

# IdeaSynth: Iterative Research Idea Development Through Evolving and Composing Idea Facets with Literature-Grounded Feedback

Kevin Pu  
Department of Computer Science  
University of Toronto  
Toronto, Ontario, Canada  
kevin.pu@mail.utoronto.ca

K. J. Kevin Feng  
Human Centered Design &  
Engineering  
University of Washington  
Seattle, Washington, USA  
kjfeng@uw.edu

Tovi Grossman  
Department of Computer Science  
University of Toronto  
Toronto, Ontario, Canada  
tovi@dgp.toronto.edu

Tom Hope  
Allen Institute  
Seattle, Washington, USA  
tomh@allenai.org

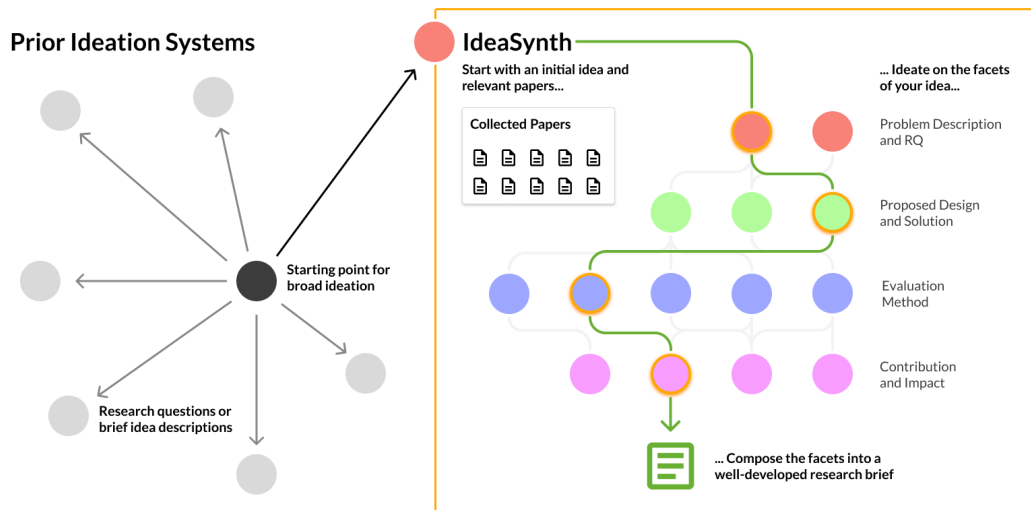
Bhavana Dalvi Mishra  
Allen Institute for AI  
Seattle, Washington, USA  
bhavanad@allenai.org

Matt Latzke  
Allen Institute for AI  
Seattle, Washington, USA  
mattl@allenai.org

Jonathan Bragg  
Allen Institute for Artificial  
Intelligence  
Seattle, Washington, USA  
jbragg@allenai.org

Joseph Chee Chang  
AI2  
Allen Institute for AI  
Seattle, Washington, USA  
josephc@allenai.org

Pao Siangliulue  
Allen Institute for AI  
Seattle, Washington, USA  
paos@allenai.org



**Figure 1:** While prior systems focused on helping users explore a set of diverse ideas broadly, we present IDEASYNTH, an ideation tool that helps users more deeply develop an initial research idea into a concrete and specific research brief. Intuitively, researchers often dissect projects into different facets to further develop them; IDEASYNTH allows users to structure their idea as faceted idea nodes, explore and evolve different variations, and, finally, compose them into a complete research brief.

## Abstract

Research ideation involves broad exploring and deep refining ideas. Both require deep engagement with literature. Existing tools focus primarily on broad idea generation, yet offer little support for iterative specification, refinement, and evaluation needed to further develop initial ideas. To bridge this gap, we introduce IDEASYNTH, a research idea development system that uses LLMs to provide



This work is licensed under a Creative Commons Attribution 4.0 International License.  
CHI '25, Yokohama, Japan  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1394-1/25/04  
<https://doi.org/10.1145/3706598.3714057>

literature-grounded feedback for articulating research problems, solutions, evaluations, and contributions. IDEASYNTH represents these *idea facets* as nodes on a canvas, and allow researchers to iteratively refine them by creating and exploring variations and combinations. Our lab study ( $N = 20$ ) showed that participants, while using IDEASYNTH, explored more alternative ideas and expanded initial ideas with more details compared to a strong LLM-based baseline. Our deployment study ( $N = 7$ ) demonstrated that participants effectively used IDEASYNTH for real-world research projects at various ideation stages from developing initial ideas to revising framings of mature manuscripts, highlighting the possibilities to adopt IDEASYNTH in researcher’s workflows.

## CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**; **Empirical studies in HCI**.

## Keywords

Research Ideation; Scientific Literature; Human-AI Collaboration

### ACM Reference Format:

Kevin Pu, K. J. Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2025. IdeaSynth: Iterative Research Idea Development Through Evolving and Composing Idea Facets with Literature-Grounded Feedback. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 31 pages. <https://doi.org/10.1145/3706598.3714057>

## 1 Introduction

Research ideation often starts with an initial spark of an idea [6]. However, the subsequent process of iteratively developing initial ideas, often high-level and under-specified, into a well-specified and structured research idea grounded by literature is equally important to the success of a research project [16, 25, 52, 54]. Unlike initial research ideation, this process is often *multi-faceted* with researchers explore and consider many variations of research questions, methods, evaluation criteria, and contributions framing. Through this process, researchers iteratively converge on a specific, well-explored, high-impact, and feasible research idea [25, 54]. However, this process is effortful, because it often involves both understanding of literature to generate variations for different facets [30, 35], and also to *compose* different idea facets together to form a coherent research idea [6].

Recent work has shown great promise of leveraging Large Language Models (LLMs) to improve the initial stages of broad research ideation [29, 36, 41, 44, 45, 49]. For example, tools that support literature sensemaking and discovery [29], generating new research questions and ideas [36, 49], or critiquing and evaluating ideas [46, 49]. However, while prior works highlight LLMs’ usefulness in inspiring broad directions during early ideation, they do not account for the *further iteration and deeper development* phase of research ideation. As a result, using LLMs to support the process of focused expansion and distilling initial research ideas into a concrete and specific research idea is left largely under-explored.

In this work, we investigate the user needs and design opportunities to support researchers in the *expansion and refinement*

*stages* of research ideation. Through a formative study with eight researchers, we identified three common pain points in researchers’ workflows. First, researchers struggled with expanding initial ideas to concretely operationalize them into an executable project, which also hindered their ability to evaluate their ideas’ novelty and feasibility. Second, researchers felt unsupported in organizing and evaluating multiple versions and iterations of their ideas. Third, researchers who have tried to use LLMs for research ideation often did not find the feedback helpful because it lacked the level of depth and specificity needed for refining and iterating on ideas.

We address these three challenges by introducing IDEASYNTH, a research idea expansion and refinement system that facilitates the user’s process of developing and experimenting with different variations of a research idea. IDEASYNTH focuses on further evolving initial ideas by complementing human ideation with literature-grounded suggestions, guiding researchers to expand, explore, and connect their ideas, thereby enhancing the ideation workflow through human-AI collaboration. The system features a canvas interface that allows users to externalize variations of different *idea facets*—such as research problem descriptions, proposed methods, and evaluation approaches—as nodes on the canvas. IDEASYNTH’s LLM-powered idea refinement leverages the context of a user-controlled scientific literature collection to help the user refine individual facets and strengthen connections between facets. Users can also create different *idea paths* that connect different facets to generate a *research brief*<sup>1</sup>

To evaluate IDEASYNTH, we conducted a controlled lab study and a field deployment study where participants used IDEASYNTH to develop research ideas. In the lab study ( $N = 20$ ), 12 HCI and 8 NLP researchers developed ideas using IDEASYNTH and a strong baseline system that shared a subset of IDEASYNTH’s features. In addition, seven of the participants also completed a subsequent field deployment study where they used IDEASYNTH for three days to develop research ideas from their own ongoing research projects. We found that IDEASYNTH allowed participants to expand and iterate on their initial ideas, with details from the literature-grounded suggestions. Participants also felt that by externalizing idea facets and their connections, IDEASYNTH supported them in considering more potential idea directions, avoiding idea fixation, and feeling more confident about their final idea. Through log analysis, we also identified behavioral differences where participants in the baseline condition spent more time on writing based on their prior knowledge and perceived the baseline AI assistance lacking for supporting ideation. Further, participants in the field deployment study used IDEASYNTH across various stages of ideation, from initial project specification to framing exploration on manuscripts. They saw it fitting into their existing workflows by lowering the barrier for ideation, and also envisioned using it collaboratively with others. We highlight the implications of our findings by discussing the need to tailor AI support for different research processes and diverse user needs. We also identify potential limitations and ethical concerns of involving LLM in the research ideation process, presenting caveats for future system designs.

