Human-Centered Artificial Intelligence Considerations and Implementations: A Case Study from Software Product Development

Tobias Komischke(✉)

Infragistics, Inc., Cranbury, NJ 08512, USA
tkomischke@infragistics.com

Abstract. This paper provides an overview of artificial intelligence (AI) and human-centered artificial intelligence (HCAI). It presents a case study of applying AI and HCAI to software product development. Considerations such as use case development, user involvement and the creation of a smart assistant are reviewed.

Keywords: Human-centered AI · AI · Machine learning · Smart assistants · Avatars · Product development · HCI · Usability engineering · UX design

1 Introduction

With recent advances, artificial intelligence (AI), technology that is capable of behaving with human-like intelligence, has become more pervasive. AI is a powerful enabler of automation. Yet, as long as humans are part of a system that ingests inputs, processes them and provides outputs, that is to say in all cases where full automation has not been reached or is not desirable, there is an interaction between humans and artificial intelligence that needs to be carefully crafted.

2 Artificial Intelligence

There are many definitions for AI. We can think of it as machines that mimic cognitive functions that humans associate with other human minds, such as learning and problem solving [1]. We can say that AI is any system that passes the Turing test [2].

It is a common misconception that AI is the same as machine learning. This lexical reduction severely limits the application areas of AI and causes confusion in persons who intend to get into this field.

AI can be thought of as a toolbox of various methods such as machine learning, rule-based systems, optimization techniques, natural language processing (NLP) and knowledge graphs. It is worth mentioning that from the perspective of a consumer of AI, the utilized AI method does not matter as long as it generates value for him or her, i.e. whether or not the AI helps in accomplishing a work task more effectively and efficiently. A rule-based system that is deterministic can be more helpful than a system that is based

H. Degen and S. Ntoa (Eds.): HCII 2021, LNAI 12797, pp. 260–268, 2021.
https://doi.org/10.1007/978-3-030-77772-2_17

on a machine learning algorithm and is probabilistic. As the toolbox metaphor indicates, the methods can be mixed, i.e. a voice command may be interpreted through NLP, then further processed through machine learning.

Although machine learning is oftentimes equated to neural networks (incl. deep learning), there are numerous machine learning methods that do not involve neural networks. These include linear and logistic regression, decision trees, random forests, support vector machines, K-nearest neighbors, K-means clustering, and more. Their advantage over neural networks is that the predictions they make are easier explainable, e.g. by breaking them down to show the impact of each of the features (independent variables) on the result (dependent variable). A neural network in this respect is more comparable to a black box. Finally, within neural networks there is a class of algorithms that use 3 or more (hidden) layers in the neural networks for computations. This is called deep neural network or deep learning.

Figure 1 summarizes the distinction in terminology.

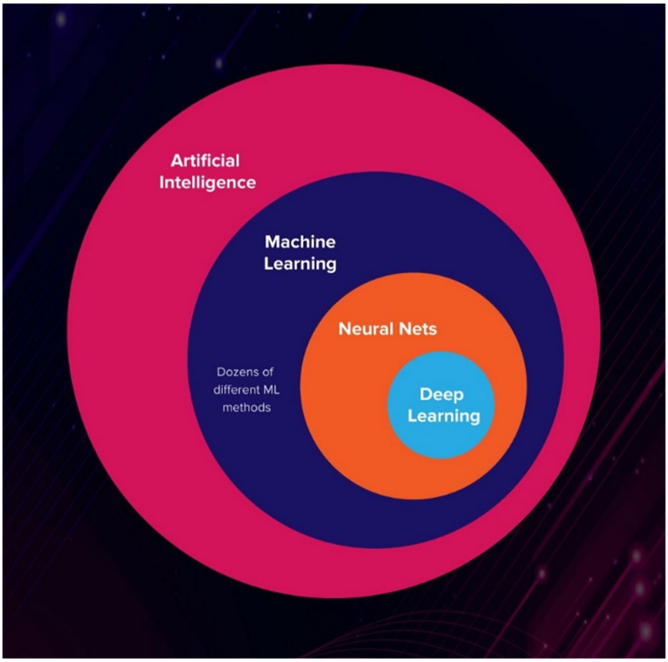

Fig. 1. Terminology of the AI space (Source: https://medium.com/ai-in-plain-english/artificial-intelligence-vs-machine-learning-vs-deep-learning-whats-the-difference-dccce18efe7f)

3 Human-Centered AI

The central question that Human-Centered AI (sometimes abbreviated as HAI, sometimes as HCAI) addresses is: how can AI systems be crafted so that they make communication and collaboration involving humans more effective, efficient and enjoyable? How can they augment human capabilities rather than straight out replacing humans? For users to trust machines, what can be done to help them better understand the strengths and weaknesses of AI? What makes AI acceptable to humans? What tasks lend themselves better for human processing and what tasks are better to be carried out by AI?

