Designing an Agentic AI Assistant for Chemical Discovery

Siya Kunde[1,*], Stephanie Houde[2] and Rachel K. E. Bellamy[1]

[1]IBM Research, Yorktown Heights, NY, USA
[2]IBM Research, Cambridge, MA, USA

Abstract

We employ a human-centered approach to design an agentic system that can address the unique challenges faced by chemists in the replacement of "*forever chemicals*". We conducted two formative studies with an aim to understand the current processes followed by chemists and the potential role of AI in augmenting and accelerating their scientific process. There's a plethora of open source tools available for use that are specifically geared towards chemical and materials discovery and yet they pose a high barrier to entry for non-coding populations. Through the design of a system that not only allows access to such tools but also additional functionalities from generative AI to domain specific features, we hope to bridge the gap between subject matter experts and AI tools and functionalities that can help them. This work benefits chemists, scientists in domains like materials and drug discovery, toxicologists, regulators and other parties that form a multi-speciality team in tackling the global problem of *forever chemicals*.

Keywords

Design, agentic systems, LLM assistant

1. Introduction

Advancements with large language models (LLMs) have created a wave of research on conversational assistants across a variety of fields like customer engagement [1], library [2, 3], robotics [4], law [5, 6, 7], software coding [8, 9] etc. Each field has its own opportunities and challenges. The task of integrating an LLM powered conversational AI agent into scientific systems is a challenging one due to the level of accuracy necessary, not just for a general task like summarization, but also for tasks specific to science like reasoning.

We can harness the power of Artificial intelligence (AI) to accelerate the scientific discovery process by using it to predict outcomes, generate new artifacts guided by desired attributes and/or make decisions. It is important to ensure a seamless integration of the technology with the human chemists in charge to engender trust in the capabilities of the AI system and ensure optimal results. In most cases however subject matter experts (SMEs) may be left guessing as to the capabilities of the new system or tool that they are asked to use. They may have to spend time familiarizing themselves with a new interface or learn an entirely new skill like coding to

Joint Proceedings of the ACM IUI Workshops 2025, March 24-27, 2025, Cagliari, Italy
*Corresponding author.
skunde@ibm.com (S. Kunde); stephanie.houde@ibm.com (S. Houde); rachel@us.ibm.com (R. K. E. Bellamy)
0000-0002-0138-3862 (S. Kunde); 0000-0002-0246-2183 (S. Houde); 0000-0002-9403-2913 (R. K. E. Bellamy)

launch and use APIs. We bridge this gap through a human-centered approach to the design of such AI systems.

We conducted two formative studies with SMEs to understand their needs, wants and expectations in their work on replacement of "forever chemicals", and the potential role of an AI assistant in supporting that work. In the first study, we interviewed seven chemists using a think-out-loud protocol while attempting to find a fluorine free superacid for photo-lithography using tools of their choice. This was followed up with questions about what role an AI assistant could play in helping them achieve these same tasks. We gained insights into chemists' methodology in tackling the discovery process, the types of tools currently used to achieve this, and how an AI assistant could fill gaps in current technologies and provide a user-friendly interface that would help experts focus on the innovation process. In the second study, we organized feedback sessions with six of the same chemists to present storyboards on various design scenarios. The vignettes showcased ideas to support individual as well as collaborative contributions, while utilizing a conversational AI assistant to search, generate, visualize, manipulate and curate solutions.

2. Background

Per- and poly-fluorinated alkyl substances (PFAS) (also known as "forever chemicals") have over 200 diverse industrial uses [10]. Due to their stability, these materials can be used in harsh environments, and offer unique capabilities. They are however persistent in the environment (estimated lifetimes of 1000 years or more) [11], bio-accumulative and toxic. On the one hand, there is strong regulatory pressure to restrict and eventually eliminate their use (in EU [12] and US [13]). On the other hand, PFAS are critical to many tools, processes and materials in semiconductor manufacturing that are expected to see growth in the near future (CHIPS and Science Act of 2022 [14]). This dissonance between a vision of growth through industrial policy, and a desire for environmental protection through regulatory policy has created an urgent need to replace, redesign and remediate the use of PFAS. In our work, we adopt this as a use case, and aim to create a tool that will help SMEs, in the field of chemistry and material science, discover PFAS replacements that the industry can use to satisfy both the need for growth and for environmental safety.

3. Related Work

In this section we present prior work that describes agentic systems, agentic applications developed in various domains, and finally the current state of AI applications/tools in the field of chemistry. Through this section we present key characteristics of agentic systems, how they have been successfully implemented in other domains and identify the research gap in chemistry domain.