In summary, this paper makes the following contributions:

<sup>1</sup>In this paper, we refer to a research brief as a document that describes the necessary facets of a research project, before any project execution. We further define the term in Section 2.1.

- A formative study that illustrates challenges researchers face during ideation, including exploring different variations of ideas, expanding on their initial insights, and eliciting feedback from LLMs that help them to evaluate and refine their ideas.
- Based on the above, we introduce IDEASYNTH, a research ideation tool that helps users expand and develop their initial ideas, with:
  - (1) actionable, literature-grounded LLM-powered feedback for idea expansion, refinement, and evaluation
  - (2) a canvas editor that allowed users to structure their research idea as facet nodes and explore different variations, alternatives, and compositions.
- A lab and a field deployment user study demonstrating how IDEASYNTH can assist users in expanding and refining research ideas and how its designs and interactions can enhance real-world research ideation workflows.

## 2 Background and Related Work

### 2.1 Scientific Research Ideation process

Scientific research ideation is a non-linear complex process consisting of multiple scientific tasks and activities [19]. To come up with a new idea, researchers have to identify and prioritize interesting research problems, form ideas to address selected problems, compare the idea with existing literature, integrate new knowledge to make sense of the relevant information space, and develop an evaluation plan [24]. These tasks appear in different phases of a research project and involve different cognitive processes.

In the beginning phase of scientific research ideation, researchers mostly engage in problem identification and prioritization. This early phase can involve “opening processes” of information-seeking which includes both intentional and serendipitous starting points [19]. Researchers generate and evaluate initial research questions [24]. Like in a typical ideation process, the initial phase is associated with the divergent stage [42] where researchers can benefit from exploring various research questions before committing to a selected few [36, 37]. The outcome of this phase is mostly a set of prioritized research questions that will be further explored.

After researchers decide on an idea to pursue, they start shaping the idea into a research brief, sometimes called a research proposal. In this paper, we define research brief as a description of a research project before any execution (e.g., running an experiment, implementing a system). A research brief helps make the research project more concrete and is often used to communicate the ideas to others for feedback. A research brief can include multiple ideas, ranging from a research question or problem, a method, a hypothesis or antithesis, and a theory [25]. For example, a research brief can have a problem, a method, and a hypothesis with an experimental design. This research-brief-forming phase involves “orientation processes” where researchers make their ideas more grounded and concrete with literature-based evidences and details [19, 24]. Researchers often consult existing literature to situate their ideas, make the problem more specific, generate idea facets (e.g., method, evaluation criteria), and evaluate idea facets and the resulting ideas [30, 35]. This phase of idea generation, typically associated with the converge stage [42], is especially important for research ideation

because of the potentially high costs of a research project. A well-thought of research brief can ensure that resources are spent on high-impact problems and facilitate idea communication.

### 2.2 Research ideation support with AI

**2.2.1 Initial Idea Generation.** Recent work on ideation informs us of the potential and limitations of leveraging AIs for ideation support. For example, prior work suggests that LLMs or humans with access to LLMs can generate more novel ideas [43, 49, 59] and higher quality ideas [20] than humans who do not use LLMs. While the majority of prior work in this space suggests that providing AI support during ideation is beneficial to users [20, 23, 24, 36, 43], a few also pointed to potential risks of incorporating AI in ideation, such as generating less original ideas [56]; less diverse ideas [2, 56]; and more fixation [56]. In sum, prior work points to how carefully design UIs and interactions may be keys to successful integration of AIs into the ideation process, and that failure to do so may incur severe adverse effects.

More directly related to our work, multiple AI systems and interventions have been tailored to support the initial broad ideation of research ideas where “opening processes” happen. CoQuest helps researchers generate research questions with an LLM-based agent [36]. Many prior systems provided users with relevant inspirations instead of generating the ideas directly. For example, helping users find relevant research analogies as inspirations [27]; generating personalized inspirations using LLMs with an evolving knowledge graph [21]. However, most of prior systems only focused on supporting initial broad ideation instead of further developing them into a concrete research brief. For example, CoQuest only supports generating research questions, but does not consider potential solutions nor how to evaluate them to demonstrate impact and contribution.

**2.2.2 Idea Convergence, Evaluation, and Refinement.** Some systems focus support in the converging stage for evaluating research ideas. Prior work has proposed ways to provide users with feedback for generated ideas. Shaer et al. explores the application of LLMs to the converging phase of brainstorming by using LLMs to help evaluate the ideas based on their novelty, relevance, and insightfulness [46]. Liang et al. and D’Arcy et al. separately shows that language models can generate helpful feedback on scientific papers [15, 33]. Lu et al. reports that automated paper evaluation has higher inter-rater agreement than human evaluation [38]. Shen et al. provides AI explanation dialogues to support human-AI scientific writing [48]. These findings suggest the potential for automatically evaluating ideas. However, there are contradicting findings that caution against relying too heavily on automated idea assessment. LLMs still cannot evaluate research ideas consistently [49]. Moreover, evaluating research ideas is a highly subjective task that can be affected by evaluators’ preferences and experience. Human experts do not always agree on the novelty of a research idea [49]. Higher inter-rater agreements of LLM ratings [38] do not mean that the generated ratings capture the intrinsic quality of an idea. Moreover, the format of research ideas that get evaluated in prior work are mostly short text description [36] or a writeup on the formal level of a published paper [11, 33]. Informed by these prior work, instead of producing evaluation scores, IDEASYNTH gives actionable suggestions that

help users decide on what to do to improve their ideas. In contrast to full research papers describing a completion work, we aim to help user produce research briefs that are less formal and more succinct compared to full papers. Yet, compared to broad research ideas that most prior work aimed to target (e.g., generate research questions), research briefs are more specific and detailed, which have different implications for evaluation.

Another research thread explores how to help researchers refine existing ideas. SCIMON compares a given research idea, specified by faceted context (e.g., methods, problems), to existing literature and automatically updates the idea if it is too similar to existing work [57]. ResearchAgent, an LLM-based research idea writing agent, generates research ideas by adding an idea facet at a time starting from problem, method, to experiment design [5]. ResearchAgent also provides review, feedback and ratings of an idea and automatically updates some ideas based on the feedback [5]. While ResearchAgent divides idea generation into facets, they have to be generated in a fixed sequence, unlike in IDEASYNTH that gives users the freedom to generate idea facets in any order. IDEASYNTH also allows users to apply context to specific facets instead of the whole idea. Acceleron, an LLM-based system provides a prompt template that helps users rewrite a research proposal based on identified gaps in literature [41]. Similarly, PaperRobot [58] incrementally generate paper titles and abstracts based on existing papers. However, Acceleron and PaperRobot generate research artifacts with an initial idea description and little human involvement, which is different from the approach of supporting researchers develop a more detailed research brief in IDEASYNTH.

### 3 Formative Study

We conducted a formative study with 8 HCI and NLP researchers (5M, 3F) both to understand the challenges in existing ideation workflows and also to identify design opportunities for enhancing research ideation with AI. The study consists of a walk-through of participant's existing artifact of a past or ongoing research project (e.g., research project proposals and running project notes documents), followed by a live ideation session using a Wizard-of-Oz (WoZ) prototype that leveraged an off-the-shelf research support tool. Finally, we conclude the study with a semi-structured post-interview.

We recruited Ph.D. students and postdoctoral researchers through our organization's Slack channels and through personal connections. We focused our recruitment on researchers who were actively involved in project ideation and could provide detailed accounts of their workflows. Demographic details of our participants can be found in Appendix A. We selected participants on a first-come, first-serve basis, and kept recruitment open during our study until we reached data saturation.

All studies were facilitated by two members of the research team and conducted virtually via Google Meet. Each study was around 60 minutes in length and was divided into 3 parts. To ground the study to real-world research projects, prior to the study, we asked participants to submit a project document—a running artifact with project updates, ideas, links, meeting notes, and more. This artifact was the central focus of **Part 1** of the study, where we

asked participants about how they sourced and developed their ideas, challenges with ideation, and their use of AI tools in research.