Shneiderman emphasizes that because human-centered AI should serve human needs, humans have to be in the center of HCAI [3]. Consequently, humans must stay in control even in highly automated scenarios. In Shneiderman's opinion, human control and automation are not mutually exclusive. This human-centered viewpoint is a continuation of the efforts that starting in the 70s and 80s have allowed broad adoption of computers and software programs. That discipline is called "Human-Computer Interaction" aka HCI.

The HCI discipline has developed rules and standards that provide guidance for the designers and developers of interactive digital systems. One of the best-known standard is ISO-9241 which consists of several parts covering a broad range of critical topics such as software ergonomics and human-centered design processes. While these guidelines are still valid and valuable today, their application to the field of HCAI is limited due to several reasons:

- Some guidelines are in direct conflict with AI. ISO 9241 Part 110 [4] states that a user-centered system shall conform with user expectations. However, because of the dynamic and probabilistic nature of many AI systems, without a proper understanding of the AI's reasoning, the user may have inaccurate or false expectations about the behavior of the systems. Don Norman would call this discrepancy the difference between a conceptual and a mental model [5].
- HCI guidelines do not assume the technology to be a human-like actor being able to pass the Turing Test. The expectations of humans interacting with an "intelligent" system that may utilize a bot or agent as an interface are noticeably higher than with a traditional, utilitarian software tool.
- HCI usability guidelines were created for graphical user interfaces, while "smart" systems have the ambition to facilitate human-system interaction in more natural and seamless ways that mimic human-to-human communication, e.g. through natural language. Research on chatbots demonstrates that for these kinds of systems proper conversation design is more critical than graphical user interface [6].

It is for these reasons that AI requires new rules to guide the development of proper human-centered AI solutions. As first sets of guidelines have emerged (e.g. [7]), case studies are helpful in verifying these rules and to further extend them.

4 Case Study

One area where AI has potential to realize the value proposition of HCAI stated in Sect. 3, are digital productivity and collaboration applications. These are software products assisting individual users, teams and managing content, planning and tracking tasks, creating data insights through dashboards, and communicating and collaborating with others.

Our company has been developing products in that market for years and for the newest product decided to incorporate AI. Some of the key considerations and activities on that journey are discussed in the following.

4.1 AI Vision

As a user-centered product development company, our vision for AI is to create value to users and their organizations through AI capabilities that help improve productivity, collaboration and data insights. This includes but is not limited to creating insights from data and events, portraying these insights through data visualizations that can be easily shared, and suggesting, supporting and automating actions.

To that end, we want to ensure that customers and users are appropriately involved during the design and development of the product. Intimate knowledge about the context of use allows us to derive the application areas for AI from concrete usage scenarios, i.e. AI is not being introduced without a clear use case that starts with the consideration of user needs.

Finally, we strive for our new "intelligent" product capabilities to be acceptable and enjoyable for users. The AI must be non-threatening. Industry survey data shows that from a user perspective the most accepted role for AI is one of an assistant. Users have a much lower acceptance for AI in a managing role or as a peer [8]. Consequently, we made the deliberate choice that our AI would not carry out actions for the users autonomously but would suggest actions to the users. The user then has the option to accept or reject (or plainly ignore) the suggestion, thus staying in control.

4.2 Approach and Process

Following a human-centered AI approach, we have made sure that customers and users are in the loop during the product design and development. We have established a small pool of private preview customers that utilize the product during its development and provide us with feedback. We also present feature ideas and design alternatives to them and incorporate their reactions and preferences.

From a conceptual standpoint, we have been following an evolutionary rather than a revolutionary approach that introduces AI without a big bang, but in a more gradual and subtle way within standard workflows that users are familiar with. This ensures acceptance of the technology. Gartner coined this approach "Everyday AI" [8]. As previously discussed, the AI integration process is based on and revolves around use cases.

