methods of ITS to the target population of interface tutoring which, contrary to usual students, is not actively engaged in learning: its primary goal is to accomplish a task rather than learning about an interface. We postulate that while frequent interventions of ITS's are prohibitive from a user perspective in this context [11], punctual interventions are not, and may be sufficient to elicit incidental learning of interfaces. Timing these punctual interventions is crucial for learning from a memory/retention perspective. There is an evident tradeoff between frequent interventions that favor learning but make the ITS less acceptable to a user. This tradeoff will be resolved by solving a POMDP formulation for ITS [18] that will have been modified and extended to encompass utilities inherent to the teaching interventions. Transferring theories of extrinsic and intrinsic motivation [15] and determining how they can drive design of an ITS to teach interfaces will also be investigated.

The second challenge is that, to our knowledge, the transition from novice to expert behavior is not a problem in the typical teaching setting for which ITS have been developed. As a result, our work will further consider models where users switch from one interaction mode to another, which we have already started developing [3].

Implemention Methods The thesis will be primarily conducted in the HCI field.

Year 1 a literature review that examines which current strategies and designs for ITS, and user models, can be transposed for incidental learning, especially focusing on models of decision-making and exploration.

Year 1 Conceiving an experiment to quantify perceived user costs to switching and teaching interventions.

Year 1/2 the formulation of the ITS providing punctual observations as a POMDP, trading off memory models with ITS intervention utilities.

Year 2/3 Empirical validation of the models and the ITS.

Year 2/3 Exploiting results of theories of intrinsic and extrinsic motivation to generate new designs for the ITS.

Year 2/3 A working prototype of the proposed ITS for evaluation in strictly controlled, but also more ecologically valid conditions.

Supervision The PhD will be supervised by Gilles Bailly, Mehdi Khamassi and Julien Gori (all 3 at CNRS and Sorbonne Université) and will be hosted in the ISIR Multi-scale interaction team (ISIR, Sorbonne Université).

Gilles Bailly is a research scientist with a strong track record at the CHI and UIST conferences (the two main avenues for publication in the field of HCI;14 papers these last five years and 6 awards). He has demonstrated an expertise in expert-to-novice transition [3]. He has used methods from AI to build models of users performance [4]. He was co-organizer of the French meeting in HCI-AI (∼100 attendees) in 2016 and 2018 respectively (`https://ihmia.afihm.org`)

Mehdi Khamassi is a research director, whose work lies at the intersection between Cognitive Science, Neuroscience, AI and Robotics, with an emphasis on various aspects of decision-making and reinforcement learning [1, 17]. He is an associate editor for several journals, including Frontiers in Decision Neuroscience.

Julien Gori is a research scientist since 2021, with an expertise in computational models in HCI [10]. Before that, he was a post-doctoral fellow at the Finnish Center for Artificial Intelligence. He has published several works at top HCI venues such as CHI and UIST.. He is the developer and maintainer of a library that facilitates computational user modeling in HCI, with an emphasis on Deep Reinforcement Learning [9].

Candidate Profile Applicants with a strong academic record in HCI or a field related to AI are encouraged to apply. Interest and/or experience in user modeling and/or models of decision-making is appreciated.

References

[1] Frédéric Alexandre, Peter F Dominey, Philippe Gaussier, Benoît Girard, **Mehdi Khamassi**, and Nicolas P Rougier. When artificial intelligence and computational neuroscience meet. In *A Guided Tour of Artificial Intelligence Research*, pages 303–335. Springer, 2020.

[2] Ali Alkhatlan and Jugal Kalita. Intelligent tutoring systems: A comprehensive historical survey with recent developments. *arXiv preprint arXiv:1812.09628*, 2018.

[3] **Gilles Bailly**, **Mehdi Khamassi**, and Benoît Girard. Computational model of the transition from novice to expert interaction techniques. *ACM Transactions on Computer-Human Interaction*, 2022.

[4] Xiuli Chen, **Gilles Bailly**, Duncan P Brumby, Antti Oulasvirta, and Andrew Howes. The emergence of interactive behavior: A model of rational menu search. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 4217–4226, 2015.

[5] Benjamin Clement, Didier Roy, Pierre-Yves Oudeyer, and Manuel Lopes. Developmental learning for intelligent tutoring systems. In *4th International Conference on Development and Learning and on Epigenetic Robotics*, pages 426–433. IEEE, 2014.

[6] Andy Cockburn, Carl Gutwin, Joey Scarr, and Sylvain Malacria. Supporting novice to expert transitions in user interfaces. *ACM Computing Surveys (CSUR)*, 47(2):1–36, 2014.

[7] Emmanouil Giannisakis, **Gilles Bailly**, Sylvain Malacria, and Fanny Chevalier. Iconhk: Using toolbar button icons to communicate keyboard shortcuts. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 4715–4726, 2017.

[8] Sonja Gievska, Rob Lindeman, and John Sibert. Examining the qualitative gains of mediating human interruptions during hci. In *Proceedings 11th International Conference on Human-Computer Interaction, Las Vegas, Nevada*, 2005.

[9] **Julien Gori**. `https://jgori-ouistiti.github.io/CoopIHC/`.

[10] **Julien Gori**, Olivier Rioul, and Yves Guiard. Speed-accuracy tradeoff: A formal information-theoretic transmission scheme (fitts). *ACM Transactions on Computer-Human Interaction (TOCHI)*, 25(5):1–33, 2018.

[11] Tovi Grossman, Pierre Dragicevic, and Ravin Balakrishnan. Strategies for accelerating on-line learning of hotkeys. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1591–1600, 2007.

[12] Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. In *2015 aaai fall symposium series*, 2015.

[13] Justin Matejka, Wei Li, Tovi Grossman, and George Fitzmaurice. Communitycommands: command recommendations for software applications. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 193–202, 2009.

[14] Aurélien Nioche, Pierre-Alexandre Murena, Carlos de la Torre-Ortiz, and Antti Oulasvirta. Improving artificial teachers by considering how people learn and forget. In *26th International Conference on Intelligent User Interfaces*, pages 445–453, 2021.

[15] Pierre-Yves Oudeyer and Frederic Kaplan. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6, 2009.

[16] Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes. *Computational interaction.* Oxford University Press, 2018.

[17] Stefano Palminteri, **Mehdi Khamassi**, Mateus Joffily, and Giorgio Coricelli. Contextual modulation of value signals in reward and punishment learning. *Nature communications*, 6(1):1–14, 2015.

[18] Anna N Rafferty, Emma Brunskill, Thomas L Griffiths, and Patrick Shafto. Faster teaching via pomdp planning. *Cognitive science*, 40(6):1290–1332, 2016.

[19] Abolfazl Zaraki, **Mehdi Khamassi**, Luke J Wood, Gabriella Lakatos, Costas Tzafestas, Farshid Amirabdollahian, Ben Robins, and Kerstin Dautenhahn. A novel reinforcement-based paradigm for children to teach the humanoid kaspar robot. *International Journal of Social Robotics*, 12(3):709–720, 2020.