In **Part 2**, participants completed an interactive activity with a WoZ design probe. We leveraged Google Docs<sup>2</sup> and Perplexity AI<sup>3</sup> to power the “system interface” and “LLM-based intelligent support” aspects of our WoZ study, respectively. Specifically, Google Docs' real-time collaboration features enabled the study facilitator to populate content to the shared document in real-time from a different computer; and Perplexity AI's ability to search the web and provide functional links for academic papers could help us better understand the design opportunities of a research ideation system connected to an academic database. The probe's interface consisted of two side-by-side Google Docs. One of the Docs (the “ideation doc”) contained three initial research ideas we pre-generated by prompting Perplexity AI with participants' research interest from the recruitment survey as well as attachments of 1–3 of their recent publications (§B). The other Doc (the “output doc”) is used as a WoZ AI-ideation support system that displayed AI outputs. During the study, participants could trigger an “AI assistant” in the ideation doc by prefacing a request with a “!” command, after which a study facilitator (i.e., one of the authors) would manually forward the request to Perplexity AI and then paste its output back into the output doc. Participants spent around 25 minutes working with the probe to further develop a research idea. They had the freedom to build on one of the three ideas presented to them, use some combination of those ideas, or ignore them altogether. This activity ended when the participant wrote a short paragraph about their idea outlining aspects that may include a research question, methods, related work, and expected outcome.

We concluded with **Part 3**, an exit interview in which participants reflected on their experience using the design probe and how it differed from their typical workflows. They also shared feedback on the probe and discussed desirable ways of interacting with AI in ideation support tools.

After the formative study, the first author analyzed the transcripts with open coding followed by thematic analysis [8, 13] to identify key themes, focusing on common challenges participants faced. Tentative themes were discussed and iterated upon with the larger research team at weekly meetings to arrive at our final set of themes. The codebook can be found in Appendix C. Below, we describe the common themes we found in our analysis:

#### 3.1 Challenge 1: Expanding Initial Ideas to Concrete Projects

While prior works on ideation emphasized on supporting broad ideation and exploring diverse directions [36, 46], participants from our study also point to scenarios where they were already settled on a broad direction but struggle to further develop them into a concrete research project. Further, many said that they often do not lack interesting broad initial research ideas to pursue. However, they pointed to how it is often more challenging to expand their initial ideas due to uncertainty over how to concretely operationalize them. Participants also highlighted difficulties in measuring the

<sup>2</sup><https://docs.google.com/>

<sup>3</sup>Perplexity AI with Academic Focus has capability of conducting web search focused on scientific publications: <https://www.perplexity.ai/>



novelty of their ideas as they often have broad topics of interests but without specific idea facets (e.g., designs and methods, how to evaluate, specific impacts, and broader contributions) to solidify the project. P2 felt *“initially there’s a lack of clarity into what the proposed approach would look like so it’s hard to talk about it in a very concrete way. Because I am not sure what the core components of my approach will be. It’s hard to even express it in writing.”* Moreover, the process of narrowing down a broad range of thoughts to an operationalized project idea is inherently difficult. P3 expressed that *“narrowing it [initial idea] down to the one operationalizable research hypothesis...is what takes the maximum amount of effort.”*

Understanding existing literature deeply is crucial for evaluating ideas but poses significant challenges. As P2 noted, starting a project without a comprehensive literature review makes it hard to confidently discuss existing gaps. The overwhelming volume of related works make it even more difficult to ensure all relevant literature is covered (P4). Researchers often feel insecure about the novelty and impact of their ideas, relying on validation from peers and advisors, who may point out limitations in impact or feasibility (P5). This presents the need for literature-grounding and feedback to contextualize and iterate during the idea expansion process.

### 3.2 Challenge 2: Exploring Variations and Iterations of Ideas

When further developing an initial idea into a concrete research project brief, participants reported a common strategy of exploring variations and different framings. However, they faced challenges in evaluating multiple versions and iterations of their ideas. P6 often developed early prototypes of several ideas, setting aside those that do not work for later reconsideration. They often present multiple variations to peers and advisors to get feedback, which allowed them to identify blind spots and pivot between directions. P5 described that they will *“try to have more than one base idea, like two or three, just so that when I do get the opportunity to talk with my advisors...if one is a dead end we can talk about others as well...I try to create a few variations of each base idea.”*

This iterative process often involves extensive documentation, such as slide decks with updates and meeting notes, long-running documents with figures and summaries, and detailed research proposals with multiple versions and planning scenarios (P4, P5). Managing the context of multiple ideas stemming from the same seed idea is particularly challenging, as researchers struggle to track various developments and inspirations that are not immediately within the current scope (P1). The spatial fragmentation of notes and the need to link external resources further complicate the organization and evaluation of ideas (P1).

### 3.3 Challenge 3: Receiving Useful Support from LLMs

Reflecting on interacting with our WoZ prototype during the study and their own experience using LLM-based tools for ideation, participants felt that existing tools do not provide a useful level of specificity and literature-relevance in their feedback. For example, P5 felt that the AI often reiterates the same points without adding substantial new insights, failing to clarify concepts or identify specific details needed to refine ideas. P4 proposed that researchers

would benefit from AI providing fair and critical feedback, asking research questions, and offering critiques rather than simply affirming ideas. Furthermore, using AI for literature review tends to lead to reiterations of known information without providing the motivational context or the “why” behind certain approaches, which is crucial for ideation (P3).

Interaction paradigm with AI is another area of concern. While most participants ( $n = 5$ ) agreed that conversational interaction with AI is useful, typing natural language requests may not always be the preferred interaction method. Instead, participants desired alternative designs where AI support could be more active and in-situ within their artifacts and workspaces but without being disruptive (P2, P5). Additionally, AI suggestions often lack literature context or history, making it hard for researchers to see and understand the relevance of the suggestions (P7). In terms of agency and ownership, participants highlighted that they prefer using AI to inspire ideas, assist in study design, and spark further development rather than generating final outputs (P3, P5).

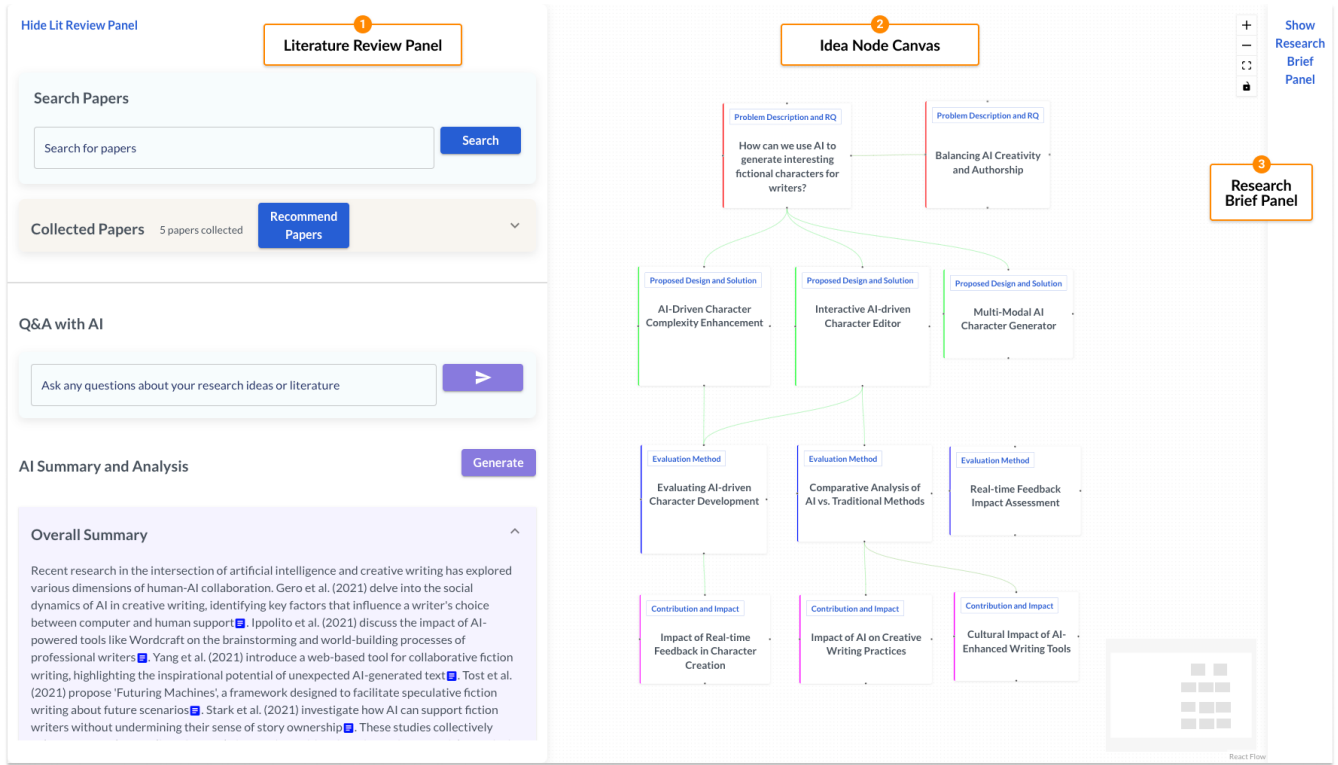
### 3.4 Design Goals

Based on the challenges identified by our formative study, we formalize three design goals corresponding to each of the challenges listed above:

