How Do Users Interact with AI Features in the Workplace? Understanding the AI Feature User Journey in Enterprise

Sarah D Hanses
Microsoft, USA
hansessarah@microsoft.com

Jennifer Wang
Microsoft, USA
jennifer.wang@microsoft.com

ABSTRACT

This paper investigates the use of AI features - intelligent attributes in products - in the workplace with enterprise users who engage with AI enabled systems through a variety of touchpoints. Oftentimes, product teams developing AI features face a siloed view of AI experiences, and this work aims to present an end-to-end understanding of the range of enterprise users and their experiences when interacting with AI in the workplace. The purpose is to identify the phases in the AI feature journey for enterprise users across their spectrum of experiences, perceptions, and technical acumen. This paper presents this journey of enterprise users working with AI features, analyzes existing challenges and opportunities within this journey, and proposes recommendations to address these areas when planning, designing, and developing AI features for business applications.

CCS CONCEPTS

• **Human-centered computing**; • **Human computer interaction (HCI)**; • **Empirical studies in HCI**;

KEYWORDS

Artificial intelligence, user journey, user experience, user perceptions, human-centered computing, practitioners, enterprise, explainable AI

ACM Reference Format:
Sarah D Hanses and Jennifer Wang. 2022. How Do Users Interact with AI Features in the Workplace? Understanding the AI Feature User Journey in Enterprise. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts), April 29–May 05, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3491101.3503567

1 INTRODUCTION

According to a McKinsey & Company global survey [1] on the state of artificial intelligence (AI) in 2020, 50% of respondents reported having adopted AI in at least one business function. Organizations who have adopted AI report revenue increases for inventory and parts optimization, pricing and promotion, customer service analytics and sales/demand forecasting. While these survey results demonstrate that businesses are increasingly focused on leveraging the power of AI to help improve efficiency and effectiveness, only a minority of companies who have adopted AI in their work processes have recognized the risks associated with AI use, with even fewer working to address those risks. Furthermore, the adoption of AI in the workplace is a trend that is expected to continue. A Mordor Intelligence industry report indicates that the enterprise AI market has registered a CAGR[1] of 52.17% during the forecast period 2021 - 2026, suggesting that enterprises are increasingly recognizing the value of incorporating AI into their business processes [2].

1.1 AI features in enterprise and their complexities

While AI technologies have impacted both consumer and enterprise applications, less research has focused on the unique considerations for AI enterprise experiences. It is important to understand these AI systems as a whole and within a broader human computer interaction framework – this involves the user, the system, the tasks and context, and how these elements combined result in various interactions and outcomes, including perceptions, attitudes, intentions, and behaviors [3]. Within products, AI features can be used across a broad range of industries and business functions, ranging from the consumption of AI insights to the configuration of models. Examples may include analyzing customer sentiment, predicting the potential to convert specific sales leads, optimizing supply chains and inventory, or managing predictive maintenance in factory equipment. These AI use cases reflect a collaborative effort across multiple enterprise users, such as those who may configure their companies' IT structures and solutions, those who analyze data and share insights, and those who use an AI feature's output to make business decisions. These disciplines must interact in a collaborative way for AI features to operate effectively and provide value to enterprise organizations.

1.2 Present study

Acknowledging the increasing use of AI in the workplace and the need to understand the related risks and challenges of AI use in the workplace so that they may be mitigated, this research sought to understand how enterprise users manage and interact with intelligent features within a situated ecosystem and across the entire lifecycle. This research also sought to uncover user motivations, aspirations, and concerns when working with AI features in business applications via extensive 1:1 interviews with representative enterprise users who ranged from consuming AI insights to setting up models for their company.

The intended outcome in researching these business scenarios and user behaviors was to identify challenges and potential opportunities in AI enterprise experiences. The application of this

[1]CAGR = Compound Annual Growth Rate, a metric for mapping growth over a period of time

CHI '22 Extended Abstracts, April 29–May 05, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-9156-6/22/04.
https://doi.org/10.1145/3491101.3503567

Table 1: Participants were selected to represent primary enterprise personas

Admins	Makers	Analysts	End Users
Technical workers who manage their companies' IT infrastructure	Semi-technical workers who build apps for their companies or clients	Data and statistics experts who analyze data and share insights with End Users	Non-technical business workers (e.g., sellers, marketers) who use insights to make decisions

research was intended to influence how product and design teams build AI features that both enable ethical design [10] and product development to move towards effective execution, while supporting the needs of enterprise users with their job efficiency and effectiveness. These intended outcomes address an emerging consensus that designing for AI is challenging, complex and emergent, and that a framework for understanding human-AI interactions can enable more effective envisioning and refinement of AI use [4]. We share the AI feature user journey in enterprise, as well as discuss its application within product development and share future considerations.

