In Procs. IJCAI 2022 workshop on Explanation in AI, Vienna, Austria, July 23, 2022.
Also, available at https://arxiv.org/abs/2206.05030

# Explanation as Question Answering based on a Task Model of the Agent's Design

**Ashok Goel**[1,2], **Harshvardhan Sikka**[1*], **Vrinda Nandan**[1], **Jeonghyun Lee**[2], **Matt Lisle**[2] and **Spencer Rugaber**[1]

[1]Design & Intelligence Laboratory [2]Center for 21st Century Universities, Georgia Institute of Technology

ashok.goel@cc.gatech.edu, {harshsikka, vrinda, jonnalee}@gatech.edu, mlisle@gmail.com, spencer@cc.gatech.edu

## Abstract

We describe a human-centered and design-based stance towards generating explanations in AI agents. We collect questions about the working of an AI agent through participatory design by focus groups. We capture an agent's design through a Task-Method-Knowledge model that explicitly specifies the agent's tasks and goals, as well as the mechanisms, knowledge and vocabulary it uses for accomplishing the tasks. We illustrate our approach through the generation of explanations in Skillsync, an AI agent that links companies and colleges for worker upskilling and reskilling. In particular, we embed a question-answering agent called AskJill in Skillsync, where AskJill contains a TMK model of Skillsync's design. AskJill presently answers human-generated questions about Skillsync's tasks and vocabulary, and thereby helps explain how it produces its recommendations.

## 1 Introduction

AI research on transparency and explanation faces a familiar conundrum. On one hand, the more complex the design of an AI agent, the larger is the need for making the inner working of the agent transparent to the user. On the other hand, the more complex the design, the more difficult it is to generate explanations of how the agent produces an output. AI faces this conundrum irrespective of the paradigm for designing the AI agent: knowledge-based or data driven, symbolic or connectionist, embodied or software, or some combination of them.

We adopt a stance towards the generation of explanations in AI agents that is both design-based and human-centered. Our stance towards explanation is design-based in that we seek to answer questions about the AI agent based on an explicit model of how the agent produces its results. In particular, we capture the agent's design through a hierarchical Task-Method-Knowledge (TMK) model that specifies the agent's tasks and goals, as well as the mechanisms, knowledge and vocabulary it uses for accomplishing the tasks [Murdock and Goel, 2008; Goel and Rugaber, 2017]. Our stance is human-centered in that the questions are acquired from data collected from humans and the answers are meant for consumption by real users. The questions are acquired through participatory design [Muller and Druin, 2012; Spinuzzi, 2005] in focus groups involving key stakeholders.

We present our approach by illustrating the generation of explanations in Skillsync [Robson *et al.*, 2022b], an AI agent that links companies and colleges to facilitate worker upskilling and reskilling. In particular, we embed a question-answering agent called AskJill in Skillsync, where AskJill contains a TMK model of Skillsync's design. AskJill is based on the Jill Watson question-answering technology developed in earlier work [Goel and Polepeddi, 2018; Goel *et al.*, 2021]. At present, both Skillsync and AskJill are under development. AskJill presently answers questions only about Skillsync's tasks and vocabulary. Our research hypothesis is that a TMK model of an agent's design both provides a scheme for classifying questions asked by users and captures the knowledge needed to answer questions that human stakeholders typically ask about the working of the agent. We present preliminary results on the evaluation of AskJill in Skillsync.

## 2 Background

### 2.1 Skillsync

The Skillsync application [Robson *et al.*, 2022a; Robson *et al.*, 2022b] helps companies address the need to reskill or upskill their employees in partnership with colleges. It also helps colleges match their continuing education and professional development programs to the needs of industry. Figure 1 illustrates the workflow in Skillsync. The application enables companies to document the needed skills in the form of training requests and send these training requests to relevant education providers. It also allows colleges to formulate training proposals in response to the requests based on their educational programs. Skillsync uses various AI techniques, including machine learning, language models, and matching algorithms, to extract knowledge, skills, and abilities (KSA) [USCDC, 2022; USVA, 2022]. from job data and to match them with corresponding courses. The job data originates from sources like the U.S. Department of Labor, industry associations, company job descriptions, and job postings provided by the National Labor Exchange (https://usnlx.com/). The course data is sourced from course catalogs of continuing

*Contact Author

education and professional development programs at universities and colleges, including technical and community colleges. The extracted KSAs are organized and prioritized in a skills framework. Skillsync helps make the process of matching jobs with educational programs both more efficient and effective. However, the potential adoption and use of Skillsync in companies and colleges requires that its results are trustworthy, that its processing is transparent, and that it can explain its design and processing.


