Computational Workflows for Designing Input Devices

YI-CHI LIAO, Aalto University, Finland

Input devices, such as buttons and sliders, are the foundation of any interface. The typical user-centered design workflow requires the developers and users to go through many iterations of design, implementation, and analysis. The procedure is inefficient, and human decisions highly bias the results. While computational methods are used to assist various design tasks, there has not been any holistic approach to automate the design of input components. My thesis proposed a series of *Computational Input Design* workflows: I envision a sample-efficient multi-objective optimization algorithm that cleverly selects design instances, which are instantly deployed on physical simulators. A meta-reinforcement learning user model then simulates the user behaviors when using the design instance upon the simulators. The new workflows derive Pareto-optimal designs with high efficiency and automation. I demonstrate designing a push-button via the proposed methods. The resulting designs outperform the known baselines. The Computational Input Design process can be generalized to other devices, such as joystick, touchscreen, mouse, controller, etc.

1 INTRODUCTION

Input devices, such as buttons, joysticks, sliders, and knobs, are ubiquitous, and they are the fundamental components of any user interface. Many research has shown that the design of the input devices highly affects user behaviors, performances, and experiences [13, 34]. For example, the design of push-buttons results in distinct typing speed and comforts [36], and the transfer functions of a mouse lead to different pointing efficiency [4]. However, making a good design of an input device is extremely challenging. The standard workflow is so-called *User-Centered Design (UCD)*, which usually consists of *design, Implement, and analyze* phases [1]. This process is inefficient and costly. The designers and developers have to craft several prototypes in order to test them with the users. User researchers then need to plan and conduct controlled experiments. The experiment results require interpretations so that the designers can make improvements for the next iteration. Every step is potentially biased by human decisions. For instance, the selection of test sets mostly relies on the designers' experience. Because of the high cost, a complete user-centered design process is hardly conducted. Unfortunately, sometimes designs are fixed based on what feels right to the designers. This is even more prone to miss good designs.

While the same difficulties exist for software-based interfaces, such as menu and layout designs, there are many computational tools for automating the design process [2, 45]. *User models* can be applied for simulating user performances

Design

Sample
w. MOBO

UCD

CID

Analyze

Implement

Evaluate
w. Meta-RL

Physical
Simulator

Fig. 1. Traditional User-Centered Design (UCD) workflow is inefficient, and the resulting designs are highly biased by the designers' decisions. My thesis proposes *Computational Input Design (CID)*, which contains several components: A multi-objective Bayesian optimization (MOBO) algorithm efficiently searches for Pareto-optimal designs across the design space; A physical simulator instantly renders the design selected by the optimizer; A meta-reinforcement learning (meta-RL) user model adapts to the design instance within a handful of trials and provides performance metrics for evaluating this instance. Jointly, the CID workflows automatically derive Pareto-optimal design instances.

[23, 32], that further allow *optimization methods* to derive the optimal designs in a *simulation environment* [29, 46]. However, those computational resources are usually not suitable for input device designs. First, the construction of user models for physical interaction is generally a challenging task, and those models are typically task-specific. Secondly, few works have attempted to create input-device simulators that allow realistic rendering of arbitrary designs. Thus, only a few papers addressed using optimization methods to automatically search for better input interface design.

I propose *Computational Input Design (CID)* workflows to provide high automation, efficiency, and robust results to the designers (Figure 1). Three innovations should be implemented to complete CID: (1) sample-efficient *multi-objective Bayesian optimization (MOBO)* algorithms to efficiently search for Pareto-optimal designs, (2) *physical simulators* for deploying designs quickly without fabricating and assembling an entire physical interface, and (3) *meta-reinforcement learning (meta-RL) user models* to simulate human's ability to quickly adapt to new devices and design instances.