- **DG1:** Better support for further developing an early stage research idea into a concrete research brief through reflection and evaluation. (**Challenge 1**). We observed the need for support with idea *expansion* rather than idea *generation*. We thus aim to provide better tools to make ideas more detailed and concrete so that researchers can operationalize them.
- **DG2:** Help users explore and evaluate different variations of an idea by analyzing their strengths and weaknesses and visualizing the connections between ideas facets (**Challenge 2**). Researchers found it difficult to manage idea iterations and fragmented notes while working in linear artifacts (e.g., documents, slides). New, non-linear interfaces may be needed to help researchers visualize and make sense of relationships between various idea components.
- **DG3:** Enhance the specificity and relevance of LLM support by grounding the response in relevant information in literature and focusing on inspiring users to develop their idea further (**Challenge 3**). Researchers often found LLM outputs too generic to be useful. This may be because LLMs did not encode the necessary domain-specific knowledge during training or require user guidance in accessing less well-represented knowledge in training data. We thus increase relevance and specificity by leveraging users’ existing collections of related literature to ground LLM responses.

## 4 IDEASYNTH

Based on our design goals, we present IDEASYNTH (Fig. 2), a research idea development system that helps scholars expand and refine their initial ideas. The system represents idea facets as linked nodes on a canvas, visualizing the logic relationships and forming a tree-like information structure for the user’s ideation context (DG2). IDEASYNTH also provides contextualized feedback at the node level and for the overall idea canvas based on the collected



**Figure 2: IDEASYNTH’s system UI screenshot.** On the left is the Literature Review Panel ①, where the user can search and add papers to the collection, and generate literature summary and analysis. In the middle is the Idea Node Canvas ②, where the user can externalize idea facets as nodes and explore and expand ideas in a tree-like hierarchical structure. The nodes are semantically zoomed out to display the idea canvas. Full UI and functionality of the node can be found in Figure 5. On the right is the Research Brief Panel ③, the user can select multiple nodes and generate a coherent research brief that describes the selected ideas facets (Expanded UI in Fig.6).

papers (DG3), enabling users to iteratively enrich ideas, explore alternatives, and connect facets to create a coherent research brief (DG1). We chose the research brief as the system’s output due to its utility in communicating and refining research projects, as it encapsulates core ideation activities such as defining a problem, method, and hypothesis.

#### 4.1 User Scenario and System Walk-through

In this section, we use a user scenario to walk through various features and design of IDEASYNTH.

**4.1.1 Literature-based broad ideation.** An HCI researcher is brainstorming a new project on *using AI to assist with fictional character development in creative writing*. She begins with a literature review, which helps refine her research question to: “*How can we use AI to generate interesting fictional characters for writers?*”. However, she faces challenges in narrowing the project scope to propose novel contributions and worries about prematurely focusing on a potentially underdeveloped idea.

**4.1.2 Decomposing initial idea into idea facet nodes and adding relevant papers.** The researcher decides to use IDEASYNTH to further develop and refine her research idea. She first created a “Problem Description and RQ” node on the canvas interface and entered her

research question (Fig.4.①). This node represents a “facet” of the idea. Then, using the paper search function, she adds the papers she collected during her literature review into the paper collection tray (Fig.3.①②). This way, IDEASYNTH could better understand her mental model and provide more contextualized feedback as she continue to develop her idea. She also uses the “Recommend papers” button to find additional relevant papers and expand her literature review (Fig.3.②). She can add as many papers as she wants as long as each paper’s content and generated summaries fit the LLM context window (around 300 pages for gpt-4-turbo).

To obtain an overview of the collected papers, she generates an AI literature summary and analysis (Fig.3.④⑤). The system provides a one-paragraph summary and a tailored analysis that draws connections between nodes on the canvas and specific sections of relevant papers. From the summary, she gains a deeper understanding of the challenges around evaluation from a prior work that also focused on AI writing assistants and explored mechanisms of using AI to drive fictional character development. Clicking on the paper icons in the summary opens the referenced papers in the collection tray for further reading. Beyond showing connections between prior work and the canvas, the literature analysis also provided actionable next step suggestions for expanding the ideas on the canvas. For example, she could drag a suggestion block onto

**1 Search Papers**

Search for papers Search

**2 Collected Papers** 5 papers collected Recommend Papers

**Social Dynamics of AI Support in Creative Writing**  
 Katy Ilonka Gero, Tao Long, Lydia B. Chilton  
 2023  
 International Conference on Human Factors in Computing Systems

**TLDR:**  
 This work asks when and why a creative writer might turn to a computer versus a peer or mentor for support, and uncovers three elements that govern a writer's interaction with support actors

**3 Q&A with AI**

Ask any questions about your research ideas or literature  
 How did these papers conduct their human evaluation? >

**AI Response** ✕

The human evaluations across the discussed papers primarily involved qualitative methods including interviews and participatory studies. 'Social Dynamics of AI Support in Creative Writing' [1] conducted interviews with 20 creative writers to understand their preferences for AI versus human support. Similarly, 'Creative Writing with an AI-Powered Writing Assistant' [2] used interviews and journals from 13 professional writers to gauge the impact of NLG tools on their creative processes. 'AI as an Active Writer' [3] implemented a web-based tool in a collaborative writing setting and collected feedback from writers through a 'finish each other's story' approach, allowing for dynamic interaction and feedback on AI-generated text. These methodologies highlight the reliance on direct user feedback to assess the utility and acceptance of AI in creative writing, aligning well with the analysis of real-time feedback mechanisms in your proposed AI character development tool.

**AI Summary and Analysis** Generate

**4 Overall Summary**

Recent research in the intersection of artificial intelligence and creative writing has explored various dimensions of human-AI collaboration. Gero et al. (2021) delve into the social dynamics of AI in creative writing, identifying key factors that influence a writer's choice between computer and human support [1]. Ippolito et al. (2021) discuss the impact of AI-powered tools like Wordcraft on the brainstorming and world-building processes of professional writers [2]. Yang et al. (2021) introduce a web-based tool for collaborative fiction writing, highlighting the inspirational potential of unexpected AI-generated text [3]. Tost et al. (2021) propose 'Futuring Machines', a framework designed to facilitate speculative fiction writing about future scenarios [4]. Stark et al. (2021) investigate how AI can support fiction writers without undermining their sense of story ownership [5]. These studies collectively enhance our understanding of AI's role in creative writing, focusing on its potential to assist in generating ideas, maintaining writer autonomy, and fostering engagement with future-oriented narratives.

**5 Analysis and Suggestions**

**Exploring Human-AI Collaboration in Fiction Writing**  
 AI as an Active Writer: Interaction Strategies with Generated Text in Human-AI Collaborative Fiction Writing [1]

**Key Section:** This research introduces a human-AI collaborative tool that allows writers to interact dynamically with AI-generated text, providing opportunities to edit, summarize, and regenerate content. The study reveals that writers often find inspiration in unexpected AI-generated text, which can be particularly useful in developing novel character traits and backstories.

**Connection to Ideas:** This study's findings support the feasibility of using AI to generate interesting character backstories by showing how AI can contribute creatively in a collaborative setting. It suggests that AI can serve as a co-creator, potentially leading to innovative character development.