2 METHODS

2.1 AI feature terminology

Recognizing that enterprise users may have varying interpretations of different terminology used to describe technical systems (e.g., system, tool, feature, application, solution), the terminology to be used was evaluated as part of a separate research study [6]. This study determined that users most consistently understood an "AI feature" to be focused on completing one specific task or function, and contained within an application, website, or system. As a result, "AI feature" was used throughout interview protocols and screening criteria to describe the AI technology used. Examples of AI features in enterprise include a cash flow forecast enabled by machine learning, real-time transcription of calls enabled by speech recognition, or customer sentiment analysis enabled by natural language processing.

2.2 Participants

In this study, 57 participants took part in semi-structured interviews that were conducted between August 2020 and January 2021. Participants were selected to represent primary personas across a range of different functional roles within enterprise. Within our product group, there are 30+ products and 90+ specific business profiles, and enterprise personas help categorize similar jobs-to-be-done across these business profiles. For this study, business profiles were aggregated into four personas representing our primary enterprise users across various verticals in which users configure, create, or consume AI experiences, and which are outlined in Table 1.

Furthermore, it was a requirement that participants have direct experience in using or configuring an AI feature in their current role. While we recognize that not all enterprise users interact with AI features, selection was limited to users who had prior experience working with AI features, so that their collective experiences could be understood and aggregated to develop the end-to-end journey of enterprise users working with AI features.

2.3 Interview protocol

Semi-structured interviews focused on both the business experience and overall perceptions of AI in the workplace, as well as specific interactions between personas when working with AI features. Participants described their end-to-end process and overall experience working with AI features, which served the purpose of both surfacing concrete examples of how AI is used in the workplace, as well as uncovering how users perceived their experiences. Participants also responded to high-level questions regarding perceptions of AI, challenges and concerns with AI, and more specific product-focused questions, such as the business case for adopting an AI feature, how they learn how to use specific AI features, how AI features influence specific business decisions, and who they consider to be accountable for AI features. The more specific questions diverged slightly by enterprise persona based on current understanding of how their roles vary in the workplace. Participants were also asked what else they may want to share about working with AI in the workplace, and additional topics raised by participants were captured.

2.4 Research synthesis, journey mapping, and visualizations

Interviews were conducted with the different personas as a four-part research series, starting with Admins, followed by Makers, then Analysts, and lastly, End Users. Findings were synthesized after each set of interviews to aggregate findings specific to that persona and were shared with product and design teams for feedback.

Following completion of all interviews, findings were aggregated across all personas to develop an overarching user journey map. A journey map visualizes the process that a user goes through and the steps they may take in pursuit of a specific goal, synthesizing a series of user actions, experiences, and feelings, to create a narrative that is visually represented [7]. This user journey map was initially created in Excel, including the major phases (defined as a series of related events) and subphases (defined as each specific event) that users encounter in their AI feature experiences, with their actions and challenges mapped by each phase and persona. This temporal approach was taken to capture the sequential way in which enterprise users interact with AI features in the workplace, with subphases happening in parallel or as iterative processes within the user journey. Findings were further analyzed to assess severity levels for common challenges and to highlight opportunities to improve the AI feature experience. This journey map was iterated over time based on feedback by product-area researchers and other internal disciplines involved in designing AI feature experiences.

Once the user journey was in a finalized state, the co-authors partnered with design teams to create assets that could be consumed by product and design teams. This ultimately resulted in the

Table 2: AI Feature User Journey: Onboard

Perceive	Identify	Decide
Interpret or understand AI in a particular way	Identify a need or technical opportunity that enables business goals	Evaluate potential technical solutions and select one

development of a high resolution Figma file that visually demonstrates the AI feature user journey in enterprise (see supplementary file to this paper) and other interactive materials that were shared with internal teams.