3.1. Agentic systems

The ability of a system to act independently based on a self-generated intention is colloquially referred to as agency. [15] states that instead of looking at the concept of agency from a binary perspective (is agentic or is not agentic), we consider it from the point-of-view of "increasing agency". This might be determined by the degree to which certain characteristics are present within the designed system and the number of them that are present. Key characteristics of agentic systems covered in literature include:

1. Goal directedness [15, 16, 17, 18] : The degree to which a system can adapt.
2. Under-specification [15, 19, 16] : The degree to which a system can accomplish the end goal without exact specification of how to do so.
3. Directness of impact [15, 16, 18] : The degree to which a system can affect the world without a human-in-the-loop.
4. Long-term planning [15, 19, 16, 18] : The degree to which a system makes temporally interdependent decisions in service of achieving the final goal vs over a long time horizon.

3.2. Agentic system applications

LLM-powered conversational assistants have been explored in a number of domains [1, 2, 3, 4, 5, 6, 7, 8, 9]. Some of these have been implemented specifically with an agentic design pattern to take advantage of one or many characteristics we have listed in section 3.1. For example, Cicero [20] is an AI agent that has demonstrated a strong capability to interact with humans in complex environments to achieve their goal using planning and reinforcement learning. LAVE [21] is another LLM-powered assistive agent, which is a language-augmented video editing tool that enables video editing either by the agent or direct user interface manipulation. Combining these modes of interaction provides flexibility and enables manual refinement of the agents actions. In this case, the use of language as a interaction medium enhances multimedia editing by a broad set of users from different backgrounds. Given that scientific discovery teams can be comprised of people from across disciplines and experiences, forming a common ground will be easily done using natural language and our study will explore this as the interaction medium. ChainBuddy [22] is an AI assistant for generating evaluative LLM pipelines built into the ChainForge platform. Through a user study, ChainBuddy was found to provide users with a good starting point for developing their own LLM pipelines. It helped the users across a wide range of tasks and use cases, reducing the effort required, accelerating their workflow, and allowing them to readily learn the platform. The qualities and performance gains demonstrated by these agentic systems [20, 21, 22] are desirable for our use case and we explore these ideas in the studies.

3.3. Tools and assistance systems for chemistry

[23] describes a compound knowledge graph based AI assistant that uses 25 AI models and enables natural language interaction between user and AI through SADL [1], while we explore the

[1]Semantic Application Design Language (SADL) [24] is a formal, structured English-like language and development environment for authoring semantic models that allows non-semantic domain experts to read, write, and/or provide

idea of a LLM powered agentic system. [25] demonstrates a ChatGPT based text mining tool for metal-organic framework enabling parsing, searching, filtering, classification, summarization, and data unification. This prototype demonstrates potential of using LLMs in text-based retriever augmented generation, but is limited in capabilities and resources. [26] is an open source package that solves reasoning-intensive chemical tasks. While it provides functionality important to chemists, it is not accompanied with a user interface, posing a barrier to entry for our non-coding SMEs. This highlights the importance of considering the needs of actual SMEs users when designing systems for them, and adopting a user centered approach increases the likelihood of designing a system that optimally supports SME workflow and doesn't introduce new barriers. Our work adopts such an approach and provides insights into how chemists can be facilitated throughout their discovery process.

4. Iterative User Centered Design

We conducted two formative interview studies to inform our design decisions and iterative lines of inquiry to further the development process. First we interviewed chemists to explore how they approach PFAS mitigation now using current methods. Subsequently we invited them back to participate and help us gather design insights from their reactions to a visualization of what that same process might be like in the future with integrated conversational AI assistance.

4.1. Study 1: Understanding the current process

To help probe how chemists currently work, the interview was structured around a search task characteristic of what a chemist might do when trying to find replacements for PFAS. The design of this task was done with the help of a professional chemist.

Task In line with the use-case we chose (described in section 2), we asked the user to envision themselves as working on the task of "searching for a Fluorine-free super-acid to replace one containing Fluorine in chip manufacturing". We provided them with a SMILES [2] string of a molecule to start off their search and observed how they completed the task, information that they focused on, their thought process and pain-points.

We conducted qualitative interviews with 7 chemists at a large international technology company. Users were first asked questions about their background such as educational qualifications, role as a chemist, and what the role entailed (including interaction with AI systems). This was followed by completing the task. Throughout this task we asked them to think out loud, so that we could have a better understanding of how they analyzed the problem and to gain insights about their decision making processes. Participants interacted with an internal pre-existing tool CIRCA [28, 29, 30] (screenshot shown in Figure 1) that allows patent and molecule search to facilitate the task execution. They were additionally encouraged to show us any other tools they would normally use for similar tasks. The interview concluded with a

feedback on ontologies without requiring extensive training in semantic technologies.
[2]SMILES [27] is a popular string representation of molecules