Figure 1: An overview of Skillsync, adapted from [Robson *et al.*, 2022b]

## 2.2 TMK Models

Task-Method-Knowledge (TMK) modeling constitutes a formal approach to building machine readable knowledge representations of AI agents with the ultimate goal of empowering the agents to reason about other agents or about themselves. By interacting with a TMK model, an AI agent is able to provide causal explanations and manipulate various aspects of the agent being modeled [Goel *et al.*, 1996]. TMK models also fulfill a separate but related goal: providing a machine readable formalism that is also interpretable by humans interacting with the AI agent in question. TMK models encode information in three different ways: Tasks represent the *why* of a system, specifying the goals of the AI agent. Methods represent the *how* of a system, describing the internal processing of the agent. Knowledge captures the *what* of the AI agent, expressing the information that the agent is operating on. TMK models are compositional, causal, and hierarchical. In particular, a method for a task decomposes the task into subtasks that have methods of their own. The tasks at the leaf level are directly accomplished by a chunk of knowledge, an action in the agent, or an interaction with a user. TMK models have been applied to multiple domains ranging from m navigational and assembly planning [Murdock and Goel, 2003] to adaptive web agents [Murdock and Goel, 1999] and design of game-playing agents [Jones *et al.*, 2009]. AskJill in Skillsync uses a TMK model of Skillsync to generate explanations of how it works.

## 2.3 Jill Watson

AskJill has evolved from the virtual teaching assistant AI agent, Jill Watson. The goal of the Jill Watson QA agent is to amplify teacher presence in online learning environments. Jill Watson was initially developed in the context of the Online Master of Science in Computer Science (OMSCS) program at Georgia Tech launched in 2014. Students in the program often interact with instructors through discussion forums. Classes in the program demonstrated the need for amplifying teacher presence quickly following launch, as students posted hundreds of questions in the associated discussion forums, resulting in difficulties for the instructional team to successfully answer them all. Jill Watson uses homegrown technology on top of IBM's Watson platform [Ferrucci *et al.*, 2010]. Conceptually, the Jill Watson Q & A agent uses a hybrid classification approach to answer questions pertaining to the syllabus of a course. First, Jill Watson uses statistical machine learning methods to categorize the underlying intent behind an incoming question from a student. Subsequently, a knowledge based classifier parses the question and structures an appropriate response from an underlying knowledge-base of relevant course related information. While the original Jill Watson answered students' logistical questions about a class, more recently, we developed a variation of Jill Watson called AskJill that answers users' content questions based on a User Guide [Goel *et al.*, 2022]. Jill Watson is now being adapted to provide explanations for systems and agents like Skillsync.

# 3 Participatory Design, Question Classification, and Requirement Analysis

## 3.1 Focus Groups and Question Generation

Guided by Spinuzzi's [2005] methodology, our research focused on the initial exploration and discovery stages of participatory design involving walkthroughs, user observations, and feedback solicitation. Specifically, most of the data came from user comments and answers to guided prompts that were designed to elicit user feedback regarding their tasks and needs. A total of five focus groups were conducted between September and December 2020 to gather feedback from potential users of Skillsync through the process of collective discovery [Muller and Druin, 2012]. By using data generated from these focus groups, we aimed to draw inferences about the possible usage of AskJill and corresponding user requirements. Specifically, we focused on identifying notable questions or issues that participants addressed while they were interacting with the Skillsync prototype.

In order to inform development and training of the knowledge base for AskJill, user questions that we extracted from the focus group data were categorized based on an Explanatory Ladder derived from the TMK framework. There are three main categories of explanation, included in Figure 2. At the lowest level of the ladder, the agent can answer questions about its vocabulary; at the next level, questions about knowledge in two subcategories (raw data and information, and inferred knowledge); and the highest level, questions about the reasoning with three subcategories (context of the current interaction, task, the process to accomplish those tasks).

Of the 52 user questions that we extracted, 31 questions were relevant to company users and the remaining 21 to college users. The results of the question classification in the Explanatory Ladder are shown in Figure 2. Our classification results suggest that both types of users actively generated task-related questions that sought clarification or guidance for task completion (e.g., *How do I add a competency? How do I select and add occupational tasks?*).