Identify Use Cases. As a first step, we researched more than 100 cross-industry case studies from 7 leading companies applying AI and specifically machine learning. We identified the use cases that they were addressing, the AI and machine learning methods and algorithms that were used, and what value was generated, e.g. increased productivity or increased prediction accuracy. The majority of machine learning methods could be categorized as classification algorithms, i.e. class labels are predicted for given examples of input data.

In the next step, we created a list of potential use cases for our own product. They were in part informed by our knowledge and understanding of our customers' context of use and their direct inputs. The teams in our organizations who defined the use cases were the product management team and the AI team. We drafted a document of 36 use cases in 11 feature areas. Each use case described future interactions and abilities for users as well as what AI or machine learning methods could be utilized to realize it – in our case mostly rule-based methods, knowledge graphs, and machine learning methods like classification and natural language processing.

Prioritize Use Cases. The document was shared with other stakeholders within the company and the embedded use cases were prioritized based on value creation for the users and the technical implementation effort. It has since been revised and updated as needed, e.g. as a consequence of insights gained directly from customer feedback or from analytics of our telemetric data gained from private preview customers.

Per Use Case, Identify Data Needs. In order for our AI to be able to understand, predict and suggest, it requires input data aka independent variables ("features" in the terminology of machine learning) from the software product usage, e.g. events triggered by user behavior such as mouse clicks, ASCII input, etc. The data requirements for a use case had to be identified and described to a level that allow effective and efficient queries to the databases and necessary data pre-processing.

Smart Assistant Development. Parallel to the use case work, we developed the face of our AI. To present our AI in a non-threatening way and thus maximize its acceptance by users, we decided to personify the AI into a relatable virtual assistant called "Emily". Through choosing an anthropomorphic representation, we hoped to establish a sense of relatedness and trust in the AI.

To develop Emily's personality in line with the product promise as well our company culture, we used "Mini-Markers" a questionnaire consisting of adjectives describing the basic human personality factors Extraversion, Agreeableness, Conscientiousness, Neuroticism, Openness – the so-called "Big Five" [9]. We created an online survey of the questionnaire and asked internal stakeholders and private preview customers to describe their view of Emily along these five personality factors.

Based on the personality traits, we created a persona describing Emily's mindset, goals, personality and demographic characteristics (Fig. 2).

The persona description helped us in setting the right tone of voice for Emily's communications with users. It also informed Emily's avatar. As Emily shows on the user interface to communicate with users, we developed an avatar. Through numerous design iterations we created design options for the display fidelity determining the details in which facial features, hairstyles, clothing and accessories are rendered (ranging from sketchy comic style to photo-realistic).

The final style was derived from internal and external feedback gained through surveys and depicts Emily as a multi-colored cartoon character created as a vector-graphic. The visual simplicity lends itself well for displaying Emily in small size on the UI: as we do not want her to dominate the screen and our product is also offered for mobile devices with small canvas sizes, in a lot of situations she has to be shown so small that details of her visual features would not be perceivable.