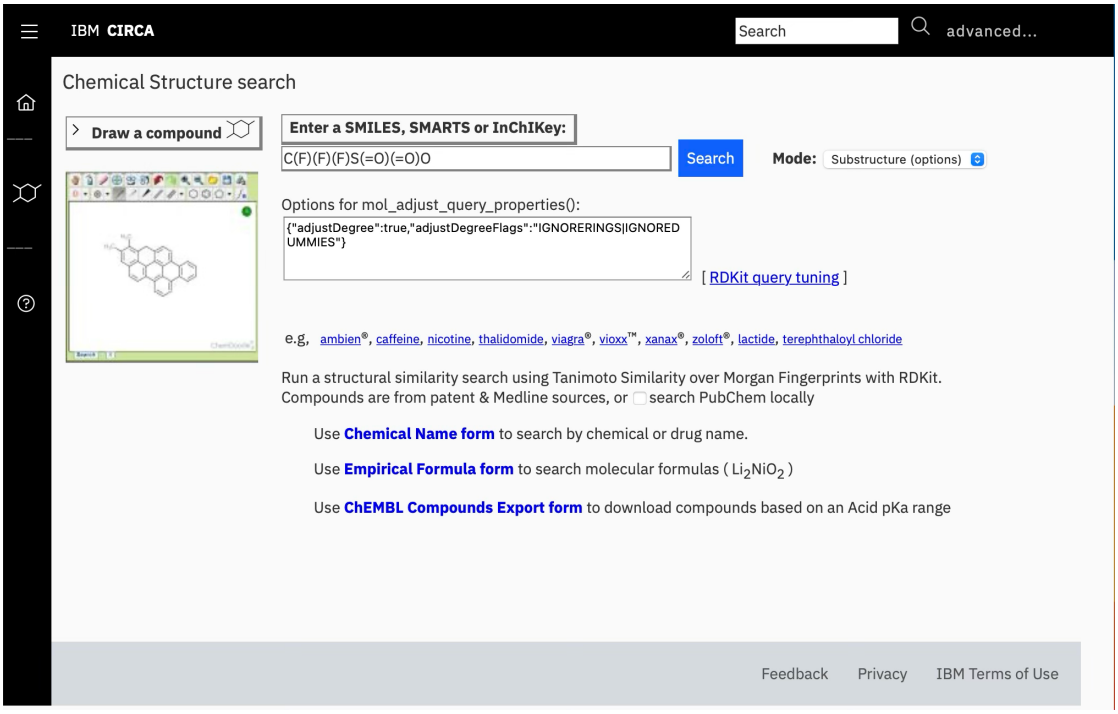

Figure 1: Search page of CIRCA, a tool to search patents and molecules.

discussion about which parts of the process they would have liked to automate with the help of a conversational AI assistant.

Following internal guidelines for human subject studies, we maintained user data privacy, and confidentiality, including obtaining consent and collecting the minimum necessary personally identifying information.

4.1.1. Analysis

Data collected during each study session included video recordings of each session, audio transcriptions, and notes taken by 1-2 researchers observed each session. Two researchers collaboratively reviewed transcriptions and notes to identify notable observations. Using reflexive thematic analysis[31] we identified the themes presented as results below.

4.1.2. Results

We identified the following themes:

1. ***Tools:*** SMEs used a variety of tools, from Google Scholar and Wikipedia to chemistry specific tools like SciFinder [32] and ChemSpider [33].
2. ***Analyzing molecular structural diagrams is a key interaction method:*** These are easily browsed and understood at a glance. The structural sub-components by themselves or in combination with others are key in imparting certain molecular properties. The

ability to view, discuss and manipulate molecular structures and sub-structures impact the task directly and is very important. As P3 put it "its easier to see an invalid structure if it is drawn out." Five participants emphasized their preference for visual searches by drawing substructures. As P2 put it, "We chemists like to work with structures."

3. ***In-lab synthesis is always needed for validation purposes:*** Users mentioned the need to go back to the lab and perform validation experiments to confirm whether a compound could be synthesized. This was because they reported that important implementation details are oftentimes obfuscated in the reference materials. P5 indicated that he would "use different search terms to figure out how to synthesize" a molecule. P4 additionally noted that the feasibility of synthesis would also impact their choices, saying "roughly how many synthetic steps [would be required] would be a consideration."
4. ***Contextual requirements are part of their discovery process:*** When we introduced the working task to a user, in their think-aloud process we discovered additional questions they wanted answered. These questions were focused around the purpose of Fluorine in the current process, if any regulations were applicable, what the regulatory time frame was, if there was a commercially available compound that will work. They posed questions around the ease of synthesis and degradation of the compounds. There were other questions about site specific requirements. P3 said that any real-world search for a new molecule would be "highly application dependent." P4 said that "it would help to always keep the context in mind of which regulations are applicable" when searching.
5. ***They use an iterative winnowing down approach to identify candidates:*** When the users started on the task, the results of initial searches were huge and there were a variety of strategies for winnowing this down to a smaller set appropriate to synthesize. These seemed to vary from chemist to chemist, from filtering for known sources in publications and looking at claims section of patents, to understanding the function that Fluorine was performing in the current application, as a guide to find solutions. P1 commented that "you were trying to reduce the set down" with iterative search filters. P6 noted the need to "whittle [the choices] down to a reasonable number.
6. ***Getting close enough with AI help is good enough:*** To arrive at a shortlist, users mentioned how getting close might be good enough (or even preferred). P2 mentioned how queries may filter out options that might actually be "close enough" for them to adapt and use, and seeing them as part of the search would be preferred. P6 mentioned how all that's needed is chemical "intuition" sparked through new ideas produced by their existing generative model. He also showed interviewers a drawing of a molecule that had an interesting way of providing a function to the molecule and explained how that had sparked ideas that were fruitful in a collaborative discussion between a small group of chemists.
7. ***Source attribution is crucial to trust in the output and decision making process:*** Many users commented that they would not trust output of patents because the implementation details that are important in the synthesis process are typically obfuscated for reasons of intellectual property protection. P2 said that it helps to see "a lot of big suppliers" for a molecule because it shows that is is viable to make or buy it. Similarly P6 included "commercial availability" among the important ways to search for viable molecules.