Overall, our participatory design study led to eliciting authentic and quality feedback. Most focus group participants will be actual users of the interface in the future. Therefore, they tended to be highly interested and motivated to understand the values of the Skillsync interface or how it would meet their specific needs. By using the question-generating

| Task Keywords | Training Request, RFP, Request | Training Proposal, Proposal |
|---|---|---|
| **Goals** | A training request is created by a company. It details an upskilling or reskilling need that a company hopes can be addressed by an educational partner. | College creates a training proposal that meets the needs detailed in the company partner's transmitted training request. |
| **Inputs** | College creates a training proposal that meets the needs detailed in the company partner's transmitted training request. | Receive the transmitted Training Request, Review Training Opportunities Catalog, Select Training Opportunities, Enter Proposal Details, Create Training Proposal Summary |
| **Outputs** | Completed Training Request | Alignment Score, Completed Training Proposal |

Table 1: Examples tasks from TMK.

| Explanatory Ladder Category | Explanatory Ladder Subcategory | # of Company Questions | # of College Questions | Total # of Questions |
|---|---|---|---|---|
| Vocabulary | Vocabulary | 2 | 0 | 2 |
| Knowledge | Inferred Knowledge | 5 | 5 | 10 |
| Knowledge | Raw Data and Information | 3 | 3 | 6 |
| Reasoning | Context | 5 | 3 | 8 |
| Reasoning | Task | 10 | 6 | 16 |
| Reasoning | Process | 6 | 4 | 10 |
| Total # of Qs | | 31 | 21 | 52 |

Figure 2: Classification of Questions from Companies and Colleges.

| Term | Definition |
|---|---|
| training objectives, desired skills | Training objectives are the intended outcome that learners will achieve after completing a training opportunity. |
| training request, RFP, request | A training request is created by a company. It details an upskilling or reskilling need that a company hopes can be addressed by an educational partner. |
| training proposal, proposal | A training proposal is a training provider's response to a training request, including a training plan, cost, and other details. |
| due date, proposal due date | Proposal due date is the date by which a company needs training proposals in order to be considered. |
| cohort | A cohort is a group of trainees who are completing a training program at the same time. |
| alignment score, alignment | Alignment score is a representation of how well a training plan covers the company's requested training objectives |

Figure 3: A subset of Skillsync's Vocabulary

activity, we had an opportunity to explicitly bring their awareness and attention to the AskJill tool prior to the pilot experiment. This allowed us to gain knowledge on what comes to a user's mind and issues that need to be addressed.

# 4 A TMK Model of Skillsync

## 4.1 Specification of Skillsync's Vocabulary

The TMK model of an agent specifies the vocabulary used by the agent. Figure 3 includes a representative list of AskJill's knowledge of Skillsync's vocabulary (terms and their definitions). In total, we have 41 terms and definitions in AskJill's knowledge base.

## 4.2 A Task Model of Skillsync

AskJill's task model of Skillsync comprises task goals, subtasks, and the inputs and the outputs of the task. This information is stored in AskJill's knowledge base and captures a hierarchy of 10 tasks that the user can accomplish on Skillsync, from high-level tasks to primitive actions. For example, to accomplish the task of creating a training plan, the company user needs to create and submit a training request and the college user needs to respond with a well-aligned training proposal. The two users can then agree to move forward with the training, thereby accomplishing their goal of creating a joint training plan. The task model hierarchically breaks down the tasks related to the two tasks until it reaches the primitive tasks. The lowest levels in the task hierarchy involve primitive user actions such as button click, text entry, and file upload. Table 1 captures the "training request" and "training proposal" tasks as examples, showing the associated task goals, subtasks, inputs and outputs.

# 5 AskJill, A Question-Answering Agent for Skillsync

## 5.1 An Overview of AskJill

AskJill is a question answering agent with the goal of providing explanations for tasks and vocabulary present in the Skillsync Platform. When a user logs into Skillsync, AskJill can be accessed via a text window. The user can type their question and can expect precise answers from AskJill within seconds. Figure 7 demonstrates AskJill answering user questions on the Skillsync Platform. In this section, we outline the architecture and underlying concepts that make up the AskJill system, including how incoming questions are classified and parsed, how answers are generated, and how AskJill agents are trained to do all of the above.