3 USER JOURNEY FINDINGS

After synthesizing findings across all enterprise personas, the following phases of the AI feature user journey were identified: *Onboard, Operationalize, Launch, Use, Maintain,* and *Evaluate*. These phases and their corresponding subphases are outlined in further detail within this section. While this journey is presented as mostly linear, much of this journey is iterative, while some phases and subphases are experienced in parallel. Using this journey map, we created a framework for identifying commonalities and differences between enterprise personas in their AI feature journey.

3.1 Onboard

In the *Onboard* phase, an enterprise user identifies a need or an opportunity to implement AI and decides to use a particular AI feature for their team or organization, as outlined in Table 2.

Perceive: Enterprise users have different degrees of AI knowledge and perceptions of AI, with individuals varying in how they understand and trust AI [11]. The degree of AI understanding and knowledge among participants were on a spectrum, with Admins and Makers having experience in model configuration and generally a deeper understanding of AI, such as understanding data requirements and model limitations. They described AI as requiring input from and being dependent on humans and as mimicking human logic and decision making to classify or categorize information based on data. End Users were relatively less technical, were less likely to differentiate between AI and machine learning, and more likely to acknowledge that they are not well versed in AI. Enterprise users shared common perceptions that AI makes work easier and more efficient, and that machine learning learns from data and experience to improve. While users noted that they are advocates of AI in the workplace, many have colleagues who are concerned about the impact AI may have on jobs, making them reluctant to adopt such technology.

Identify: In this subphase, an opportunity to implement efficiencies or reduce errors is identified. In some cases, a specific business problem presents a need to seek out an AI feature to provide certain efficiencies. For example, a finance team may request a predictive model to enable business forecasting. In other use cases, an enterprise user such as a Maker may become aware of the AI feature first, and then identify a way to apply it to an existing business problem. For example, one enterprise user read in Reddit about a computer vision tool that reads business documents, researched the feature in more detail, and ultimately secured approval to launch it within their organization to streamline processing of contractual documents. Similarly, enterprise users may also leverage AI features that they used in previous roles or companies and find ways to bring those experiences and features with them when they transition into new roles.

Decide: Here enterprise users evaluate potential technical solutions and select one. These users, typically Admins and Makers, research and compare AI features by considering multiple factors such as cost, ease of implementation, and ongoing support required. Whereas Admins may be involved in evaluating these decisions on behalf of an entire company, Makers are more involved in demonstrating the potential worth of an AI feature to their organization or team to secure approvals as needed. Makers may also decide to adopt an AI feature that is already available to them within a currently used system, and then drive adoption of that feature within their team, often beyond the scope of their role. End Users are less involved in identifying potential AI features but may provide feedback on potential features regarding how effectively they may be applied in their roles. Other individuals or stakeholder groups within their organizations were also noted to have oversight of these decisions, such as technology boards, senior management and company VPs.

3.2 Operationalize

The *Operationalize* phase includes planning and gathering systems requirements, preparing data, building models, and then testing and tuning those models, outlined in Table 3.

Plan: In this subphase, Admins gather systems integration and data requirements, while Makers and Analysts outline overarching business questions and understand what data exists and its format. End Users may provide feedback on inputs and outputs needed for a specific AI feature, such as what product options are needed in a pricing strategy feature, while Admins collect and organize these requirements.

Prepare Data: Here it is primarily Makers and Admins who are involved in collecting data, removing restricted data (e.g., personally identifiable information, medical records), creating data sets from publicly available data sources (if needed), and determining how the data will be moved around and shared. The collection process may involve connecting to external data sources, also requiring setup and configuration. Tasks may involve re-formatting data from certain file formats, inspecting data for errors and missing information, and following up with clients or internal teams for supplementary information. Finally, users must inspect the quality of the data and assess if the data is representative of the reality it is predicting, to mitigate potential bias or model predictions with low confidence values. Users also cited other individuals or stakeholder groups involved in these processes, such as legal, compliance and data security teams.