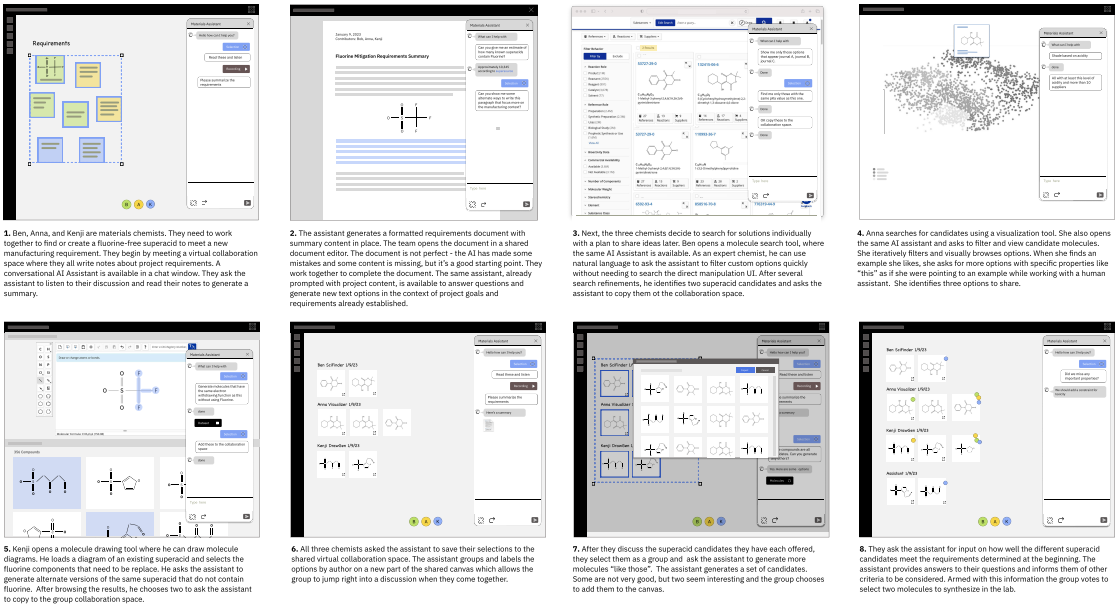

1. Ben, Anna, and Kenji are materials chemists. They need to work together to find or create a fluorine-free superacid to meet a new manufacturing requirement. They begin by meeting a virtual collaboration space where they all write notes about project requirements. A conversational AI Assistant is available in a chat window. They ask the assistant to listen to their discussion and read their notes to generate a summary.

2. The assistant generates a formatted requirements document with summary content in place. The team opens the document in a shared document editor. The document is not perfect - the AI has made some mistakes and some content is missing, but it's a good starting point. They work together to complete the document. The same assistant, already prompted with project content, is available to answer questions and generate new text options in the context of project goals and requirements already established.

3. Next, the three chemists decide to search for solutions individually with a plan to share ideas later. Ben opens a molecule search tool, where the same AI Assistant is available. As an expert chemist, he can use natural language to ask the assistant to filter custom options quickly without needing to search the direct manipulation UI. After several search refinements, he identifies two superacid candidates and asks the assistant to copy them ot the collaboration space.

4. Anna searches for candidates using a visualization tool. She also opens the same AI assistant and asks to filter and view candidate molecules. She iteratively filters and visually browses options. When she finds an example she likes, she asks for more options with specific properties like "this" as if she were pointing to an example while working with a human assistant. She identifies three options to share.

5. Kenji opens a molecule drawing tool where he can draw molecule diagrams. He loads a diagram of an existing superacid and selects the fluorine components that need to be replace. He asks the assistant to generate alternate versions of the same superacid that do not contain fluorine. After browsing the results, he chooses two to ask the assistant to copy to the group collaboration space.