**Next Steps** (Drag to canvas to create nodes based on suggestion)

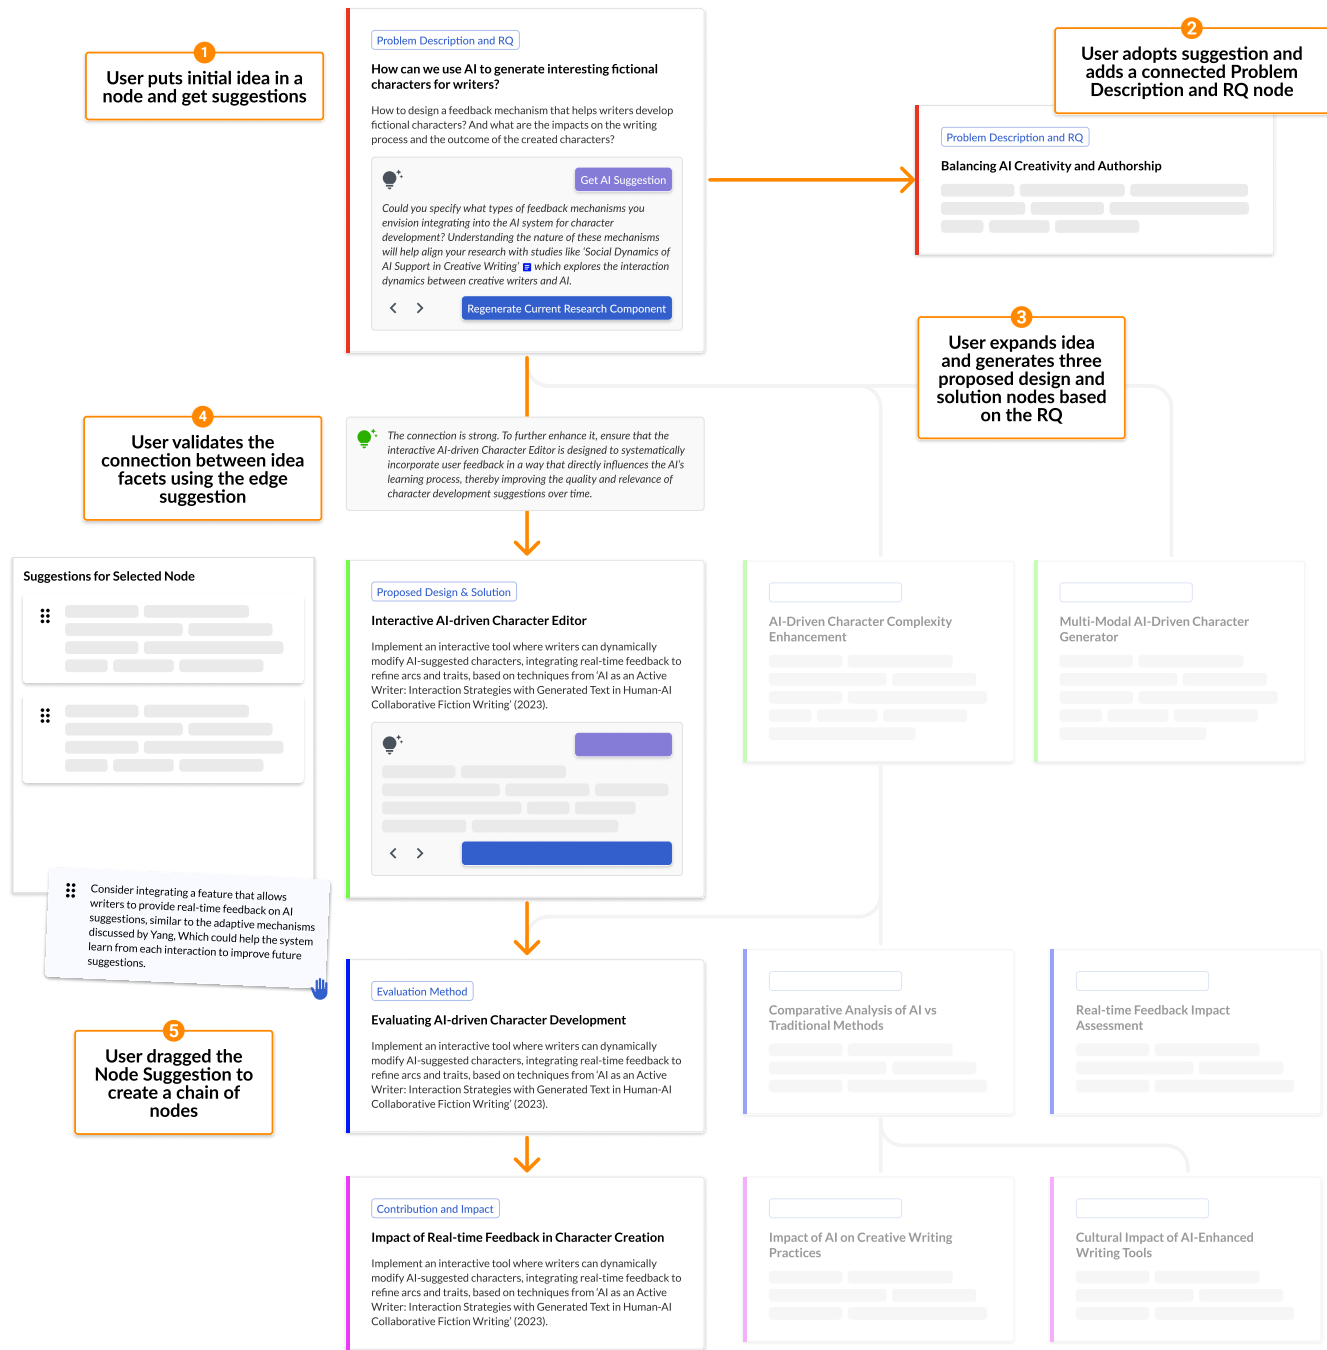
- Develop a prototype AI tool based on these findings that focuses specifically on generating and refining character backstories in collaboration with human writers.
- Test the tool with a group of writers to evaluate its effectiveness in supporting creative character development and gather feedback for further improvements.

**AI-Enhanced Creativity in Speculative Fiction**  
 Futuring Machines: An Interactive Framework for Participative Futuring Through Human-AI Collaborative Speculative Fiction Writing [2]

**Figure 3: Literature Review Panel.** This is a scrollable side-panel. The figure expands the panel to two columns (flow from left to right) to display all components. The user can ① search and add papers, and use Recommend Papers to expand the collection ②. They can ③ pose targeted questions about details in the paper or inquire about their idea nodes. On the right column, the user can generate an overall literature summary based on the paper collection ④, and an analysis that connects key paper sections to the idea nodes on the canvas ④. In our evaluation, the baseline system contains only the functionalities in the left column (Section 5.1.2).

the canvas to generate a chain of nodes that represent a new idea. For now, she decides to first explore and expand upon her initial research question.

**4.1.3 Further developing the idea by exploring facet node variations.** Returning to her initial RQ facet node, she writes down two more detailed research questions in the node: "How to design a feedback mechanism that helps writers develop fictional characters?" and



**Figure 4: Node Canvas UI demonstrating the user scenario workflow.** The HCI researcher starts by entering her initial idea in a blank node ①. Then, she uses Get AI Suggestions to receive feedback. IDEASYNTH prompts the researcher to clarify the problem description and scope, and also proposes an alternative research question. She adopts the suggestion and generate another node ②. To further expand on her idea, the researcher explored other idea facets and generated three nodes on proposed design and solution ③. She evaluated the nodes and used the edge to validate the connection between the design and the problem ④. To examine the particular node, the researcher expands the left arrow menu to view node-based literature analysis that connects the idea facets to key sections from the collected papers. The node literature analysis also offers actionable suggestions, which she uses to drag and create a chain of subsequent nodes to ideate on the evaluation and contribution ⑤.

“What are the impacts on the writing process and the outcome of the created characters?” (Fig.4.①). She uses “Get AI Suggestions” to receive feedback. IDEASYNTH first presents a clarifying question to ask the researcher to specify what aspects of the writing process she would like to focus on, with examples such as *initial character creation* or *continuous character development throughout writing*. The suggestions also reference relevant papers with findings on how writers seek feedback.

Based on this suggestion, she iterates on her research question to center around “*character creation aligned with a plot narrative vision in early writing*”. She also reads a suggested paper in-depth and gains insights on existing writer feedback mechanisms. She then toggles to view another AI suggestion. This time, IDEASYNTH proposes an alternative research question to consider *how to balance AI creativity and authorship*, a key issue illustrated by a paper she had previously added to the system. The researcher finds the research question inspiring as she has previously not considered the trade-off in authorship. This time, she uses “Generate Alternatives” feature to create new nodes that explore variations of this suggested research question. She incorporates the generated node on *investigating the author’s creative control and authorship* and connects it with her first node (Fig.4.②).

**4.1.4 Creating connections between nodes to explore their relations.** She then uses the downward arrow menu at the bottom of the two “Problem Description and RQ” nodes to explore other idea facets (i.e. “Proposed Design and Solution”, “Evaluation Method”, and “Contribution and Impact”). First, she uses IDEASYNTH to generate multiple “Proposed Design and Solution” nodes with the prompt “*interactive AI system for character building with author feedback*” (Fig.4.③). The generated nodes are color-coded by idea facet and linked to their original nodes. The researcher evaluates each design: she finds the first two compelling—one is a character editor with real-time AI feedback, and the other enhances character complexity by analyzing literary databases. She finds the third, a multi-modal character generator, interesting but beyond the project’s scope and deletes it. She then examines the connecting edges between nodes to validate their relevance to the original research questions, checking the logical strength of the node connection and receiving suggestions for improving it (Fig.4.④). The researcher sees that the connections to the two proposed designs are strong, and adopts one of the edge suggestions by further specifying how the AI learns and incorporates writer’s feedback.