In my thesis, I demonstrate designing a push-button upon the proposed framework. Buttons are transducers that register a discrete event from physical motion [14, 18, 33], and arguably the most basic input component for any interfaces. Interestingly, each button design is unique in its haptic response characteristics. Gamers, programmers, and hobbyist groups have a keen interest in tactility, which is associated with sensory experience and performances. The design of push-button is a costly process where the developers have to craft physical prototypes for further conducting studies. The CID process can greatly reduce the cost. In my previous works [20–22], I introduced a novel FDVV models to parameterize button design. A sophisticated physical simulator is introduced to render arbitrary tactility according to the given design parameters. A MOBO algorithm is applied to efficiently explore the design space and search the Pareto-optimal designs. The user model based on meta-RL algorithms is trained on various example buttons that allow fast adaptation to new design instances.

1.1 Contribution

I have introduced a novel FDVV button model allowing realistic button simulation and interactive design [20–22]. As shown in Figure 3, an end-to-end simulation pipeline covers capturing the button tactility, modeling, controlling, and rendering. The guidelines of this work can go beyond buttons to general haptic-rendering devices. The other

paper, MOBO for user interface design, is under submission (an overview of the algorithm is illustrated in Figure 2). In this work, I introduced a novel workflow allowing designers to efficiently identify Pareto-optimal design candidates using Bayesian optimization. I demonstrated and assessed the MOBO workflow across three representative interface design problems and the resulting designs outperform all the baselines. The final component of my dissertation is a meta-RL-based model able to quickly adapt to various input components, including the unseen design instances. This paper will close the loop of *Computational Input Design* process. The contributions of my thesis are as follow:

- Proposing *Computational Input Design* workflows and identifying the major components to realize the concept, which further allows a highly efficient and robust input design process;
- Three novel implementations: the multi-objective Bayesian optimizer, the physical simulator, and the meta-RL algorithm, can be useful for other general design tasks;
- Demonstrating designing a button via the CID workflows.

1.2 Background

I am a third-year Ph.D. student in the School of Electrical Engineering, Aalto University, working on Computational Interaction, advised by Professor Antti Oulasvirta. My research focus is building input devices and touch interfaces utilizing computational methods, such as modeling, Bayesian optimization, reinforcement learning, and meta-learning. My ultimate goal is to bring HCI one step closer toward a principled design process and optimized physical interactions.

2 KEY RELATED WORKS

Computational interaction is one branch of the modern HCI field where researchers strive for constructing user models and applying optimization methods to automatically find optimal designs [30]. Here, we review the development of computation interaction and its limitation on input devices. Then, we discuss the methods applied in my thesis.

2.1 Computational Methods on Input Device Design

In contrast to the traditional User-Centered Design research process[1], computational interaction aims for improvements in modeling, optimization, and inference. Better user models allow more precise prediction of user behavior and performances [23]. Promising optimization methods enable faster search for the optimal designs [31]. Those computational tools have been applied on various kinds of interfaces, such as sketch, menu, and layout design, visualization of scatter plots, etc. [2, 25, 45]. However, only a few computational methods are applied for assisting input component design. Some notable exceptions include tuning the gain function [19] and investigating the optimal design of mouse sensor position [15]. These works only applied optimization algorithms for selecting parameters, but no user models nor configurable physical components were involved. My thesis is, for the first time, proposing a holistic Computational Input Design framework in which models directly interact with the optimizer and the physical simulator.

2.2 Bayesian Optimization

A part of our paper is built upon Bayesian optimization, a machine learning method that aims to efficiently explore a black-box function and further identify an optimal point. The approach is suitable for applications in which the functions are expensive to evaluate in terms of the required time or effort [42]. Bayesian optimization has been applied in games to assign parameter values to maximize the user engagement [12]. Conducting crowd-sourcing Bayesian optimization studies is another approach to minimize the required time period [7, 12]. Brochu et al. [3] demonstrated a technique for

allowing designers to quickly determine appropriate values for animation rendering. Koyama et al. [16, 17] attempted to assist users on editing of photographs to achieve a promising visual appearance. Bayesian optimization has also been used for customizing interface settings for individuals [28, 43]. However, the application of Bayesian optimization to HCI design problems has largely been limited to the optimization for tasks with only a single objective. In practice, most interface design problems are characterized by a complex interplay between multiple, often competing, objectives, e.g., speed versus accuracy. In my thesis, I leverage the MOBO algorithms introduced by Shah and Ghahramani [40] and demonstrate its applicability to input device design process.