Table 3: AI Feature User Journey: Operationalize

Plan	Prepare Data	Build	Test	Tune
Organize and arrange for system building and integration	Collect, format, and make ready the data to be used	Create and assemble the necessary components for an AI model	Measure and check quality/performance	Improve the feature by making small changes

Table 4: AI Feature User Journey: Launch

Deploy	Adopt	Learn
Make the AI feature ready for use	Begin using an AI feature and integrate within workflows	Acquire the necessary skills to use an AI feature

Build: In this subphase, typically Makers create and assemble the necessary components for an AI model. The build stage is specific to enterprise users who are building custom AI features, versus those who are setting up or using out-of-the-box AI features. In this subphase, users build AI models in various tools, most usually in R or Python, and create workflows or templates.

Test: In this subphase, users check the quality and performance of the AI feature before launching it more broadly. This encompasses trialing different algorithms based on business requirements, testing the AI feature using a certain percentage of available data, and measuring the model's success by comparing the feature's outputs to known actuals. This is a collaborative process with Makers testing an AI model's accuracy and with End Users trialing the feature to evaluate how it works for their business scenarios and providing feedback on its accuracy and effectiveness.

Tune: Here Admins and Makers improve the feature by making small changes until a certain accuracy level is achieved. They may re-evaluate training data to ensure it is representative, update the model with new parameters or variables to improve accuracy, follow up with external clients to request more detailed or refined data inputs, or integrate feedback received from End Users to improve the AI feature's outputs.

3.3 Launch

In the *Launch* phase, the AI feature is deployed and users begin to adopt and learn how to use the AI feature. The launch phase and its encompassing subphases are outlined in Table 4.

Deploy: It is primarily Admins and Makers who prepare the AI feature for use in this subphase, while Analysts and End Users become aware of the AI feature. Admins are involved in managing permissions following specific criteria and running user acceptance testing to ensure a broader launch will not disrupt other systems. Both Admins and Makers are involved in preparing learning materials and developing awareness of how an AI feature will be used and when it will be launched. Additional enterprise personas are also involved, such as communications teams who may oversee change management notifications and learning teams who partner with Admins and Makers to develop relevant learning materials.

Adopt: This subphase encompasses correlating business purposes and added efficiencies within one's role and adhering to new operating processes incorporating AI features. Admins support or partner with other teams on change management communication, while Makers may evangelize for feature use, often beyond the scope of their current role. For Analysts, AI feature adoption is motivated by understanding and visualizing technical analyses and large amounts of data more easily efficiently. End Users adopt AI features usually as an advised or determined part of their workflows.

Learn: This subphase encompasses acquiring the necessary skills or knowledge to use an AI feature. While Admins partner with learning teams to communicate AI functionality and deliver learning resources for other users, Makers, Analysts, and End Users leverage a broad and varied range of resources to understand and learn AI features. This includes internally developed business or technical documentation, online resources (e.g., blogs, articles, tutorial videos and learning platforms), and trial and error within the AI feature itself. Multiple users across personas said most of their AI learning was self-taught and that this process was necessary to learn how to use the feature. End Users were the most likely to learn how to use the AI feature more informally from supervisors or other colleagues.

3.4 Use

After the AI feature has been launched, it is primarily Analysts and End Users who begin using an AI feature within their workflows. The Use phase and its encompassing subphases are outlined in Table 5.

Setup: In this subphase, enterprise users configure an AI feature's output for a particular use. Makers create workflows or template in this subphase, while Analysts and End Users set parameters and define variables within an AI feature based on business questions and define what data the AI feature can access. Admins may help other users troubleshoot issues, such as issues with available data.

Interpret Output: Here Analysts and End Users aim to understand and examine an AI feature's output. They rely on explanations provided by the AI feature to understand its outputs, as well as representations that illustrate how well a machine learning model performed. They also rely on their own judgement and reasoning

Table 5: AI Feature User Journey: Use

Setup	Interpret Output	Validate Output	Inform
Configure the elements of an AI feature for particular use	Examine and understand an AI feature's output	Ensure output is viable for use	Use AI feature's output to support a business decision

Table 6: AI Feature User Journey: Maintain

Monitor	Update	Resolve
Ensure AI feature continues to work as expected	Refine AI feature and its outputs	Find solutions to technical issues or errors

Table 7: AI Feature User Journey: Evaluate

Assess	Reflect
Measure effectiveness of AI feature	Consider how AI feature impacts work

to make sense of an output, such as understanding if an output matches other trends or is consistent with their experience.