## 5.2 Data Flow of an Answer to a Question

When a user asks a question to the AskJill system, it is parsed and then sent to 2D hybrid classification system. In 2D classification, AskJill makes use of two approaches in tandem: a natural language intent classifier that is trained using machine learning and a set of semantic rules that process classified intents into structured queries. Following classification, these queries are sent to the Skillsync knowledge base, where a response is generated based on querying the domain knowledge base. With an appropriately high confidence exceeding a tuned threshold, the response is passed through the dialogue management system, which converts the answer to a natural language response styled after humanlike conversation. This response is sent back to the Skillsync system and displayed on the text interface. After answering, AskJill prompts the user to provide feedback, asking "Was this answer helpful", and stores the user feedback in its database. That feedback is subsequently used for retraining the agent. If AskJill is unable to answer a question, it can gently redirect the conversation to its domain of competence by suggesting alternate topics associated with the questions it is trained on. Figure 4 includes an overview of the data flow and hybrid classification step when the AskJill runtime answers a question.

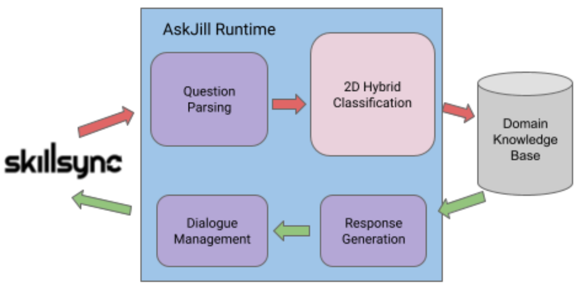


Figure 4: AskJill Dataflow, adapted from [Goel *et al.*, 2022]

### 5.3 Answering Example Questions

Let us now consider a couple of examples. First, say a user asks, "*What is an alignment score?*" via the AskJill text interface on the Skillsync website. AskJill parses the question and sends it to the 2D classification system, where the first layer, the natural language intent classifier, determines that the question relates to vocabulary. Then, the second layer, the rule based classifier, creates a structured query referring to "alignment score". The query is sent to the Domain Knowledge base, implemented in MySQL, that holds the definitions for various terms in the Skillsync glossary. AskJill retrieves the definition of "alignment score" from the knowledge base and formats its response, "Alignment score is a representation of how well a training plan covers the company's requested training objectives". This response is formatted into natural language and sent to the user via the same text window. Similarly, when a user asks "*What is the reason for completing a training request?*"AskJill extracts the intent as "goals" and retrieves the goal associated with a training request. In that case, AskJill responds, "A training request is created by a company. It details an upskilling or reskilling need that a company hopes can be addressed by an educational partner". In both cases, AskJill requests user feedback, asking "Was this answer helpful?", and stores the user feedback in its database. Illustrative examples are provided in Figure 5.

### 5.4 Training AskJill with Machine Teaching

AskJill agents can be rapidly created to support a variety of domains [Goel and Polepeddi, 2018; Goel *et al.*, 2021; Goel *et al.*, 2022], with limited labor and time costs through the use of an interactive machine teaching environment known as Agent Smith. The first step in enabling the Agent Smith system to rapidly train an AskJill Agent is to map the domain to the unstructured and structured databases. As mentioned earlier, the current focus of AskJill in the Skillsync domain is answering questions about Skillsync's tasks and the vocabulary a user may encounter on the platform.

**Mapping Vocabulary and Tasks to the Knowledgebase**

To capture relevant vocabulary from Skillsync in the AskJill knowledge base, key terms were identified based on results from participatory design studies of potential users outlined in Section 3, as well as the conceptual relevance of the term to the domain. These terms and their associated definitions are mapped to the domain knowledge base, which has two distinct tables: the structured and unstructured table. The former has a set of specific pieces of information associated with different kinds of terms within the broader vocabulary and TMK categories, while the latter contains text responses associated with the category. These tables are used during the training process, and again by the AskJill agent during runtime when answering user questions. Tasks were mapped in a similar way, with different actions and interactions taken by users mapped to a Task model, also contained in the same structured and unstructured tables contained in the domain knowledge base.

