Introduction to Intelligent User Interfaces

Albrecht Schmidt
LMU Munich
Munich, Germany
albrecht.schmidt@ifi.lmu.de

Sven Mayer
LMU Munich
Munich, Germany
info@sven-mayer.com

Daniel Buschek
daniel.buschek@uni-bayreuth.de
Research Group HCI + AI,
Department of Computer Science,
University of Bayreuth
Bayreuth, Germany

ABSTRACT

Recent advancements in artificial intelligence (AI) create new opportunities for implementing a wide range of intelligent user interfaces. Speech-based interfaces, chatbots, visual recognition of users and objects, recommender systems, and adaptive user interfaces are examples that have majored over the last 10 years due to new approaches in machine learning (ML). Modern ML-techniques outperform in many domains of previous approaches and enable new applications. Today, it is possible to run models efficiently on various devices, including PCs, smartphones, and embedded systems. Leveraging the potential of artificial intelligence and combining them with human-computer interaction approaches allows developing intelligent user interfaces supporting users better than ever before. This course introduces participants to terms and concepts relevant in AI and ML. Using examples and application scenarios, we practically show how intelligent user interfaces can be designed and implemented. In particular, we look at how to create optimized keyboards, use natural language processing for text and speech-based interaction, and how to implement a recommender system for movies. Thus, this course aims to introduce participants to a set of machine learning tools that will enable them to build their own intelligent user interfaces. This course will include video based lectures to introduce concepts and algorithms supported by practical and interactive exercises using python notebooks.

CCS CONCEPTS

• **Human-centered computing** → **User interface design**; **Human computer interaction (HCI)**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

intelligent user interface, natural language interfaces, recommender systems, novel interaction techniques

ACM Reference Format:
Albrecht Schmidt, Sven Mayer, and Daniel Buschek. 2021. Introduction to Intelligent User Interfaces. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '21 Extended Abstracts), May 8–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3411763.3445021

CHI '21 Extended Abstracts, May 8–13, 2021, Yokohama, Japan

ACM ISBN 978-1-4503-8095-9/21/05.
https://doi.org/10.1145/3411763.3445021

1 INTRODUCTION

Today's user interfaces include more and more intelligent behavior. Human-computer interaction and artificial intelligence are tightly linked and incorporated in many application domains. New ways in the collaboration between humans, computers, and data are researched, e.g. as outlined in the the vision of "Interactive Human Centered Artificial Intelligence" [8] with a great potential for empowering users. Already now, this is evident from various examples of interactive systems which would not work without an "intelligent" component, that is, using computational methods, Machine Learning (ML) or other techniques from Artificial Intelligence (AI). Here, we see a more tight coupling of the ML development and the user interface (UI) design part [4], keeping the user during the design cycle of the ML part in the loop, too. This includes, for example, analyzing and modeling fundamental input behavior patterns [2], embedding inference into user interfaces [1], and interaction with AI in-the-loop, for example, for text entry and composition [3]. This allows us to now build whole new systems that have not been possible without the support of AI, e.g., [5]. It also creates conceptual challenges as systems make decisions and act partly autonomously, which may lead to different interaction paradigms, e.g., intervention user interfaces [9]. Intelligent User interfaces (IUIs) have gained significance beyond the research domain. Entire companies are built around IUIs or got successful after embedding IUIs. For instance, online shopping is hardly usable whiteout great search features, and similarly, online movie and music streaming services need a good recommendation system.

1.1 Motivating Examples of IUIs

IUIs today already cover many topics and use cases with diverse motivations and addressed issues. Here, we selected a few practically relevant examples for learning about concepts in IUIs in the context of this course (also see Figure 1):

- *Shopping, music, videos:* There are often too many options to potentially present to users to design and build a single user interface that allows users to browse all of them. Hence, we need ways of narrowing down the search space and systems that recommend things, implicitly or explicitly. The course covers both recommendation systems and their resulting adaptivity in UIs.
- *Interaction with services and content:* Increasingly, interactive technologies are integrated into our daily life using natural language as an interface – from questions in search engines to calling a service hotline at the energy provider. Here, chatbots (textual language) and voice agents (spoken language) play a big role and illustrate the importance of combined HCI

Figure 1: Examples of intelligent user interfaces in everyday life. *Left:* More and more people use smart speakers and voice assistants at home. *Center:* Recommendations have become a ubiquitous element in many user interfaces, for example, on movie platforms. *Right:* Modern smartphone keyboards decode even sloppy typing, correct words, and suggest next words. The course covers concepts and techniques relevant to these and further interfaces and examples.

and Natural Language Processing (NLP) work, as demonstrated in this course with relevant examples.

- *Text:* Writing is central to knowledge work, and ubiquitous in many. Hence, enabling efficient and effective interaction with texts is vital for many user interfaces and tasks, such as search, summarization, information extraction. Creating text is also a prominent creative activity: From work documents to papers, communication, and art – intelligent text input has become an expected standard (e.g., mobile keyboards). NLP and input modeling play important roles in this context, which we cover in this course to recreate and explore smart keyboard features.
- *UI Design:* Interface design often provides a virtually unlimited number of options. How do we choose the best interface? Optimization is a core approach to create good or even optimal UIs, according to a formalized metric [6, 7]. This course illustrates this approach for layout optimization.