Validate Output: Analysts and End Users then validate that an AI feature's output is viable for use. If the output does not match users' expectations in the previous subphase - either because the feature did not perform well or did not match trends consistent with user experience - then users will inspect the feature. This entails checking for missing or inaccurate data or changing parameters until the user determines that the output is valid.

Inform: Lastly, Analysts share an AI feature's output with other enterprise users in a variety of mediums that visually represent the data. End Users will often use an AI feature's output to support a specific business decision, such as on a marketing strategy or pricing strategy or how to allocate business resources.

3.5 Maintain

In the Maintain phase, the AI feature is monitored, updated and technical issues are inspected and resolved. The Maintain phase and its encompassing subphases are outlined in Table 6.

Monitor: In this subphase, Admins and Makers ensure the feature continues to work as expected. This includes preventing data breaches, ensuring data remains secure, and monitoring model performance. For Admins, it also encompasses checking in with users on how the feature is working and collecting feedback.

Update: Makers refine the AI feature and its outputs in this subphase and evolve the feature as business requirements change. This includes checking if new data is available, updating models, and adapting the feature to changing business requirements. End Users provide feedback to relevant teams, either to teams internal to their organization or the AI feature provider if the feature provides unexpected results.

Resolve: In this subphase, Admins and Makers work to find solutions to technical issues. This includes gathering additional details on such errors, assessing the impact of technical issues, and implementing fixes.

3.6 Evaluate

In the Evaluate phase, enterprise users measure the effectiveness of an AI feature and reflect on how it impacts their work. The Evaluate phase and its encompassing subphases are outlined in Table 7.

Assess: In this subphase, Admins measure the effectiveness of an AI feature to see if it enabled the anticipated efficiencies. They analyze key performance metrics to see if costs or resource management was improved because of a specific AI feature. Other enterprise personas may provide feedback as part of this assessment. This phase is important as it is one basis for enterprises to determine if they will continue using the feature within their workflows, and if they will continue working with a specific AI feature or decide if other solutions would be more effective.

Reflect: Here enterprise users consider how an AI feature impacts their work. Some enterprise users reflected on the overall complexity of AI, while other users noted that they continue to build trust in certain AI features over time. For many enterprise users, this reflection is of positive sentiment and focused on how AI is beneficial to the workplace; in contrast, some users also express concern over its potential impact to job reductions.

4 OPPORTUNITIES FOR PRODUCT TEAMS

This AI journey research uncovered multiple opportunities related to using AI features in the workplace. The primary pain points and opportunities expressed by users include the following: 1) data quality and quantity (*Operationalize* phase), 2) security and privacy (through all phases), 3) lack of adoption and AI comprehension (*Launch* phase), and 4) maintenance and evaluation of the AI (*Use* and *Maintain* phases).

4.1 Data quality and quantity

The most consistently cited challenge from enterprise users was around the quality and quantity of the data that is used to train AI models. Users in the Maker persona reiterated how important it is that the data set used to train AI features be representative, and how challenging this often is in the workplace. In some instances, relevant data sets do not exist, and users need to create data sets themselves from publicly accessible data sources, which may be limited or carry other constraints. Many enterprise users were aware that predictions enabled by AI would only be as good as the data it is trained on, and so, limited data may result in outputs with lower confidence values, or potentially biased results if data is not representative of the reality it is predicting.

Challenges also exist in aggregating, formatting, and cleaning accessible data. This may require multiple data transformations, transferring data between different format types, or enriching data with missing inputs. AI features may or may not alert users when their data is lacking sufficient representation or when AI features are providing outputs whose confidence levels are low due to inconsistencies in or lack of data used to train an AI model. There is an opportunity to create consistent data quality standards in the industry and enable easier connection to data sources, in addition to enabling easier formatting and transformation of data.

4.2 Security and privacy concerns

Another pain point identified in the AI feature user journey across personas centers on security and privacy of AI features. Users consistently pointed to data being stored in the cloud, and the need to ensure that data is not leaked or forgotten when leveraging within AI features. Some users also noted that AI features have automated workflows that could potentially be updated to redirect sensitive information externally, without the awareness of those who have feature oversight. There is an opportunity here to alert individual users, for example Admins and Makers, when features are updated in this way, or to enable users to monitor what an AI feature is doing or has done, such as via an activity log.