2.3 Physical Simulator and Haptic Rendering

A part of my thesis is built upon physical simulators that aligned with haptics research pursuing the creation of rich and realistic sensations [11]. Research has looked at advanced factors affecting haptic perception, such as friction and texture [9]. Force simulators have been introduced, for example, the Phantom device is a 6-DOF pen-type force-rendering device capable of emulating the softness of deformable objects [24]. However, a low operating rate (60 Hz), excessive degrees of freedom, and a lack of vibrotactile stimulus limit its use to simulate input devices accurately. Doerrer and Werthschuetzky enabled users to edit the force-displacement profile of a push-button in software [5]. Nonetheless, the force-displacement model is known to be incomplete, and there are no experimental results to validate this approach. Hence, there is a need to explore more effective methods for constructing physical simulator and haptic rendering.

2.4 Meta-Reinforcement Learning Algorithms

Lastly, the user model in my thesis is largely based on meta-reinforcement learning (meta-RL) algorithms. Reinforcement learning (RL) is an area of machine learning concerned with how software agents take actions in an environment in order to maximize the cumulative reward function Sutton and Barto [44]. Recently, various approaches have been introduced for reinforcement learning with neural network function approximators. The well-known algorithms are deep Q-learning [26], trust region policy gradient methods [38], proximal policy optimization [39], and soft actor-critic [10]. On the other hand, Meta-Learning concerns the question of learning to learn. The goal of meta-learning is to train a model that can quickly adapt to a new task using only a few data points and small training iterations [6, 8]. The model is previously trained during a meta-learning phase on a set of similar tasks, such that the trained model can quickly adapt to new unseen tasks. There are several variation of meta-RL implementations, such as LSTM meta-learner [35] and gradient-based model-agnostic meta-learners [8, 27]. A recent development is based on gradient-based meta-learner combined policy-gradient methods shows a good fit for continuous state, action environments [37]. I plan to follow this approach to build a meta-RL model to simulate continuous user's hand motions.

3 RESEARCH PROGRESS

My current publications have contributed to the vision towards state-of-the-art MOBO for interaction design and realistic button simulation.

3.1 Multi-Objective Bayesian Optimization for Interface Design

I introduced Bayesian optimization as an efficient solution to multi-objective design challenges related to interaction techniques (Figure 2). Specifically, I described and demonstrated an interactive workflow for identifying *Pareto-optimal design candidates* for HCI design tasks using Bayesian optimization. I extended *Bayesian optimization* [41], a powerful approach for efficient identification of optimal designs in studies involving human data. The value of the approach lies

Human-in-the-loop optimization
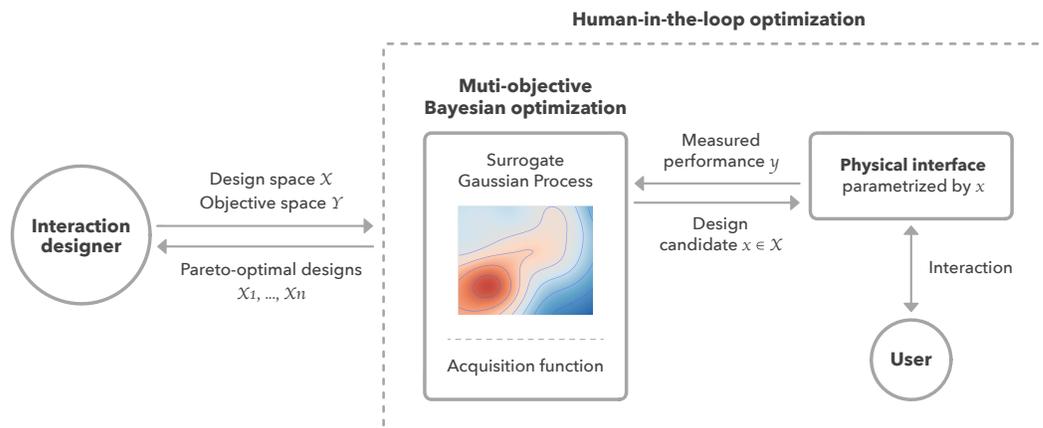
Fig. 2. Multi-objective Bayesian optimization (MOBO) is a novel workflow for efficiently identifying a set of informative designs. The designer defines design parameter space (X) and objective function space (Y). The optimizer will generate design candidates $x \in X$ to search for the Pareto-optimal points. The interface will render according to each x. The user's interaction will then be translated into objective values y. The optimizer updates its proxy model, a Gaussian Process model defined by the observed $\{x, y\}$ sets. Then, the proxy model generates a new design candidate, and sends it to the interface.