6. All three chemists asked the assistant to save their selections to the shared virtual collaboration space. The assistant groups and labels the options by author on a new part of the shared canvas which allows the group to jump right into a discussion when they come together.

7. After they discuss the superacid candidates they have each offered, they select them as a group and ask the assistant to generate more molecules "like those". The assistant generates a set of candidates. Some are not very good, but two seem interesting and the group chooses to add them to the canvas.

8. They ask the assistant for input on how well the different superacid candidates meet the requirements determined at the beginning. The assistant provides answers to their questions and informs them of other criteria to be considered. Armed with this information the group votes to select two molecules to synthesize in the lab.

Figure 2: A vision scenario of AI support for creative exploration and collaboration. A large language model powered AI Assistant, fine tuned for the materials science domain, offers contextual support throughout a creative search process. A group of chemists use the assistant in both independent and collaborative scenarios to co-creatively define requirements, generate alternatives, and shortlist candidates for synthesis.

Beyond our thematic analysis we additionally prompted participants to think about a future where an AI assistant would help them in some or all parts of the process, users mentioned:

1. Lab automation (help synthesizing chemicals in the lab)
2. Conversational access to tools
3. Generating and identifying productive paths to pursue
4. Data analysis process (help create visualizations)
5. Maintain contextual requirements of the task and act as guardrails

4.2. Study 2: Design scenarios for the future

The themes identified from Study 1 together with user feedback on how AI assistant could help them in their process, we designed vision scenarios (as shown in Figure 2). We conducted a second round of interviews with six of the same chemists from Study 1. In this study we presented users with the vision scenarios of AI support for exploration and collaboration for chemical discovery. The vignette guided users through a process somewhat similar to what they had followed in the first study, strengthened with AI support. Additionally we had incorporated their visions from Study 1 for how an AI assistant might help them in the discovery process. Through this study we obtained feedback on concrete representations of what the users had discussed with us in the first study.

As in the first study, data collected in the second study included one hour video recordings, audio transcriptions, and observer notes. Our analysis of data from the second study was more agile and informal than the first study, consisting primarily of interactive theme identification and discussion between 3 researchers during 30 minute debrief sessions immediately following each session in the manner described by Krug[34] for purposes of informing near-term prototyping efforts. At the conclusion of all sessions the most salient observations, listed below, were agreed upon.

Features in the vision scenario assessed to be beneficial by users included:

1. AI transcription and summarization of meeting notes and discussions helps team focus on tasks and content instead of documentation and formatting.
2. Presence of the same assistant across tools facilitates in-context searching, question-answering, maintaining a detailed record of project requirements in which they are evaluating options.
3. Ability to easily refer to a selection from the chat. Understanding the contextual references (for e.g. when one might say "what is this molecule's acidity?" not only saves time (from re-drawing structures), but is a common and efficient way during in-person discussion for SMEs to reference images and items mentioned previously.
4. Ability to express search criteria using domain specific language. This can mean either using an LLM that understands chemistry specific terms and can produce scientifically accurate responses, or it could be drawing a molecule using a tool that allows such functionality.
5. Ability to combine direct manipulation with natural language requests allows for a natural interaction (like pointing in a discussion with a colleague).
6. Generation of alternatives provides new ideas and directions to explore.
7. AI generation of grouping, labeling, and placement of artifacts (like reports and molecules) provides a default level of organization.
8. AI generation of alternatives based on some criteria (like property similarity for example) in the shared space allows the group to view and consider ideas that teammates had not thought of.

5. Discussion and Conclusion

In this paper we present the findings of two formative studies conducted with chemists, the first to understand their current discovery process when looking for molecules with certain properties, and second to gauge the Chemists reactions to design vignettes visualizing a future AI system that might assist them in this task in the future. From the first study, we were able to understand their approach to the task, their current tool use and the constraints within which they need to operate. The second study provided feedback on specific features of the prospective AI system.

Our studies highlight the benefits of incorporating AI technologies in Chemist's discovery systems. Not only did the users ideate features that were specific to their task, but they also

pointed out how integration of additional general AI features would benefit them in their workflow. The more general AI features mentioned were natural language interfaces, use of generative models, automating transcription, summarization and organization of artifacts, contextual referencing, searching and question/answering. Chemistry related features included integration of specialized tools for drawing and visualizing molecules, and relevant knowledge bases. Other insights from the study were related to 1) the need for source attribution as this allowed them to infer how well tested the molecule was, and 2) the need for any new tool to integrate into the Chemist's existing workflow. We hypothesize these would be well supported by an agentic AI system style. We discuss below the design of an agentic AI assistant for chemistry from a user interaction and experience perspective.