The researcher can also use the left arrow menu to review an automatically generated literature analysis specific to the current node (Fig.4.⑤, Fig.5.⑥). Similar to the literature analysis on the side panel, this analysis extracts relevant paper sections that connect to the existing idea facet and proposes next step actions. This node-level analysis provides more granular feedback on each idea facet, helping the user iterate on the idea and receive inspirations for expanding their ideas based on prior works. The researcher finds the suggestion to evaluate AI character creation across different writing styles compelling. She drags this suggestion to create a chain of “Evaluation Method” and “Contribution and Impact” nodes, and links them to existing idea nodes. (Fig.4.⑤).

At this point, she wonders how previous works conducted their human evaluations, so she types a targeted literature query in

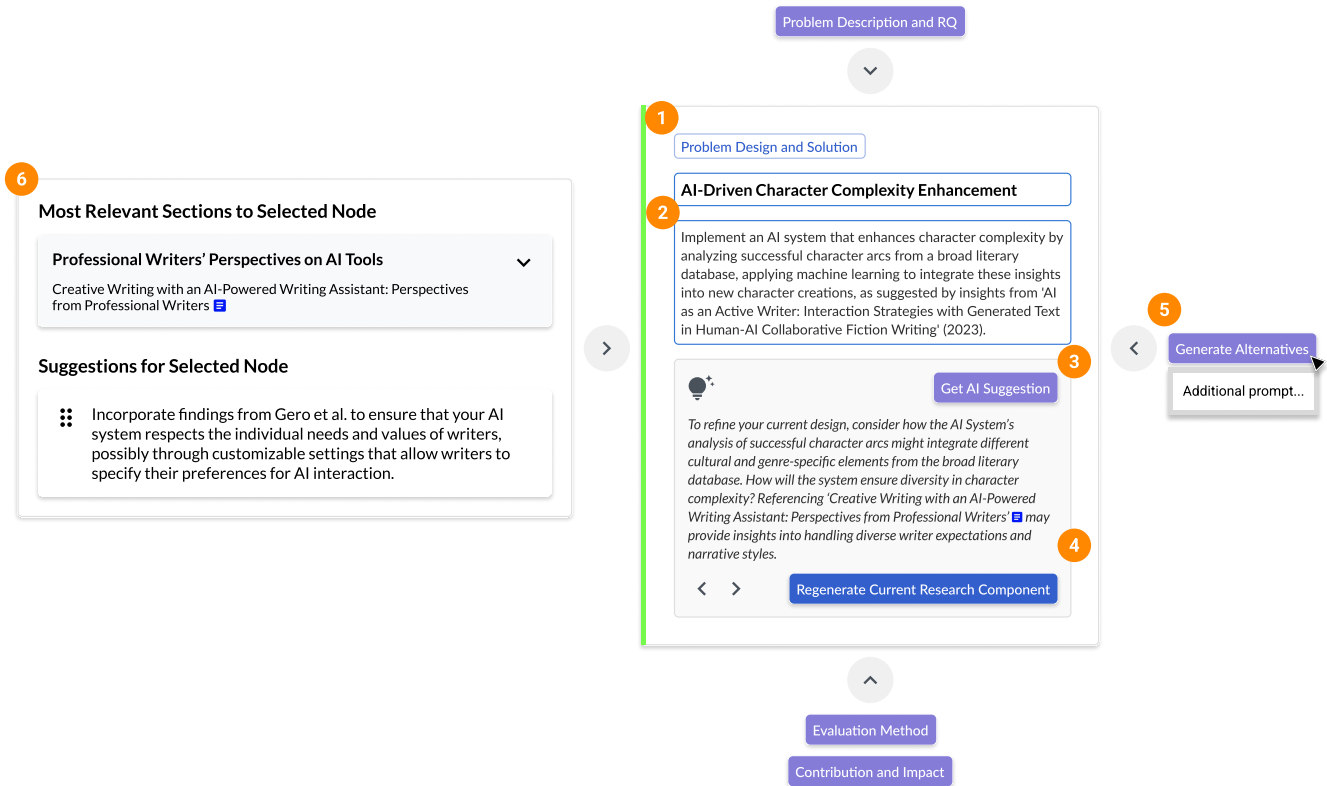
the “Q&A with AI” interface and receives a response (Fig.3.③). IDEASYNTH synthesizes the response based on the evaluation methods it summarized from each paper. The response lists out the human evaluation method from the collected papers: one paper included an interview with 20 writers on their preferences for AI versus human support, while another paper developed a web application and collected feedback from writers through a “finish each other’s story” approach, among others. The researcher notes down these literature-driven insights as she progresses in her project ideation.

**4.1.5 Generating research briefs by exploring different node compositions.** Using IDEASYNTH’s tools, the researcher refines node content with AI suggestions, deepens her literature understanding through summary and Q&A capabilities, and broadens her research ideas by exploring alternative approaches via distinct idea facets, ultimately resulting in a more well-defined project. After decomposing her initial idea into different facet nodes and generating variations for each, she feels she has sufficiently explored the problem space. She then selects a *path of idea facet nodes* that encapsulates the research project to generate a *research brief* on the right-side panel. The brief compiles these facets into a cohesive document, with the corresponding nodes highlighted (Fig.6). The researcher shares the research brief with her collaborators to further discuss this project idea. She can also return to the canvas to develop multiple clusters of ideas, creating multiple research briefs to continue ideation. The canvas space, then, functions as a mind map and preserves the ideation context within this broad topic.

## 4.2 Node-based Idea Canvas

IDEASYNTH is built around a node-based canvas, where each node represents a distinct idea facet. We specifically define four types of idea facets: *Problem Description and Research Question*, *Proposed Design and Solution*, *Evaluation Methods*, and *Contribution and Impact*. These idea facets are chosen as prior works have commonly identified them as critical components that represent the research idea in various formats, such as research paper [10, 26, 39], research proposal [52], or abstract [3], across different scientific domains. We omitted other potential facets such as evaluation results, or limitations as they cannot be concretely described in early research ideas without execution. Users can select the different idea facet types via a drop-down-menu (Fig.5.①) as they see fit.

The node representation decomposes the research idea into a digestible format for iteration and evaluation. Each node on the canvas offers specific functionalities designed to facilitate interaction. Users can freely edit the node title and content (Fig.5.②). They can also link nodes together to form logical connections between them. Each edge is color-coded to represent the strength of the connection (color gradient: red=weak, green=strong) and can be expanded to display text suggestions to enhance the coherence between the connected nodes (§1.5). Users can manually create idea facets by right-clicking on the canvas. Or they can use IDEASYNTH to generate new idea facets by adopting node suggestions (Fig.5.④⑥) or manually selecting a facet type (or the same type via “Generate Alternatives”) at the node menu (Fig.5.⑤). A new node is generated based on the context of the current node and any directly linked nodes. The user can use an additional prompt to steer the



**Figure 5: Idea Node UI.** In IDEASYNTH, idea facets are represented as nodes. The user can select the type of idea facet ① using a drop-down menu, the left border is color-coded based on the facet type. The user can enter a node title to represent the key theme of this node, and describe the idea in more details in the node content text field ②. The user can view node-level AI suggestions in the form of clarifying questions, idea alternative, and idea expansion ③. They can manually adopt the suggestion to use the action button to let IDEASYNTH execute ④. On the four sides of the node, there are expandable menus. User can use the purple buttons to generate idea facet nodes in a hierarchical structure. The user can optionally enter a prompt to steer the generation ⑤. The left expandable menu reveals the node-level literature analysis, which connects the idea facet to relevant paper sections and provide suggestions ⑥.

generation. Informed by our formative study finding, IDEASYNTH allows researchers to start their ideation with any node type, giving the user control to flexibly navigate the ideation process to diverge, converge, and iterate on any idea facets based on their needs. IDEASYNTH also features semantic zooming to hide node details when users zoom out. The tree-like vertical node structure reflects the progression of a research concept from initial stages to a more refined, comprehensive research brief.

After sufficient idea expansion and exploration, users can select multiple nodes and generate a research brief (Fig.6, §1.6). This brief integrates the content of each selected node while maintaining the hierarchical connections between idea facets, ensuring that the decomposed ideas and their relationships are preserved. By doing so, the system helps users translate structured representations of ideas into a text-based one for further development or presentation.