in its efficient exploration of the design space, with the technique particularly well suited to black-box optimization with noisy observations and its evaluation is expensive. An intuitive way to think about it is that it can help the designers to *avoid* testing design instances that would be uninformative, for instance, parameters associated with poor usability or some measurable objectives. I demonstrated using the MOBO workflow for tackling representative input device design problems. In all case studies, the results indicated that MOBO design approach enables the efficient identification of promising design candidates.

3.2 Button Simulation and Design

My previous publication investigated the simulation and interactive design of push-buttons [20–22]. Accurate simulation of the button-pressing experience is challenging. While pre-existing simulators can render tactile and linear buttons, I found that the typical force–displacement (FD) model cannot accurately render button tactility. Also, no methods have been offered to help designers and engineers exploit such simulations. To address these challenges, I proposed an extended model and an end-to-end simulation pipeline around it (Figure 3). This approach allows simulating more button types than previously, including tactile-type buttons and buttons with different click reactions and various travel ranges. Furthermore, it permits the analysis and editing of buttons. Our work centers on the *Force–Displacement–Vibration–Velocity* (FDVV) model and an end-to-end simulation pipeline. The model adds vibration response and velocity-dependence on top of the FD model. In our implementation, vibration is sampled through a microphone during a button press, and multiple FD curves are sampled at several speeds. I solved several engineering challenges connected with ambitions to capture and simulate buttons via FDVV models. For rendering it, I presented a novel simulator construction for FDVV models. This simulator is capable of detecting displacement to μm precision at a high sampling rate (1 kHz) and can produce a wide range of force (up to 4.4 N) and vibration (50 Hz – 20 kHz) feedback.
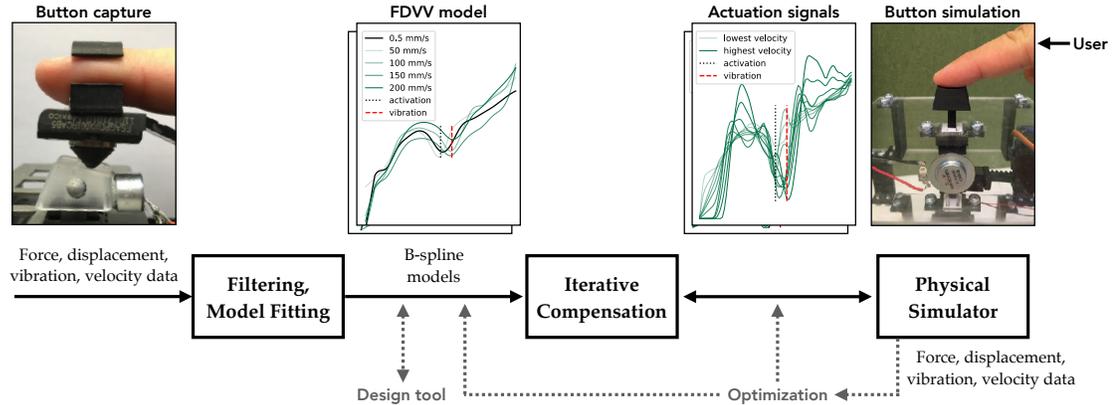
Fig. 3. An overview of my previous works on button simulation and design [20–22]. In this series of works, I introduced an end-to-end approach for button simulation and design. To capture an FDVV model of a button, sensors are placed on the finger, and the button is pressed multiple times. The resulting force, displacement, vibration, and velocity data are filtered and fitted with the lower-parametric B-splines models based on the Bayesian Information Criteria (BIC) values. A designer can edit the model produced. Before rendering the FDVV model on our simulator, an iterative compensation process computes how to cancel the simulator's transfer function. The resulting actuation signals drive the simulator.