**Capturing the Form of Potential Questions**

The next important step is to collect a set of questions that will likely be asked about Skillsync as outlined in Section 3.1. These questions are then mapped to a general list of template questions and their associated intents. Template questions seek to capture general forms of questions that occur within categories in the domain, representing a large set of potential questions. These templates are designed by extracting common patterns in the participatory design work outlined in Section 3.1, and classifying them as outlined in Section 3.2. In Table 2, we demonstrate example template questions for different categories in the explanatory ladder. These templates are stored in their own database.

| Type | Example Template |
|---|---|
| Vocabulary | What is {object}? |
| Inputs (tasks) | What inputs do I need to complete this {object}? |
| Goals (tasks) | What is the goal of {object}? |
| Outputs (tasks) | What is the expected outcome of {object}? |
| Subtasks (tasks) | What are the steps to accomplish {object}? |

Table 2: Examples of training templates used to generate training data.

**Dataset Generation**

After the creation of the databases underlying the knowledge base, as well as the template questions database, Agent Smith

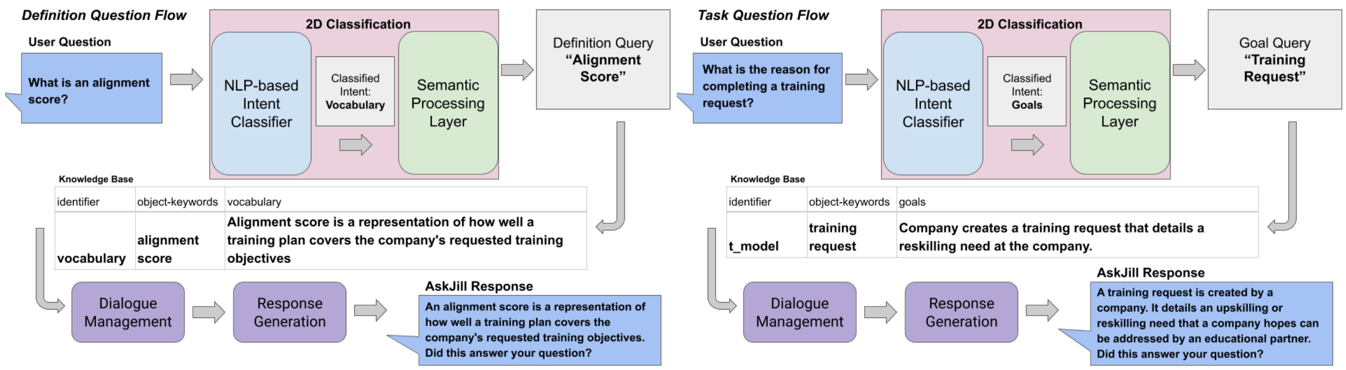


Figure 5: Examples of AskJill QA process for Vocabulary and Task related questions.

can be used to create a large dataset of example question-answer pairs to train an AskJill agent. This is accomplished through combinatorially connecting template questions with various structured and unstructured keywords that represent vocabulary and TMK related concepts.

## 6 Evaluation of AskJill in Skillsync

As part of AskJill's development process, we compiled a glossary of terms (or vocabulary) and developed a task model for Skillsync. We added this information to AskJill's knowledge domain and subsequently trained AskJill to answer a large set of human-generated questions pertaining to vocabulary (relates to "knowledge" in TMK) and users' tasks in Skillsync. We proceeded to gather both in-vitro and in-situ data from AskJill users between June 2021 and September 2021. This period included Skillsync's pilot trial, conducted with participants from Skillsync's partner colleges and companies. During this time, users asked questions about Skillsync's glossary of terms (vocabulary) and the tasks they can complete on Skillsync (task goals, inputs, outputs, subtasks). Figure 6 captures the general categories of questions that AskJill can answer. AskJill was able to answer questions that fall into all five categories (or user intents): vocabulary, task goals, inputs, outputs, subtasks. We validated AskJill's question-answering abilities in-vitro, using its own training dataset of both real user questions and anticipated questions based on known templates. AskJill correctly answered all 1511 questions in the training dataset.