There are several advantages to designing this conversational assistant as an agentic system. One is the ability of the agent to manage the contextual requirements of the goal. For example, the scientist does not need to go and search for regulations across geographic locations, but can incorporate the necessity to abide by regulations as part of the goal given to the system. Some of these requirements can potentially be specified as part of goal setting when a team, individual or organization starts to work with the system, while others will develop over time. Here we can utilize the goal-directedness characteristic of an agentic system [15, 16, 17, 18]. Having the system manage such contextual requirements can benefit cross-domain, multi-stakeholder, cross-functional teams with members having varying types and levels of experience, so individuals do not have to fully understand all the contextual requirements, but can have the system manage them, and explain them at the appropriate time. Exactly how to design such a system so it is effective remains a challenge.

Another advantage of agentic systems is under-specification [15, 19, 16], i.e. their ability to act upon a high-level goal specification. Being able to use the appropriate tools to arrive at the solution further enables the agent in pursuing an under-specified goal. While there are always security and vulnerability concerns with software agents accessing tools [15, 35], given appropriate safeguards agentic systems allow the many tools currently used by users (like PubChem [36], RDKit [37], GT4SD [38], RXN for Chemistry [39]) along with the many open source models (like [40, 41, 42]) and databases (like [43]) to be integrated in the system as resources to be invoked by the agent when needed. This can eliminate the barrier to entry for non-coding users to access tools that might typically require some amount of coding for setup and use.

Sequential reasoning using chain-of-thought thinking [44] is an important capability for planning (in reference to the characteristic of long-term planning [15, 19, 16, 18]) in agentic systems and very useful in emulating the process followed by scientists such as our users. This might mean executing one small query that requires gathering information from multiple tools/sources, or something more longer term like detection of PFAS used in current products, suggesting actions for remediation and following though on the discovery process by searching for replacement materials.

The in-lab synthesis of materials for validation purposes was another feature of their workflow deemed important by users we interviewed. Efforts are underway to automate this process in the real world [45, 46, 47], and despite this emphasis on automation, users will want to remain in-the-loop to avoid unintended consequences like synthesizing a toxic material or by-products. Similarly, there will need for a high degree of involvement from scientists in

other parts of the agentic system workflow. Scientists may need to draw molecules, assess the shortlist of candidate molecules proposed by the agent, etc. The scientists will also need a level of transparency from the system. For example, the scientists in our study mentioned the importance of source attribution in their decision making process, and would want the system to be transparent about where information came from. While working with a complex agentic system, scientists might additionally seek greater level of transparency in other parts of the system like the agents plan, model performance metrics, and tools and the data used. Due to the potential for harm without the supervision of an expert, we propose the system needs human oversight prior to acting in the physical world, having a lower level of autonomy than for less risky actions [15, 16, 18].

We conducted this study with a very small sample size as we were limited by the available number of SMEs in the company. Despite these limitations we gained insight into how to design our system's first prototype. This situation will potentially ease up a little in the next phase of evaluation as we plan to recruit users who are not just expert chemists but also other people who play a role in the PFAS remediation process. These user personas will include, but are not limited to, machine learning specialists and model builders, toxicologists, scientists working on various applications of PFAS, and regulation and policy experts. Though our focus in this study was exclusively on PFAS replacement, our results are highly relevant to other sub-fields in chemistry like drug discovery and this case-study provides insights that might also be valuable for other's designing AI tools for SMEs.

It should be noted that a recent study of materials scientists using AI in their daily work for two years found that while innovation was significantly increased, 82% of the scientists reported reduced job satisfaction when using the AI, and expressed concerns like skill under-utilization, repetitive nature of tasks, credit allocation and complexity of the AI tool [48]. Considering human factors other than innovation and productivity when designing and evaluating future agentic systems, such as job satisfaction, will be critical to their broad success.

Common findings across multiple case studies [49, 50, 51, 48], such as the one described in this paper, provide rich domain-specific findings that can in-turn inform human-centered AI more generally.

References

[1] A. G. Larsen, M. Skjuve, K. Kvale, Llm-powered conversational ai in customer service: Users' expectations and anticipated use (2024).

[2] M. Aboelmaged, S. Bani-Melhem, M. Ahmad Al-Hawari, I. Ahmad, Conversational ai chatbots in library research: An integrative review and future research agenda, Journal of Librarianship and Information Science (2024) 09610006231224440.

[3] A. J. Adetayo, Conversational assistants in academic libraries: Enhancing reference services through bing chat, Library Hi Tech News (2023).

[4] L. Villa, R. Hervás, D. Cruz-Sandoval, J. Favela, Design and evaluation of proactive behavior in conversational assistants: approach with the eva companion robot, in: International Conference on Ubiquitous Computing and Ambient Intelligence, Springer, 2022, pp. 234–245.