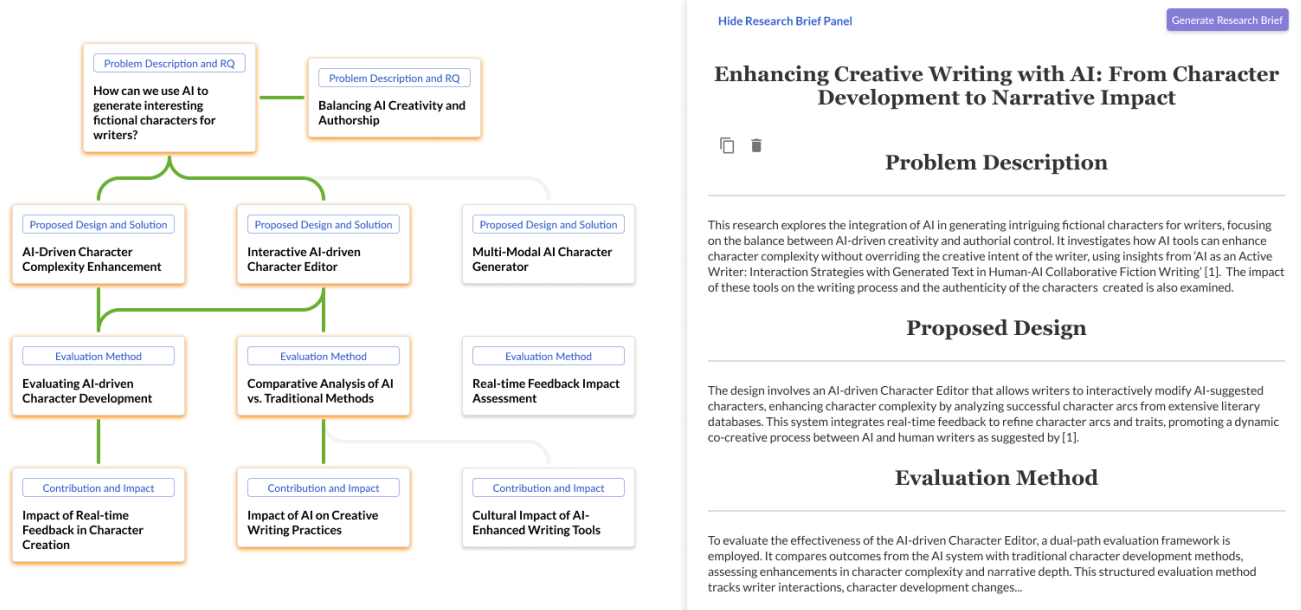
### 4.3 Methods for Generating Literature-Grounded Feedback

To provide contextualized feedback for the user's ideas like shown in the scenario, IDEASYNTH implements an LLM-powered system that grounds its generation based on the collected papers (DG3). The system allows users to construct a paper collection through a search bar (Fig.3.①) connected to the Semantic Scholar API<sup>4</sup>, which retrieves the details and metadata (title, authors, TL;DR, abstract, etc.) of the paper results based on a search query. The search result contains real papers from an open scientific literature dataset. Once the user collects an initial set of papers, they can use the "Recommend Paper" function (Fig. 3.②), which uses the Paper Recommendations API<sup>5</sup> to suggest more relevant papers based on a literature knowledge graph using the existing collection as positive examples, as per [12].

<sup>4</sup><https://api.semanticscholar.org/api-docs/>

<sup>5</sup>[https://api.semanticscholar.org/api-docs/recommendations#tag/Paper-Recommendations/operation/post\\_papers/](https://api.semanticscholar.org/api-docs/recommendations#tag/Paper-Recommendations/operation/post_papers/)





**Figure 6: Research Brief Panel.** The user can select multiple nodes on the canvas and generate a research brief on the right panel. The research brief takes content from idea facets and connections and transform them into coherent writing. When focused on the brief, the included nodes and edges are highlighted on the canvas. The user can create multiple research briefs in separate tabs.

To ensure that IDEASYNTH generates useful suggestions grounded by the collected papers, each paper added by the user goes through additional processing. The system first scrapes the paper PDF file (if open-access) and extracts the full text of the paper using GROBID [1], an open source library specialized to extract information from technical and scientific publications. We then provide the full text of the paper and perform retrieval-augmented generation (RAG) using gpt-3.5-turbo to query and summarize specific facets of the paper, including the research question, the proposed design and solution, the evaluation method, the contribution and impact, and the limitations and future works (§I.2). If the paper full text is not available, the system falls back to the abstract and TL;DR summary to use as context. We use a lower-tier LLM for RAG with paper-length text to reduce cost and response delay. To mitigate LLM hallucination, we delegate the task of finding related paper to their ideation to the researcher, with the support of algorithm-based search engine (i.e. Semantic Scholars) that yields real paper information. The user is able to collect any interesting papers relevant to their ideation topic, including works from their own discipline or other domains. IDEASYNTH only processes and summarizes paper information from the list of papers added by the user using IDEASYNTH and never refer to papers beyond this collection to avoid false source of information.

Based on the paper collection enhanced with summaries of specific facets, IDEASYNTH dynamically incorporates relevant papers or facets to guide generations for idea refinement (DG1) and explore the information space (DG2). IDEASYNTH generates LLM suggestions to scaffold the research ideation process at two abstraction levels. Overall, at the node canvas level, IDEASYNTH allows the user to generate a one-paragraph literature review summary based on

all the collected papers (Fig.3.④, §I.8). The system additionally analyzes the literature using the current ideation progress, including all the node and edge contents on the canvas as context (§I.9). The literature analysis generates several comparisons between collected papers and idea facets, each identifying the key section (i.e. facet) in a particular paper and highlighting the connection to the idea facets (Fig.3.⑤). IDEASYNTH also provides actionable suggestions for the user to iterate on existing idea facets or generate new ones, such as dragging the suggestion element onto the canvas to automatically generate a chain of nodes that executes on the suggestion (§I.4). Additionally, the Q&A interface allows users to ask targeted questions about specific details in the collected papers or inquire about ideas represented on the canvas (Fig.3.③, §I.7).

As for the idea nodes, IDEASYNTH offers three types of node suggestions: idea iteration, idea alternatives, and idea expansions (Fig.5.③, §I.3). To help refine and iterate on the user's idea (DG1), the system asks clarifying questions to prompt deeper reflections and more detailed descriptions, encouraging users to address potential challenges based on existing works. To facilitate idea exploration (DG2), IDEASYNTH suggests alternative variations for the current idea facet to broaden the user's exposure to research questions, designs, evaluation methods, and more. It also recommends ways to expand *between* idea facets, such as proposing solution designs for a given research question or deriving research questions linked to an evaluation method. Users can manually adopt these suggestions or let IDEASYNTH generate them automatically, with options to customize the prompts. For each node, the system generates a node-specific literature analysis based on its content and type, leveraging relevant collected papers and suggesting actions for further iteration or expansion (Fig.5.⑥, §I.10). Through

these methods, IDEASYNTH dynamically grounds LLM generation in the user’s collected literature at both the canvas and node levels. This provides users with actionable and contextually relevant suggestions for refining, expanding, and iterating on research ideas.

#### 4.4 Implementation Details

IDEASYNTH is implemented as a web application. The backend was implemented in Python using Flask. The frontend used the React framework in Typescript. We used LangGraph<sup>6</sup> to instantiate an LLM instance with memory of past interactions and generations. We used gpt-4-turbo with default parameter values supported by LangGraph (temperature= 0.7). However, IDEASYNTH is designed to be model-agnostic and our key contribution is on the human ideation workflow enhanced by incorporating LLM assistance. We adopted common prompt strategies like chain-of-thought and few-shot prompting, which yielded higher quality results from our testing. All LLM prompts used in the system are included in Appendix I.

### 5 Study 1: Comparative Lab Study

To evaluate the usability and utility of IDEASYNTH, we first conducted a within-subjects laboratory study where participants developed research ideas using IDEASYNTH and a baseline system. We formulated the following research questions:

- **RQ1:** How does the node-based canvas interface allow users to explore and connect different idea facets?
- **RQ2:** How does the externalization of idea facets help the expansion and refinement of initial research ideas?
- **RQ3:** How do LLM-generated suggestions grounded on collected papers help the user understand the literature space and evaluate their research ideas?
- **RQ4:** How does the user experience of research ideation using IDEASYNTH compare to using a standard LLM-based tool?

#### 5.1 Study Design

**5.1.1 Participants.** We recruited 20 researchers (11M, 8F, 1 non-binary) who are pursuing or have completed their Ph.D. degrees (i.e., graduate students and postdocs) in Computer Science through posting recruiting messages in Slack channels of several research institutes. We targeted scholars who were actively engaging in research ideation and project development. In total, twenty participants were recruited, of which 12 were in the field of HCI and 8 in the field of NLP. Thirteen participants have 2–5 years of research experience, six have more than 5 years of research experience, and one has less than 1 year of experience. From the entry survey, all participants described literature review and discussion with peers/advisors as a part of their research ideation process. Twelve participants mentioned writing a research plan or proposal as a part of their ideation, utilizing different tools for documenting ideas and related work notes (e.g. Google Doc, Notion, Obsidian, etc.) Participant demographics can be found in Appendix D.