Figure 7 shows examples of human generated questions and AskJill's agent-generated responses to the questions directly from the AskJill text interface on Skillsync's UI. Notice that AskJill is able to provide answers both vocabulary terms and tasks (goals, inputs, outputs and subtasks) on Skillsync. In addition, we gathered a small dataset of in-situ observations. These observations were collected from college and company users interacting with AskJill, embedded in Skillsync, on the platform's website. Figure 8 depicts a comparison of data collected from seven unique users who interacted with AskJill via a text window embedded in the Skillsync UI. Altogether, they asked 219 questions (of which 106 were unique). We validated that AskJill correctly answered 200 questions (91%). Out of the 19 missed questions, 1 had a major language error, 10 were outside of AskJill's competence

| Categories | Example Question | Response |
|---|---|---|
| Vocabulary | What is publish? | Publish is the action taken when the user wants to share work with other companies or training providers. |
| Goals | What is the goal of a training request? | Company creates a training request that details a reskilling need at the company. |
| Inputs | What are the inputs to the trainee profile? | Skills by job title. Skills by Department. Specific Skills search. Skills from uploaded job description. Custom skills entered by company. |
| Outputs | What are the outputs of the training plan? | Accepted Training Proposal between college and company. Denied Training Proposal between college and company. |
| Subtasks | What are the steps to complete a training request? | To accomplish that goal, the following tasks will need to be completed: ... |

Figure 6: Categories of questions AskJill can answer.

(e.g. What is the weather today?) and 8 referred to old UI terms that had been removed from the Skillsync UI. We also assessed user satisfaction directly using a feedback question (Did this answer your question?) that is integrated into the agent's type-in box in the UI. We confirmed that the users indeed provided positive feedback to the correct responses (although occasionally users did not share any feedback).

## 7 Discussion

Mueller et al. [2019] provide a recent and comprehensive review of this research. One of the key ideas to emerge out of this early research was the importance of the explicit representation of knowledge of the design of an AI system [Chandrasekaran and Swartout, 1991; Chandrasekaran *et al.*, 1989], which enables generation of explanations of the tasks it accomplishes, the domain knowledge it uses, as well as the methods that use the knowledge to achieve the tasks. This raised the question of how this design knowledge can be identified, acquired, represented, stored, accessed, and used for generating explanations. One possible answer was to endow the AI agent with meta-knowledge of its own design (e.g., [Goel *et al.*, 1996]).

Although much of the research on expert systems collapsed in the mid-nineties, explanation continued to attract attention in some schools of AI such as case-based reasoning [Leake and Mcsherry, 2005; Schank *et al.*, 2014] and intelligent tutoring systems [Aleven and Koedinger, 2002; Woolf,

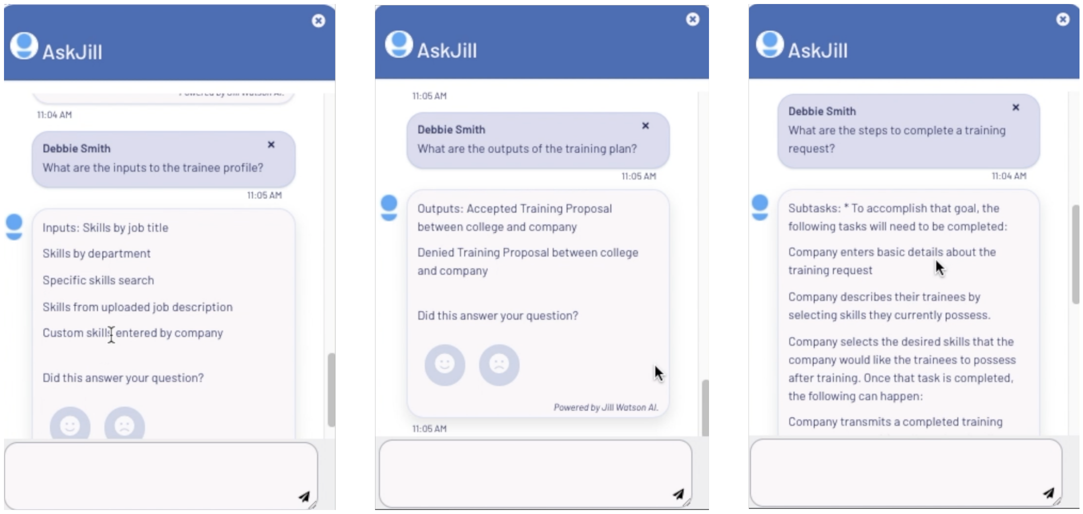


Figure 7: Examples of AskJill's response

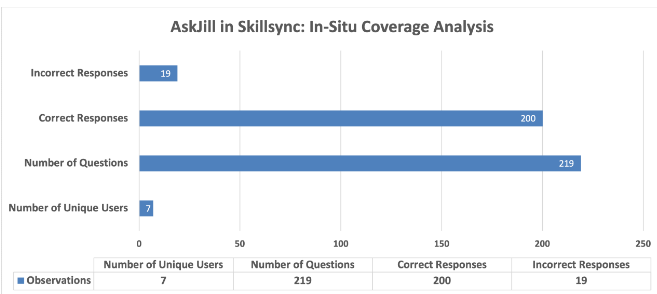


Figure 8: In-Situ Coverage Analysis of AskJill.