[5] B. Liu, Y. Hu, Y. Wu, Y. Liu, F. Zhang, C. Li, M. Zhang, S. Ma, W. Shen, Investigating conversational agent action in legal case retrieval, in: European Conference on Information Retrieval, Springer, 2023, pp. 622–635.
[6] F. Amato, M. Fonisto, M. Giacalone, C. Sansone, An intelligent conversational agent for the legal domain. information 2023, 14, 307, 2023.
[7] H.-T. Nguyen, A brief report on lawgpt 1.0: A virtual legal assistant based on gpt-3, arXiv preprint arXiv:2302.05729 (2023).
[8] S. I. Ross, F. Martinez, S. Houde, M. Muller, J. D. Weisz, The programmer's assistant: Conversational interaction with a large language model for software development, in: Proceedings of the 28th International Conference on Intelligent User Interfaces, 2023, pp. 491–514.
[9] A. Contreras, E. Guerra, J. de Lara, Conversational assistants for software development: Integration, traceability and coordination., in: ENASE, 2024, pp. 27–38.
[10] J. Glüge, M. Scheringer, I. T. Cousins, J. C. DeWitt, G. Goldenman, D. Herzke, R. Lohmann, C. A. Ng, X. Trier, Z. Wang, An overview of the uses of per-and polyfluoroalkyl substances (pfas), Environmental Science: Processes & Impacts 22 (2020) 2345–2373.
[11] D. J. Ivy, M. Rigby, M. Baasandorj, J. B. Burkholder, R. G. Prinn, Global emission estimates and radiative impact of c 4 f 10, c 5 f 12, c 6 f 14, c 7 f 16 and c 8 f 18, Atmospheric Chemistry and Physics 12 (2012) 7635–7645.
[12] Restriction of per- and polyfluoroalkyl substances (pfas) under reach, 2020. URL: https://echa.europa.eu/-/restriction-of-per-and-polyfluoroalkyl-substances-pfas-under-reach.
[13] Key epa actions to address pfas, 2024. URL: https://www.epa.gov/pfas/key-epa-actions-address-pfas.
[14] Chips and science, 2022. URL: https://new.nsf.gov/chips.
[15] A. Chan, R. Salganik, A. Markelius, C. Pang, N. Rajkumar, D. Krasheninnikov, L. Langosco, Z. He, Y. Duan, M. Carroll, et al., Harms from increasingly agentic algorithmic systems, in: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, 2023, pp. 651–666.
[16] T. R. Sumers, S. Yao, K. Narasimhan, T. L. Griffiths, Cognitive architectures for language agents, arXiv preprint arXiv:2309.02427 (2023).
[17] Z. Kenton, R. Kumar, S. Farquhar, J. Richens, M. MacDermott, T. Everitt, Discovering agents, Artificial Intelligence 322 (2023) 103963.
[18] Y. Shavit, S. Agarwal, M. Brundage, S. Adler, C. O'Keefe, R. Campbell, T. Lee, P. Mishkin, T. Eloundou, A. Hickey, et al., Practices for governing agentic ai systems, Research Paper, OpenAI, December (2023).
[19] Y. Ruan, H. Dong, A. Wang, S. Pitis, Y. Zhou, J. Ba, Y. Dubois, C. J. Maddison, T. Hashimoto, Identifying the risks of lm agents with an lm-emulated sandbox, arXiv preprint arXiv:2309.15817 (2023).
[20] M. F. A. R. D. T. (FAIR)†, A. Bakhtin, N. Brown, E. Dinan, G. Farina, C. Flaherty, D. Fried, A. Goff, J. Gray, H. Hu, et al., Human-level play in the game of diplomacy by combining language models with strategic reasoning, Science 378 (2022) 1067–1074.
[21] B. Wang, Y. Li, Z. Lv, H. Xia, Y. Xu, R. Sodhi, Lave: Llm-powered agent assistance and language augmentation for video editing, in: Proceedings of the 29th International Conference on Intelligent User Interfaces, 2024, pp. 699–714.