4.3 Lack of adoption and lack of AI understanding

For enterprise users who are less technical or have limited experience with AI (generally End Users), this lack of experience or knowledge may contribute to lack of AI feature adoption in the workplace. Understanding and adoption of AI may go hand-in-hand. Some users noted that AI can be overly complex to learn, and in some cases, may not be aware that tools they use rely on AI. Other reasons inhibiting adoption and implementation of AI include concerns that adopting AI may result in job reductions, limitations based on ability to secure additional licenses, or lack of clearly defined ownership for evangelizing AI feature use within an organization.

One opportunity includes leveraging in-context guidance and documentation to make AI features and their capabilities comprehensible and transparent. This first encompasses incorporating AI explainability techniques to support user understanding and trust in AI features, such as providing balanced information about a model's data, performance, and addressing key user questions about AI to bridge the gap between complex models and user understanding [14]. It also encompasses addressing real-world concerns and limitations related to potential impact of adopting AI. Related research has also suggested that organizations explore these topics early, adopt new key performance indicators, and develop a workforce that balances experience with creative and social intelligence all as steps to success in driving and enabling the adoption of AI within their workforces [8].

4.4 Maintenance and evaluation of AI systems

The maintenance and evaluation of AI challenges also presents challenges in enterprise. The larger challenges in this area include ambiguous accountability for feature management and informal or inconsistent feedback mechanisms. Enterprise users widely varied in how they defined accountability for AI features and in describing who they perceived to be accountable for managing them. Many enterprise users perceived accountability for an AI feature to mean ensuring that the feature is live and operating, while some users include ensuring unbiased inputs and outputs in their accountability considerations. Enterprise users did not convey a clear set of roles or responsibilities as it related to managing an AI feature in the workplace [12].

We recommend that enterprise users adopt clear sets of ownership and accountability around AI features, including but not limited to who owns the training of the model, validation of its outputs, feedback on the outputs, and ensuring that data sets are representative. To improve evaluation of AI systems, recommendations include enabling more systematic processes to check in with End Users and enabling users to submit feedback within the AI system itself [13]. Previously defined human-machine teaming capabilities also highlight the need to facilitate the interactions and balance the authority between people and machines in order to be effective [9].

5 CONCLUSION

5.1 Summary recap

It is increasingly critical to organizations overall and to individual users that AI technologies are effectively transparent and applicable within the workplace. This research illustrates that the use of AI features is cyclical in nature, and that AI features require ongoing oversight and collaboration across many business functions. Within the AI feature user journey in enterprise, there are multiple opportunities to improve the experience of working with and using AI features in the workplace. It is critical that these opportunities be addressed, by both individuals and teams using AI features and by those building them, to enable appropriate use of AI in business applications - appropriate in this context being defined as both meaningful for business scenarios and ethical in application.

5.2 Product development application

Previous research has explored how user needs and understanding of AI features differ based on different stages of interaction with AI features and has acknowledged the renewed attention to understanding these differences given the significant influence they may have in human lives and their increased interdependencies [5]. To support understanding of these differences, the results of this research were shared out through a series of presentations to internal product teams that develop AI features, to provide recommendations and to raise awareness regarding opportunities within the AI feature journey for enterprise users. These research findings have been applied by these products' development teams in their planning processes, as product teams have begun using the AI feature user journey as a framework for proposing improvements to AI features and capabilities.

For example, one product team highlighted that explainability of AI features is a challenge for users in multiple phases of the user journey (specifically *Onboard*, *Use*, and *Maintain* phases). The product team stressed that model results can be difficult to interpret by users, that End Users are not always well versed in AI, and that data quality may be poor with little guidance to the user on how to troubleshoot. Therefore, the team prioritized certain capabilities in their planning, such as improving model level explainability for existing out-of-box models within product, via specific explainability mechanisms. This team also mapped their priorities according to the severity of the challenges as indicated within the AI feature user journey. We expect that product teams will continue to leverage this AI feature journey in their product development and planning, and we expect to update the AI user journey based on feedback received and on how teams leverage the journey, to provide a holistic view across product spaces, scenarios, and user groups.