2010], often in the form of question-answering [Graesser *et al.*, 1996]. Over the last few years, explanation has again become important in mainstream AI research (e.g., [Gunning and Aha, 2019]). This is in part because of advances in machine learning, such as deep learning, that have refocused attention on the need for interpretability and explainability of internal representations and processing in AI agents. In the future, we plan to extend the TMK model so that AskJill can answer user questions related to the knowledge and methods used in Skillsync. Given the significant changes to the Skillsync platform and user interface based on the two user trials over the development period, we also plan to host additional user focus group sessions to extract additional questions about Skillsync. This participatory design approach enables us to maximize AskJill's question answering ability and enables AskJill to provide explanations for many more real time user generated questions. Another limitation of the current version of AskJill is that it does not afford explanations of specific instances of reasoning and action by the AI agent. Thus, this approach likely has to be complemented with an episodic approach that relies on specific cases of decision making [Langley *et al.*, 2017]. In our own earlier work along these lines, we used meta-cases to capture derivational traces in an earlier interactive learning environment and used the meta-cases to explain the agent's decision making [Goel *et al.*, 1996]. A future version of AskJill may keep a derivational trace of Skillsync's decision making and augment its explanatory capability based on a replay of the derivational trace.

## 8 Summary and Conclusions

A responsible, trustworthy and transparent AI agent must be able to explain how it works and produces its results. Skillsync is an AI agent that links companies and colleges for worker upskilling and reskilling. While Skillsync is useful, its adoption in practice likely will depend in part on its explainability. Thus, we have embedded AskJill, a question-answering agent, into Skillsync. AskJill answers questions about Skillsync's tasks and vocabulary. From the perspective of human-centered AI, we collected and classified questions about the working of Skillsync through participatory design by focus groups. From the viewpoint of design-based explanation, we captured the design of Skillsync through a task model that explicitly specifies its tasks and goals, as well as the vocabulary it uses for accomplishing the tasks. The main takeaway from this work is the usefulness of an explicit model (TMK) of an agent's design for (1) classifying users' questions and (2) question-answering mechanisms for explaining the agent's tasks, goals and vocabulary. In future work, we will complete the TMK model of Skillsync so that AskJill can also answer questions about the methods and knowledge Skillsync uses to complete its tasks.

## Acknowledgements

Our work on AskJill in Skillsync is sponsored by the US National Science Foundation through a Convergence Accelerator grant (#2033578). We thank our collaborators at Eduworks who developed Skillsync for their help with AskJill including Robby Robson, Elain Kelsey, Kristin Wood, and Alan LaFleur. AskJill in Skillsync uses IBM's Watson platform for intent classification. We thank IBM for its support for our work. However, the authors alone are responsible for the contents of this paper.

## References


[Aleven and Koedinger, 2002] Vincent AWMM Aleven and Kenneth R Koedinger. An effective metacognitive strategy: Learning by doing and explaining with a computer-based cognitive tutor. *Cognitive science*, 26(2):147–179, 2002.

[Chandrasekaran and Swartout, 1991] B Chandrasekaran and William Swartout. Explanations in knowledge systems: the role of explicit representation of design knowledge. *IEEE expert*, 6(3):47–49, 1991.

[Chandrasekaran *et al.*, 1989] B Chandrasekaran, Michael C Tanner, and John R Josephson. Explaining control strategies in problem solving. *IEEE Intelligent Systems*, 4(01):9–15, 1989.

[Ferrucci *et al.*, 2010] David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, et al. Building watson: An overview of the deepqa project. *AI magazine*, 31(3):59–79, 2010.

[Goel and Polepeddi, 2018] Ashok K Goel and Lalith Polepeddi. Jill watson: A virtual teaching assistant for online education. In *Learning engineering for online education*, pages 120–143. Routledge, 2018.

[Goel and Rugaber, 2017] Ashok K Goel and Spencer Rugaber. Gaia: A cad-like environment for designing game-playing agents. *IEEE Intelligent Systems*, 32(3):60–67, 2017.