1.2 Scope of the Course

Beyond specific examples, the course highlights the broad applicability and fundamental motivations for IUIs today. We give a brief overview and discuss the key aspects across application contexts and domains, including personalizing UIs, narrowing down (content) options, reducing UI complexity while keeping comprehensive sets of functionality, automation of interactions, understanding (uncertain) user input, and enabling new interfaces and interaction concepts.

With the described example topics and generalized takeaways, the *scope of the course* is compactly characterized as an applied HCI perspective, as follows:

- Providing basic terms and understanding in ML and AI.
- Focuses on examples that combine AI and UI .
- Uses simple algorithms to communicate the basics, new libraries to explore further.
- Provides a more detailed look into one or more key areas.
- Demonstrates concrete applications of implemented ML algorithms to create IUIs.
- Explicitly does *not* provide in-depth purely ML-focused content, e.g., regarding specific algorithm implementations or model architectures.

2 INTENDED AUDIENCE

This course addresses anyone interested in intelligent user interfaces, particularly researchers and practitioners in human-computer interaction and machine learning. Besides the conceptual foundations, the course includes many examples and hence is suitable for participants interested in the practical aspects of developing and implementing IUIs. In particular, the course does not assume or require prior knowledge in computational methods, Machine Learning, or Artificial Intelligence. Participants can expect to concretely benefit from the course by gaining an overview and basic understanding of intelligent user interfaces and their possibilities, for example, for further investigation in their own research and practical projects. To facilitate this, the course provides concrete takeaways and material, including practical examples and pointers to further relevant work.

3 PREREQUISITES AND PREPARATION

This is a self-contained course. General programming knowledge is required to follow the coding examples (Python). Course participants are expected to prepare for the course by watching the provided intro lecture videos (see next section).

4 COURSE OVERVIEW

The course is structured into four modules with a duration of 75 min per module. The four modules are:

(1) Overview of IUI, AI and ML terms, concepts, and tools
(2) Recommender Systems and Adaptive UIs
(3) Natural Language Processing and Bots
(4) Intelligent Text Input and Optimization

To optimally support remote participation, the instructors will provide recorded lecture videos to introduce the topics for each

module. This is followed by live discussions via Zoom and interactive coding in Python (notebooks will be provided). The following subsections describe the content of each module in more detail.

4.1 Overview of AI and ML Terms, Concepts and Tools

This is the basic module of this course. We discuss what intelligent user interfaces are and how they relate to HCI as well as to AI. We introduce basic terms and concepts in Machine Learning and AI, thus, providing the foundation for the other slots. In particular, this introduction and overview focus on examples that *combine* topics from HCI and AI for intelligent user interfaces: Concretely, simple algorithms are used first to communicate AI basics and fundamental aspects and challenges of their integration into UIs. This is then complemented by exploring further state-of-the-art methods and pointers to relevant libraries for using them in one's own research and practical projects.

All course participants are expected to attend this first module. Thereafter, we suggest participants choose between the other modules based on their interests.

4.2 Recommender Systems and Adaptive UIs

In this module, we first look at challenges that arise from interaction with massive amounts of information. In many domains, the number of choices is too large to let the users decide. Examples are online stores with thousands (or millions) of items for sale. Algorithms are needed to narrow down what the user should look at. Similar challenges arise in music and video streaming or content presentation in social networks or news portals.

We look at different approaches, including collaborative filtering, content-based filtering, and filtering based on contextual and demographic factors. We explain how the basic algorithms work and show a practical example of how to create a movie recommendation system. In the course, we also discuss how the information required for recommender systems can be acquired implicitly and explicitly. This is complemented by discussions on how such approaches can become the basis for adaptive user interfaces.

4.3 Natural Language Processing and Bots

The primary interaction between people is using natural language – be it in spoken or written form. When interacting with computer systems and data, using natural language offers many new opportunities, for example, as an alternative modality or in the context of a multimodal UI. In this session, we examine where natural language processing can be used to improve human-computer interaction. We further discuss the main challenges when moving from a graphical user interface to a text or speech-based interaction model.

First, we look at basic approaches of processing natural language text (e.g., tokenization, normalization, lemmatization, named-entity recognition). We then discuss and explain the usage of algorithms that can discover topics and concepts in texts, summarizing texts (to different lengths), and identify the sentiment of a text or phrase. We conclude this by exploring chatbots and their applications. In the practical session, we introduce different libraries and toolboxes.

4.4 Intelligent Text Input and Optimization

This module demonstrates how computational methods can be used to personalize and adapt user interfaces to individual users, based on observed user behavior. Building on this, we further introduce methods to infer user intention from such personalized behavior models. In particular, we introduce a probabilistic perspective on these aspects and take modern "smart" mobile keyboards as a concrete application context.