5.3 Future considerations

Future research should consider exploring the experiences of enterprise users who do not interact with AI features, to understand their perspectives and potential barriers to entry, as well as other stakeholders such as legal teams involved in approval processes or learning teams who develop training materials to enable the launch and adoption of AI features. It is important also to explore the perspectives of individuals who feel they have been impacted by the use AI in the workplace, to evaluate potential concerns and biases that may inhibit the use of AI. Future research may also inspect the differences across AI features that are out-of-box vs. custom built, as well as segmenting the journey further based on the type of AI technology used (i.e. computer vision, natural language processing, machine learning, speech recognition). It will also be important to utilize quantitative approaches to better assess pain point severity. Lastly, we expect that adoption, awareness and knowledge of AI features in enterprise will grow over time, making this a dynamic journey that will need to be continually updated.

ACKNOWLEDGMENTS

This research would not have been possible without all the participants who shared their experiences, product specific researchers who shared feedback on study plans and research findings, internal design partners who created visualizations to showcase the AI feature user journey, and product teams and research leads who provided feedback on this paper.

REFERENCES

[1] Tara Balakrishnan, Michael Chui, Nicolaus Henke. 2020. Global Survey: The State of AI in 2020 | McKinsey. https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/global-survey-the-state-of-ai-in-2020

[2] Mordor Intelligence. 2020. Enterprise AI Market - Growth, Trends, Covid-19 Impact and Forecasts (2021-2026). https://www.mordorintelligence.com/industry-reports/enterprise-ai-market

[3] Christine Rzepka, Benedikt Berger. 2018 User Interactions with AI-enabled Systems: A Systematic Review of IS research. In International Conference on Information Systems (ICIS '18), San Francisco, CA, USA, 17 pages. https://www.researchgate.net/publication/329269262_User_Interaction_with_AI-enabled_Systems_A_Systematic_Review_of_IS_Research

[4] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In *CHI Conference on Human Factors in Computing Systems (CHI '20), April 25–30, 2020, Honolulu, HI, USA*. ACM, New York, NY, USA 17 Pages. https://doi.org/10.1145/3313831.3376301

[5] Shipi Dhanorkar, Christine T. Wolf, Kun Qian, Anbang Xu, Lucian Popa, and Yunyao Li. 2021. Who needs to know what, when?: Broadening the Explainable AI (XAI) Design Space by Looking at Explanations Across the AI Lifecycle. In *Designing Interactive Systems Conference 2021 (DIS '21), June 28-July 2, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3461778.3462131

[6] Sarah Hanses. 2020. AI Terminology: How do enterprise users understand and differentiate between AI terms? Internal report.

[7] Sarah Gibbons. 2020. Journey Mapping 101. https://www.nngroup.com/articles/journey-mapping-101/

[8] Vegard Kolbjørnsrud, Richard Amico, Robert J. Thomas. 2016. How Artificial Intelligence Will Redefine Management. https://www.pega.com/system/files/resources/2018-05/hbr-how-ai-will-redefine-management.pdf

[9] Maria Jesus Saenz, Elena Revilla, Cristina Simon. 2020. Designing AI Systems With Human-Machine Teams. https://sloanreview.mit.edu/article/designing-ai-systems-with-human-machine-teams/

[10] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In Proceedings of CHI Conference on Human Factors in Computing Systems (CHI '19), May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA 14 Pages. https://doi.org/10.1145/3290605.3300233.

[11] Leslie Gaines-Ross. 2016. What Do People — Not Techies, Not Companies — Think About Artificial Intelligence? https://hbr.org/2016/10/what-do-people-not-techies-not-companies-think-about-artificial-intelligence?registration=success

[12] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where Responsible AI meets Reality: Practitioner Perspectives on Enablers for Shifting Organizational Practices. Proc. ACM Hum.-Comput. Interact. 5, CSCW1, Article 7 (April 2021), 23 pages. https://doi.org/10.1145/3449081

[13] Juliana Bidadanure, Ge Wang. 2019. Humans in the Loops: The Design of Interactive AI systems. https://hai.stanford.edu/news/humans-loop-design-interactive-ai-systems

[14] Q. Vera Liao, Daniel Gruen, Sarah Miller. 2021. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *CHI Conference on Human Factors in Computing Systems (CHI '20), April 25–30, 2020, Honolulu, HI, USA*. ACM, New York, NY, USA 15 Pages. http://dx.doi.org/10.1145/3313831.3376590