[Goel *et al.*, 1996] Ashok Goel, Andrés Gómez de Silver Garza, Nathalie Grué, J William Murdock, Margaret Recker, and T Govindaraj. Explanatory interface in interactive design environments. In *Artificial intelligence in design'96*, pages 387–405. Springer, 1996.

[Goel *et al.*, 2021] Ashok Goel, Vrinda Nandan, Eric Gregori, Sungeun An, and Spencer Rugaber. Explanation as question answering based on design knowledge. *arXiv preprint arXiv:2112.09616*, 2021.

[Goel *et al.*, 2022] Ashok K Goel, Harshvardhan Sikka, and Eric Gregori. Agent smith: Machine teaching for building question answering agents. In *AAAI Spring Symposium: MAKE*, 2022.

[Graesser *et al.*, 1996] Arthur C Graesser, William Baggett, and Kent Williams. Question-driven explanatory reasoning. *Applied Cognitive Psychology*, 10(7):17–31, 1996.

[Gunning and Aha, 2019] David Gunning and David Aha. Darpa's explainable artificial intelligence (xai) program. *AI magazine*, 40(2):44–58, 2019.

[Jones *et al.*, 2009] Joshua Jones, Chris Parnin, Avik Sinharoy, Spencer Rugaber, and Ashok Goel. Adapting game-playing agents to game requirements. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 5, pages 148–153, 2009.

[Langley *et al.*, 2017] Pat Langley, Ben Meadows, Mohan Sridharan, and Dongkyu Choi. Explainable agency for intelligent autonomous systems. In *Twenty-Ninth IAAI Conference*, 2017.

[Leake and Mcsherry, 2005] David Leake and David Mcsherry. Introduction to the special issue on explanation in case-based reasoning. *The Artificial Intelligence Review*, 24(2):103, 2005.

[Mueller *et al.*, 2019] Shane T Mueller, Robert R Hoffman, William Clancey, Abigail Emrey, and Gary Klein. Explanation in human-ai systems: A literature meta-review, synopsis of key ideas and publications, and bibliography for explainable ai. *arXiv preprint arXiv:1902.01876*, 2019.

[Muller and Druin, 2012] Michael J Muller and Allison Druin. Participatory design: the third space in human–computer interaction. In *The Human–Computer Interaction Handbook*, pages 1125–1153. CRC Press, 2012.

[Murdock and Goel, 1999] J William Murdock and Ashok K Goel. Towards adaptive web agents. In *14th IEEE International Conference on Automated Software Engineering*, pages 335–338. IEEE, 1999.

[Murdock and Goel, 2003] J William Murdock and Ashok K Goel. Localizing planning with functional process models. In *ICAPS*, pages 73–81, 2003.

[Murdock and Goel, 2008] J William Murdock and Ashok K Goel. Meta-case-based reasoning: self-improvement through self-understanding. *Journal of Experimental & Theoretical Artificial Intelligence*, 20(1):1–36, 2008.

[Robson *et al.*, 2022a] Robby Robson, Elaine Kelsey, Ashok Goel, Lauren Egerton, Sazzad M Nasir, Matt Lisle, Alan LaFleur, and Elliot Robson. Making ai work for skill-based training. In *International Training Technology Exhibition and Conference*, 2022.

[Robson *et al.*, 2022b] Robby Robson, Elaine Kelsey, Ashok Goel, Sazzad M Nasir, Elliot Robson, Myk Garn, Matt Lisle, Jeanne Kitchens, Spencer Rugaber, and Fritz Ray. Intelligent links: Ai-supported connections between employers and colleges. *AI Magazine*, 43(1):75–82, 2022.

[Schank *et al.*, 2014] Roger C Schank, Alex Kass, and Christopher K Riesbeck. *Inside case-based explanation*. Psychology Press, 2014.

[Spinuzzi, 2005] Clay Spinuzzi. The methodology of participatory design. *Technical communication*, 52(2):163–174, 2005.

[USCDC, 2022] USCDC. The importance of KSAs. https://www.cdc.gov/hrmo/ksahowto.htm, 2022. Accessed: 2022.

[USVA, 2022] USVA. What are KSAs? http://www.va.gov/jobs/hiring/apply/ksa.asp, 2022. Accessed: 2022.

[Woolf, 2010] Beverly Park Woolf. *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann, 2010.