4 FUTURE DIRECTIONS

As part of future work, I am working on a general user model that can learn policies for interacting with physical devices and quickly adapt to new design instances. Such a user model can be used for simulating user behavior and thus reduce the efforts of conducting rounds of user studies. For modeling interactions with input devices, such as pressing a button, reinforcement learning is so far the most general and prominent framework [44]. Meta-RL further provides a means that allows the model to quickly identify an optimal policy for a new task by leveraging the previous learning experiences. The meta-RL models bring us one step closer to the human-like behaviors when encountering a new input device. While various meta-RL algorithms have been introduced, a particular algorithm may be suited to the needs of modeling continuous human hand motions, which is based on gradient-based meta-learning and policy-gradient methods [37]. I intend to follow this approach to build a general-purpose user model. The goals of this paper are: (1) constructing a general-purpose model for the physical interaction of wide-ranging input devices, and (2) by merging the meta-RL model into other proposed workflows, we can learn the efficacy of a fully automated *computational input design* process. The first design task is deriving a push-button via the CID workflows: The physical button simulator renders button design based on the MOBO's suggestion (as mentioned in section 3), and a meta-RL model applied on a robot will act as various users to interact with the button simulator. Jointly, a series of Pareto-optimal button designs will be derived automatically. Other input devices, such as joystick and mouse, can potentially be designed by the proposed CID process.

5 ACKNOWLEDGEMENT

REFERENCES

[1] Chadia Abras, Diane Maloney-Krichmar, Jenny Preece, et al. 2004. User-centered design. *Bainbridge, W. Encyclopedia of Human-Computer Interaction. Thousand Oaks: Sage Publications* 37, 4 (2004), 445–456.

[2] Gilles Bailly, Antti Oulasvirta, Timo Kötzing, and Sabrina Hoppe. 2013. MenuOptimizer: Interactive Optimization of Menu Systems. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) *(UIST '13)*. Association for Computing Machinery, New York, NY, USA, 331–342. https://doi.org/10.1145/2501988.2502024

[3] Eric Brochu, Tyson Brochu, and Nando de Freitas. 2010. A Bayesian interactive optimization approach to procedural animation design. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '10)*. Eurographics Association, Goslar, DEU, 103–112.

[4] Géry Casiez, Daniel Vogel, Ravin Balakrishnan, and Andy Cockburn. 2008. The Impact of Control-Display Gain on User Performance in Pointing Tasks. *Human–Computer Interaction* 23, 3 (2008), 215–250. https://doi.org/10.1080/07370020802278163

[5] C. Doerrer and R. Werthschuetzky. 2002. Simulating Push-Buttons Using a Haptic Display: Requirements on Force Resolution and Force-Displacement Curve.