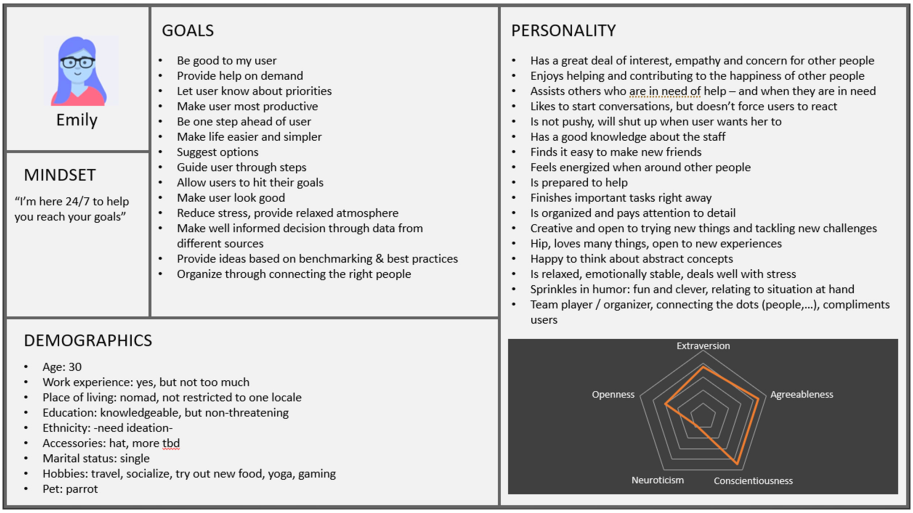

Fig. 2. "Emily" persona with personality traits

How would Emily share insights or suggest certain actions to users? We prototyped options where her avatar would appear and relay her message. We also prototyped the alternative to this push approach, where there may be a visual indicator that Emily has a message, but unless and until a user would access the message, Emily would not show up and potentially interrupt user interaction flows or superimpose on critical screen content. To keep the user in control as much as possible and thus following one of the HCI dialog principles from [4], we decided for the latter approach.

4.3 Selected Applications of AI

In the following, we are highlighting two examples of applying AI techniques. One is a rule-based implementation of user onboarding and microlearning. The other is a machine learning method to predict the number of users of our product, based on time-series.

Onboarding and Microlearning. Users and their organizations have a need of becoming proficient with a new tool as fast as possible. We therefore created a user experience that guides users from defined starting points like digital invitations, organic web searches, marketing web pages, etc. into the product, utilizing contextual information about the users to personalize the experience.

During the first product feature explorations, we provide training guides in the form of information bubbles that explain features and functions on the present screen (see Fig. 3a).

When a customer uses the product, we collect and analyze their interaction data such as click streams. We then utilize our digital assistant Emily to suggest actions to the user. For example, Emily may point out to a user, that he or she can share a newly created asset with co-workers. Refer to Fig. 3b.

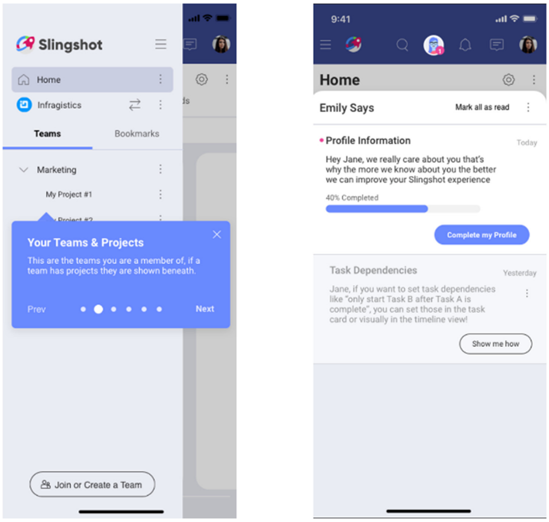

Fig. 3. (a) Information bubbles (left), (b) Emily messages (right)

In total, we identified approximately 40 scenarios where users may profit from guidance. We then distributed these scenarios into two categories: one would be solved through the information bubbles, the other through Emily. Principally, the former would relate to items shown on the present screen, while the latter would be assistance for things that may not be on the present screen. For each scenario we defined the conditions that need to be met in order to trigger the micro-learning feature as well as the verbiage that is being conveyed in the bubble or through Emily.

Forecasting Number of Users. The use of machine learning is not limited to product features, but also to other critical areas. One of them is marketing where clustering user and customer data into distinct sub-groups helps in gaining a deeper level understanding about their characteristics. Another is product management where forecasting the number of users informs about adoption. Time-series predictions can be achieved through various machine learning models. Through our private preview program, we have introduced the tool to a select number of organizations which in turn invited some of their colleagues to use the product to collaborate. For these private preview users, we don't have a long history yet, so the number of data points are limited.

Using 5 months of data reporting on the daily number of registered users, we utilized a SARIMA (Seasonal Auto-Regressive Integrated Moving Average) model. We split our data into a training set (124 days) and a test set (30 days). Training the SARIMA model on the training data, we predicted the remaining 30 days and compared the outcome with the empirical data from the test data. Figure 4 shows the number of private preview users during those 30 days. As a measure of prediction accuracy, we have used the Root Mean Square Error (RMSE). The RMSE is 4.07, implying that the prediction is off by 4 users. Given a private preview user population between 180 to 200 during the period in scope, we consider the prediction accuracy as satisfactory.