[22] J. Zhang, I. Arawjo, Chainbuddy: An ai agent system for generating llm pipelines, arXiv preprint arXiv:2409.13588 (2024).
[23] K. S. Aggour, A. Detor, A. Gabaldon, V. Mulwad, A. Moitra, P. Cuddihy, V. S. Kumar, Compound knowledge graph-enabled ai assistant for accelerated materials discovery, Integrating Materials and Manufacturing Innovation 11 (2022) 467–478.
[24] A. Crapo, A. Moitra, Toward a unified english-like representation of semantic models, data, and graph patterns for subject matter experts, International Journal of Semantic Computing 7 (2013) 215–236.
[25] Z. Zheng, O. Zhang, C. Borgs, J. T. Chayes, O. M. Yaghi, Chatgpt chemistry assistant for text mining and the prediction of mof synthesis, Journal of the American Chemical Society 145 (2023) 18048–18062.
[26] A. M. Bran, S. Cox, O. Schilter, C. Baldassari, A. D. White, P. Schwaller, Chemcrow: Augmenting large-language models with chemistry tools (2023). arXiv:2304.05376.
[27] D. Weininger, Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules, Journal of chemical information and computer sciences 28 (1988) 31–36.
[28] T. Griffin, E. Louie, S. Boyer, L. Anderson, K. Schmidt, D. Sanders, Circa: Your cheminformatics assistant, in: American Chemical Society (ACS) 258th National Meeting, 2019.
[29] T. Griffin, K. Schmidt, A system for understanding molecular formula patterns for non-stoichiometric, variable inorganic compounds., in: American Chemical Society (ACS) Fall Meeting, 2022.
[30] Circa, 2019. URL: https://circa.res.ibm.com.
[31] V. Braun, V. Clarke, Can i use ta? should i use ta? should i not use ta? comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches, Counselling and psychotherapy research 21 (2021) 37–47.
[32] S. W. Gabrielson, Scifinder, Journal of the Medical Library Association: JMLA 106 (2018) 588.
[33] H. E. Pence, A. Williams, Chemspider: an online chemical information resource, 2010.
[34] S. Krug, Rocket Surgery Made Easy: The Do-It-Yourself Guide to Finding and Fixing Usability Problems, 1st ed., New Riders Publishing, USA, 2009.
[35] R. Khan, S. Sarkar, S. K. Mahata, E. Jose, Security threats in agentic ai system, 2024. URL: https://arxiv.org/abs/2410.14728. arXiv:2410.14728.
[36] S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B. A. Shoemaker, P. A. Thiessen, B. Yu, et al., Pubchem 2023 update, Nucleic acids research 51 (2023) D1373–D1380.
[37] G. Landrum, Rdkit documentation, Release 1 (2013) 4.
[38] M. Manica, J. Born, J. Cadow, D. Christofidellis, A. Dave, D. Clarke, Y. G. N. Teukam, G. Giannone, S. C. Hoffman, M. Buchan, et al., Accelerating material design with the generative toolkit for scientific discovery, npj Computational Materials 9 (2023) 69.
[39] I. R. for Chemistry team, rxn4chemistry: Python wrapper for the ibm rxn for chemistry api, 2024. URL: https://github.com/rxn4chemistry/rxn4chemistry.
[40] E. Soares, V. Shirasuna, E. V. Brazil, R. Cerqueira, D. Zubarev, K. Schmidt, A large encoder-decoder family of foundation models for chemical language, arXiv preprint arXiv:2407.20267 (2024).

[41] E. Soares, I. Priyadarsini, E. V. Brazil, V. Y. Shirasuna, S. Takeda, Multi-view mixture-of-experts for predicting molecular properties using smiles, selfies, and graph-based representations, in: Neurips 2024 Workshop Foundation Models for Science: Progress, Opportunities, and Challenges, 2024.

[42] I. Priyadarsini, S. Takeda, L. Hamada, E. V. Brazil, E. Soares, H. Shinohara, Self-bart: A transformer-based molecular representation model using selfies, arXiv preprint arXiv:2410.12348 (2024).

[43] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, V. Pande, Moleculenet: a benchmark for molecular machine learning, Chemical science 9 (2018) 513–530.

[44] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al., Chain-of-thought prompting elicits reasoning in large language models, Advances in neural information processing systems 35 (2022) 24824–24837.

[45] M. Abolhasani, E. Kumacheva, The rise of self-driving labs in chemical and materials sciences, Nature Synthesis 2 (2023) 483–492.

[46] G. Tom, S. P. Schmid, S. G. Baird, Y. Cao, K. Darvish, H. Hao, S. Lo, S. Pablo-García, E. M. Rajaonson, M. Skreta, et al., Self-driving laboratories for chemistry and materials science, Chemical Reviews 124 (2024) 9633–9732.

[47] N. Yoshikawa, M. Skreta, K. Darvish, S. Arellano-Rubach, Z. Ji, L. Bjørn Kristensen, A. Z. Li, Y. Zhao, H. Xu, A. Kuramshin, et al., Large language models for chemistry robotics, Autonomous Robots 47 (2023) 1057–1086.

[48] A. Toner-Rodgers, Artificial intelligence, scientific discovery, and product innovation, 2024. URL: https://arxiv.org/abs/2412.17866. arXiv:2412.17866.

[49] G. E. Gonzalez, D. A. S. Moran, S. Houde, J. He, S. I. Ross, M. J. Muller, S. Kunde, J. D. Weisz, Collaborative canvas: A tool for exploring llm use in group ideation tasks., in: IUI Workshops, 2024.

[50] Y. Liu, M. Sra, Exploring ai-assisted ideation and prototyping for choreography, in: Companion Proceedings of the 29th International Conference on Intelligent User Interfaces, 2024, pp. 11–17.

[51] S. Gupta, Y.-C. Chen, C. Tsai, Utilizing large language models in tribal emergency management, in: Companion Proceedings of the 29th International Conference on Intelligent User Interfaces, 2024, pp. 1–6.