[6] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL2: Fast reinforcement learning via slow reinforcement learning. (2016). arXiv:1611.02779

[7] John J. Dudley, Jason T. Jacques, and Per Ola Kristensson. 2019. Crowdsourcing Interface Feature Design with Bayesian Optimization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300482

[8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. arXiv:1703.03400 [cs.LG]

[9] K. FUJITA. 2001. A New Softness Display Interface by Dynamic Fingertip Contact Area Control. *5th World Multiconference on Systemics, Cybernetics and Informatics, 2001* (2001), 78–82. https://ci.nii.ac.jp/naid/10031028472/en/

[10] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. arXiv:1801.01290 [cs.LG]

[11] Krueger L. (Ed.) Katz D. 1989. *The World of Touch*. New York: Psychology Press. https://doi.org/10.4324/9780203771976

[12] Mohammad M. Khajah, Brett D. Roads, Robert V. Lindsey, Yun-En Liu, and Michael C. Mozer. 2016. Designing Engaging Games Using Bayesian Optimization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 5571–5582. https://doi.org/10.1145/2858036.2858253

[13] Jeong Ho Kim, Lovenoor S. Aulck, Michael C. Bartha, Christy A. Harper, and Peter W. Johnson. 2014. Differences in typing forces, muscle activity, comfort, and typing performance among virtual, notebook, and desktop keyboards. *Applied ergonomics* 45 6 (2014), 1406–13.

[14] Sunjun Kim, Byungjoo Lee, and Antti Oulasvirta. 2018. Impact Activation Improves Rapid Button Pressing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, Article 571, 8 pages. https://doi.org/10.1145/3173574.3174145

[15] Sunjun Kim, Byungjoo Lee, Thomas van Gemert, and Antti Oulasvirta. 2020. Optimal Sensor Position for a Computer Mouse. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376735

[16] Yuki Koyama, Issei Sato, and Masataka Goto. 2020. Sequential gallery for interactive visual design optimization. *ACM Transactions on Graphics* 39, 4 (July 2020), 88:88:1–88:88:12. https://doi.org/10.1145/3386569.3392444

[17] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential line search for efficient visual design optimization by crowds. *ACM Transactions on Graphics* 36, 4 (July 2017), 48:1–48:11. https://doi.org/10.1145/3072959.3073598

[18] Byungjoo Lee, Sunjun Kim, Antti Oulasvirta, Jong-In Lee, and Eunji Park. 2018. Moving Target Selection: A Cue Integration Model. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, Article 230, 12 pages. https://doi.org/10.1145/3173574.3173804

[19] Byungjoo Lee, Mathieu Nancel, Sunjun Kim, and Antti Oulasvirta. 2020. AutoGain: Gain Function Adaptation with Submovement Efficiency Optimization *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376244

[20] Yi-Chi Liao, Sunjun Kim, Byungjoo Lee, and Antti Oulasvirta. 2020. Button Simulation and Design via FDVV Models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376262

[21] Yi-Chi Liao, Sunjun Kim, Byungjoo Lee, and Antti Oulasvirta. 2020. Press'Em: Simulating Varying Button Tactility via FDVV Models. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–4. https://doi.org/10.1145/3334480.3383161

[22] Yi-Chi Liao, Sunjun Kim, and Antti Oulasvirta. 2018. One Button to Rule Them All: Rendering Arbitrary Force-Displacement Curves. In *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings* (Berlin, Germany) *(UIST '18 Adjunct)*. ACM, New York, NY, USA, 111–113. https://doi.org/10.1145/3266037.3266118

[23] I. Scott MacKenzie. 1992. Fitts' Law as a Research and Design Tool in Human-Computer Interaction. *Human–Computer Interaction* 7, 1 (1992), 91–139. https://doi.org/10.1207/s15327051hci0701_3

[24] Thomas H. Massie and J. K. Salisbury. 1994. The PHANToM haptic interface: A device for probing virtual objects. In *Proceedings of the ASME Dynamic Systems and Control Division*. 295–301.

[25] L. Micallef, G. Palmas, A. Oulasvirta, and T. Weinkauf. 2017. Towards Perceptual Optimization of the Visual Design of Scatterplots. *IEEE Transactions on Visualization and Computer Graphics* 23, 6 (2017), 1588–1599.

[26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.