Within this context, we particularly examine two problems to be addressed with computational methods, which cover both design time and use: First, we apply computational optimization to the classic problem of keyboard layouting, which involves basic HCI models to inform navigation of a formalized design space towards an "optimal" design solution. Second, we apply probabilistic modeling and inference to keyboard decoding, that is, working out what a user intended to type, even if they do so very sloppily. Here, we will cover a range of modern keyboard features, including key penalization, auto-correction, and word prediction.

5 PRACTICAL WORK

In the hands-on session, participants will be introduced to a set of tools allowing them to get started when building IUIs. The tools used in the course mainly rely on python; for this, we will use the provided Jupyter Notebooks, which will guide participants through the particle tasks.

6 RESOURCES

All of the presentation material for this course, additional videos, and Jupyter Notebooks will be made available at: http://iui-lecture.org/. All materials are made available under the Creative Commons Attribution-Share Alike 4.0 (CC BY-SA).

7 INSTRUCTORS

Albrecht Schmidt is a full professor for Human Centered Ubiquitous Media in the Computer Science Department at LMU Munich. The focus of his work is on novel user interfaces to enhance and amplify human cognition. He is working on interaction techniques and intelligent interactive systems in the context of ubiquitous computing. In 2018 he was elected to the ACM CHI Academy.

Sven Mayer is an assistant professor for HCI at the LMU Munich. Prior to that, he was a postdoctoral researcher at the Carnegie Mellon University in the Future Interfaces Group and received his Ph.D. from the University of Stuttgart in 2018. He uses machine learning tools to design, build, and evaluate future human-centered interfaces in his research. This allows him to focus on hand- and body-aware interactions in contexts, such as mobile scenarios, augmented and virtual reality, and large displays.

Daniel Buschek leads a junior research group at the intersection of Human-Computer Interaction and Machine Learning / Artificial Intelligence at the University of Bayreuth, Germany. Previously, he worked at the Media Informatics Group at LMU Munich, where he had also completed his doctoral studies, including research stays at the University of Glasgow, UK, and Aalto University, Helsinki. In his research, he combines HCI and AI to create user interfaces that enable people to use digital technology in more effective, efficient, expressive, explainable, and secure ways.

ACKNOWLEDGMENTS

This course is partly funded by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt) and by the European Union as part of the ERC Project AMPLIFY (no.683008) and the HumanE-AI-Net Project (AMD-952026-4).

REFERENCES

[1] Daniel Buschek and Florian Alt. 2017. ProbUI: Generalising Touch Target Representations to Enable Declarative Gesture Definition for Probabilistic GUIs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, CO, USA) *(CHI '17)*. New York, NY, USA. https://doi.org/10.1145/3025453.3025502

[2] Daniel Buschek, Julia Kinshofer, and Florian Alt. 2018. A Comparative Evaluation of Spatial Targeting Behaviour Patterns for Finger and Stylus Tapping on Mobile Touchscreen Devices. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 4, Article 126 (Jan. 2018), 21 pages. https://doi.org/10.1145/3161160

[3] Daniel Buschek, Martin Zürn, and Malin Eiband. 2021. The Impact of Multiple Parallel Phrase Suggestions on Email Input and Composition Behaviour of Native and Non-Native English Writers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Yokohama, JP) *(CHI '21)*. ACM, New York, NY, USA.

[4] Huy Viet Le, Sven Mayer, and Niels Henze. 2021. Deep Learning for Human-Computer Interaction. *Interactions* (2021). https://doi.org/10.1145/3436958

[5] Huy Viet Le, Sven Mayer, Maximilian Weiß, Jonas Vogelsang, Henrike Weingärtner, and Niels Henze. 2020. Shortcut Gestures for Mobile Text Editing on Fully Touch Sensitive Smartphones. *ACM Trans. Comput.-Hum. Interact.* 27, 5, Article 33 (Aug. 2020), 38 pages. https://doi.org/10.1145/3396233

[6] A. Oulasvirta, N. R. Dayama, M. Shiripour, M. John, and A. Karrenbauer. 2020. Combinatorial Optimization of Graphical User Interface Designs. *Proc. IEEE* 108, 3 (March 2020), 434–464. https://doi.org/10.1109/JPROC.2020.2969687 Conference Name: Proceedings of the IEEE.

[7] Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes (Eds.). 2018. *Computational Interaction.* Oxford University Press.

[8] Albrecht Schmidt. 2020. Interactive Human Centered Artificial Intelligence: A Definition and Research Challenges. In *Proceedings of the International Conference on Advanced Visual Interfaces* (Salerno, Italy) *(AVI '20)*. Association for Computing Machinery, New York, NY, USA, Article 3, 4 pages. https://doi.org/10.1145/3399715.3400873

[9] Albrecht Schmidt and Thomas Herrmann. 2017. Intervention User Interfaces: A New Interaction Paradigm for Automated Systems. *Interactions* 24, 5 (Aug. 2017), 40–45. https://doi.org/10.1145/3121357