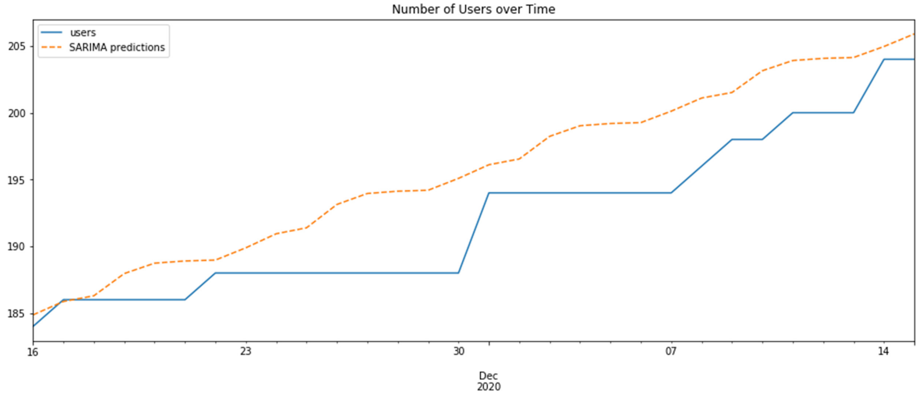

Fig. 4. Actual (solid) and predicted (dotted) number of users

5 Conclusion and Outlook

Artificial intelligence is a toolbox with many varied tools. This provides options as to how results and value can be achieved. The two AI applications presented above are examples of how we leveraged specific tools in the AI toolbox. Both could have been solved using other methods. For instance, we could have also utilized a linear regression or a recurrent neural network to predict user numbers. Different factors play into the decision which method to use. One is the resulting prediction accuracy; another one is the volume of data available. It is important to experiment, comparing methods, and for a specific method, explore the impact of parameter settings on the results.

As mentioned above under "Smart assistant development", classic HCI guidelines like Controllability from [4] still provide value as we design and develop a product with users and the AI-powered system interacting.

Reflecting on the set of HCAI guidelines provided in [7], we find a high degree of concordance. Although the guidelines in that publication were mainly applied to non-business type systems like music recommender systems and social networking, we have found a large set of matches between guidelines and our product concepts. For example, the guideline "Scope services when in doubt" calls for adjusting the AI services when it is not clear what the user's intent is. In a Natural Language Processing (NLP) feature that we prototyped and that interpreted informal written user commands to then suggest the next action, we not only provided the suggestion for the interpretation with the highest certainty, but also offered the next two highest actions for the user to choose from.

One instance where we don't adhere to a guideline yet is for the one that states "Support efficient invocation". It calls for an easy way for users to request the AI system's services when needed. At this point, we are focusing on a push approach where the AI is being triggered implicitly by user actions and other situational factors and not explicitly by user commands. In other words: the communication is one-directional. We do plan to have bi-directional communication in the future, offering both written and voice inputs. This is just one of the many AI features from our backlog that we plan to continue to realize.

Although from a user perspective, a rule-based system can be as good if not better than a machine learning algorithm, machine learning has the advantage that the system is dynamic and can adjust to changes, while rule-based systems are static and need intervention for tuning. For this reason, we expect to increasingly shift the weight of AI methods utilized from rule-based systems to machine learning algorithms. At the end of the day, however, customers and users will determine what approaches provide the highest value. Consequently, HCAI necessitates the involvement of users the same way as we as a human-centered design community have successfully been doing in the field of HCI for decades.

References

1. Russel, S., Norvig, P.: Artificial Intelligence: A Modern Approach, 3rd edn. Prentice Hall, Upper Saddle River (2009)
2. Turing, A.: Computing machinery and intelligence. Mind **59**(236), 433–460 (1950)
3. Shneiderman, B.: Human-centered artificial intelligence: three fresh ideas. AIS Trans. Hum. Comput. Interact. **12**(3), 109–124 (2020)
4. ISO 9241-110. Ergonomics of human-system interaction, Part 110: Dialogue principles (2006)
5. Norman, D.: The Design of Everyday Things: Revised and Expanded. Basic Books, New York (2013)
6. Følstad, A., Brandtzæg, P.B.: Chatbots and the New World of HCI. ACM Interact. **14**(4), 38–42 (2017)
7. Amershi, S., et al.: Guidelines for human-AI interaction. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, pp. 1–13 (2019)
8. Roth, C.: Work Everyday AI into offerings to stay competitive. Gartner report (2020)
9. Saucier, G.: Mini-markers: a brief version of goldberg's unipolar big-five markers. J. Pers. Assess. **63**(3), 506–516 (1994)