[27] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On First-Order Meta-Learning Algorithms. arXiv:1803.02999 [cs.LG]

[28] Jens Brehm Bagger Nielsen, Jakob Nielsen, and Jan Larsen. 2015. Perception-Based Personalization of Hearing Aids Using Gaussian Processes and Active Learning. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 23, 1 (Jan. 2015), 162–173. https://doi.org/10.1109/TASLP.2014.2377581 Conference Name: IEEE/ACM Transactions on Audio, Speech, and Language Processing.

[29] A. Oulasvirta. 2017. User Interface Design with Combinatorial Optimization. *Computer* 50, 1 (2017), 40–47.

[30] A. Oulasvirta. 2018. *Computational Interaction*. Oxford University Press. https://books.google.fi/books?id=-r9EDwAAQBAJ

[31] Antti Oulasvirta, Niraj Ramesh Dayama, Morteza Shiripour, Maximilian John, and Andreas Karrenbauer. 2020. Combinatorial optimization of graphical user interface designs. *Proc. IEEE* 108, 3 (2020), 434–464.

[32] Antti Oulasvirta, Samuli De Pascale, Janin Koch, Thomas Langerak, Jussi Jokinen, Kashyap Todi, Markku Laine, Manoj Kristhombuge, Yuxi Zhu, Aliaksei Miniukovich, Gregorio Palmas, and Tino Weinkauf. 2018. Aalto Interface Metrics (AIM): A Service and Codebase for Computational GUI Evaluation. In *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings* (Berlin, Germany) *(UIST '18 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 16–19. https://doi.org/10.1145/3266037.3266087

[33] Antti Oulasvirta, Sunjun Kim, and Byungjoo Lee. 2018. Neuromechanics of a Button Press. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, Article 508, 13 pages. https://doi.org/10.1145/3173574.3174082

[34] Robert G. Radwin and One-Jang Jeng. 1997. Activation Force and Travel effects on Overexertion in Repetitive Key Tapping. *Human Factors* 39, 1 (1997), 130–140. https://doi.org/10.1518/001872097778940605 PMID: 9302885.

[35] Sachin Ravi and Hugo Larochelle. 2016. Optimization as a model for few-shot learning. (2016).

[36] David Rempel, Elaine Serina, Edward Klinenberg, Bernard J. Martin, Thomas J. Armstrong, James A. Foulke, and Sivakumaran Natarajan. 1997. The effect of keyboard keyswitch make force on applied force and finger flexor muscle activity. *Ergonomics* 40, 8 (1997), 800–808. https://doi.org/10.1080/001401397187793 PMID: 9336104.

[37] Jonas Rothfuss, Dennis Lee, Ignasi Clavera, Tamim Asfour, and Pieter Abbeel. 2018. Promp: Proximal meta-policy search. (2018). arXiv:1810.06784

[38] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. 2017. Trust Region Policy Optimization. arXiv:1502.05477 [cs.LG]

[39] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347 [cs.LG]

[40] Amar Shah and Zoubin Ghahramani. 2016. Pareto Frontier Learning with Expensive Correlated Objectives. In *International Conference on Machine Learning*. PMLR, 1919–1927. http://proceedings.mlr.press/v48/shahc16.html ISSN: 1938-7228.

[41] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* 104, 1 (2015), 148–175.

[42] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* 104, 1 (Jan. 2016), 148–175. https://doi.org/10.1109/JPROC.2015.2494218 Conference Name: Proceedings of the IEEE.

[43] Jasper Snoek. 2013. *Bayesian Optimization and Semiparametric Models with Applications to Assistive Technology*. Ph.D. Dissertation. University of Toronto, Toronto, Ontario, Canada.

[44] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.

[45] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and Explore with a Layout Optimiser. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) *(DIS '16)*. Association for Computing Machinery, New York, NY, USA, 543–555. https://doi.org/10.1145/2901790.2901817

[46] Shumin Zhai, Michael Hunter, and Barton A. Smith. 2002. Performance Optimization of Virtual Keyboards. *Human–Computer Interaction* 17, 2-3 (2002), 229–269. https://doi.org/10.1080/07370024.2002.9667315