**5.1.2 Conditions.** The study adopts a within-subjects design with two conditions: IDEASYNTH and a strong baseline. Simulating a

common practice reported by participants in the formative study, the baseline system implements a word editor<sup>7</sup> where participants can develop their research ideas. To further control the task goal, the baseline editor provided an initial writing template containing section headings corresponding to the idea facet categories in IDEASYNTH. The baseline condition also contained a subset of features in IDEASYNTH. Specifically, a literature review panel (Fig.3.left column) with identical paper search and recommendation features to support literature understanding. The baseline system also provided LLM assistance through the same Q&A interface. The user can ask targeted inquiries regarding the collected papers or free-form questions about their research ideas similar to using tools like ChatGPT. The two systems share identical underlying LLM setup and system prompt. The LLM in the baseline is supplemented with the same paper metadata and the writing editor content as context. Additionally, the baseline contains an “AI Assist Writing” button in the editor, which will generate a revised research brief based on the user’s writing and attached it at the end of the editor. This function uses the same prompt (§1.6, but with editor content as context) as the “Generate Research Brief” function in IDEASYNTH that transforms idea facet nodes into a coherent writing.

**5.1.3 Tasks.** Each participant was asked to develop research ideas based on two provided broad topics. For HCI researchers, we identified a main theme of human-AI interaction and two task topics: “AI-augmented Tool for Creativity” and “Human-AI Collaboration in Data Analysis”. For NLP researchers, we chose the theme of LLM evaluation, with the two task topics being “LLM-as-a-Judge” and “Biases in LLMs”. The task themes cover a wide range of concepts and specific domains to provide sufficient freedom for researchers in different areas to develop research ideas. We did not let participants ideate on their existing research projects as there might be variance in the stage of ideation between projects and participants.

**5.1.4 Procedure.** We control for the learning and ordering effects by counterbalancing the system condition and task topic order (2×2) independently for the HCI participant group ( $n = 12$ ) and the NLP participant group ( $n = 8$ ). The participants were first given a brief introduction on the goal of the study and a task instruction sheet (Appendix E). For each task, they first watched a short tutorial of the assigned system (IDEASYNTH or baseline), then spent 30 minutes to develop their research idea on the assigned task topic. When there were 5 minutes left, participants were reminded to make final changes to their idea and converge on a research brief as the task deliverable. After each task, participant filled out a survey to evaluate the quality of their research idea and record their perception about the use experience and the task workload. At the end of both tasks, we conducted a semi-structured interview where participants reflected on their overall experience and compared the ideation process using both prototypes. The study was conducted remotely using video conferencing software. Each session lasted around 90 minutes, and participants were compensated for \$75 USD. The study was approved by our internal review board.

<sup>6</sup><https://github.com/langchain-ai/langgraph>

<sup>7</sup><https://quilljs.com/>



## 5.2 Measures and Analysis Procedure

We collected both users' perception and behavioral data to conduct a comprehensive analysis on the usage of IDEASYNTH. Quantitatively, we analyzed the post-task survey responses. The survey consists of 7-point Likert scale questions on the user's perception towards their level of literature understanding, idea refinement, and idea expansion they were able to achieve during the task. Additionally, the survey includes questions from NASA-TLX [22] on perceived workload, trust, and five questions asking the participants self-assess the quality of their final research idea in five dimensions: novelty, impact, specificity, feasibility, and relevance, defined by prior work [17]. We pre-selected four key metrics from the survey for statistical testing using a MANOVA test to account for multiple comparisons. The four key metrics were selected because they were either common and important aspects when evaluating research proposals (i.e., *impact* and *feasibility*) or closely related to our design goals based on our formative study (i.e., *confidence in literature understanding* and *sufficiently explored and evaluated alternatives*). The full list of survey questions and results are presented in Appendix G. In addition, we analyzed the interaction logs from the sessions to measure the quantity and types of actions taken (e.g. add a paper to collection, get AI suggestion, user typing, etc.) and capture behavioral differences. We used Wilcoxon signed-rank tests for pairwise two-condition comparisons for independent Likert scale survey questions (Fig. 8) and T-test for temporal data. For the qualitative data, the first author performed thematic analysis on the interview transcripts and iterated on the themes with the other authors through discussion. We report the synthesis of our themes, denoting HCI participants from H1–H12 and NLP participants from N1–N8; our full themes can be found in Appendix F.

## 5.3 Results

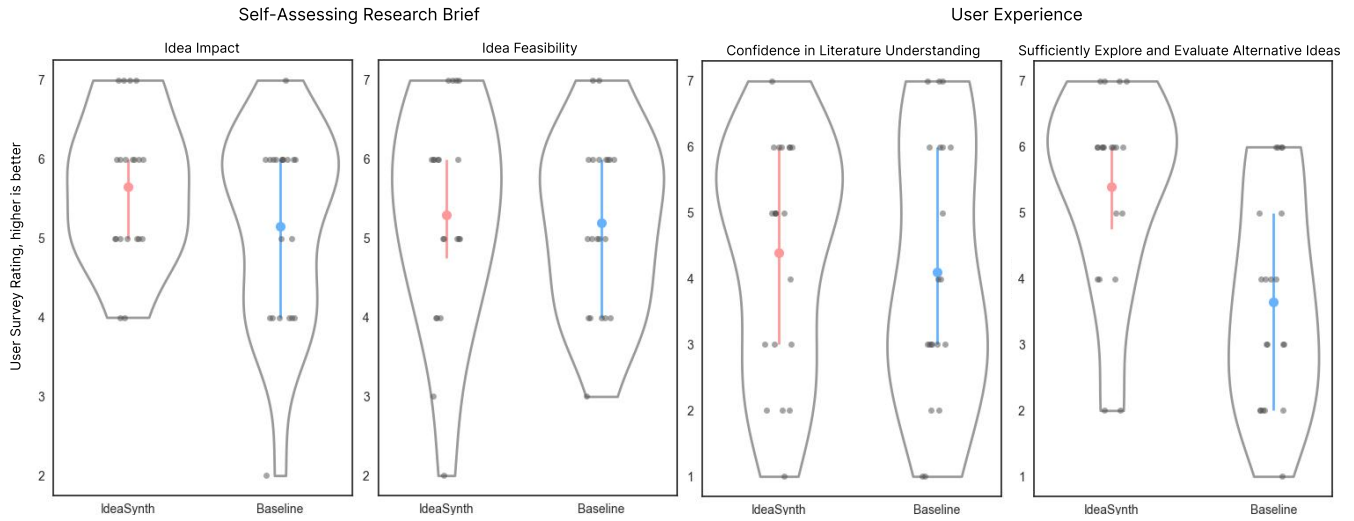
**5.3.1 IDEASYNTH helped users better explore alternative ideas and navigate the idea space (RQ1).** Based on the survey response, participants felt they more sufficiently explored and evaluated alternative versions of their research ideas in IDEASYNTH ( $mean = 5.40, SD = 1.50$ ) than the baseline ( $mean = 3.65, SD = 1.60, p < 0.01$ ), as shown in Fig 7. H3, H10, H12, and N3 stated that the baseline system, due to its editor-like interface, did not encourage exploring alternative ideas but pushed them to develop their initial idea in depth. However, the lack of exploration could lead to idea fixation (H3). In contrast, fourteen participants mentioned that the support for suggesting and generating alternative nodes in IDEASYNTH led them to explore more adjacent and complementary idea facets. H1, H2, H6, and H11 attributed the enhanced process of exploration to the structured layout with horizontal levels of nodes for each idea facet, visualizing the connections between ideas. N6 thought IDEASYNTH help them *"find the connection between two nodes and this feature can help you discover something you didn't recognize before."* H7 demonstrated the process of expanding and developing an abstract idea into a concrete research brief. They initially created a "Contribution and Impact" node with a general idea to help people understand their health data. Utilizing node-level literature suggestions, H7 generated multiple separate paths of nodes with other complementing facets. Then, H7 connected idea facets that were interesting to them to synthesize their proposed

solution in a research brief, including designing an AI health assistant across multiple sensing devices and integrating staged-based model to structure interactions across data collection, integration, reflection, and action.

The ability to evaluate multiple generated idea facets based on collected papers broadened participants' perspectives (H1, H12). H1 claimed that *"since [IDEASYNTH] actually generated new ideas and suggested some papers based on some of my prompts, it actually broadened my understanding on the area."* Sometimes, the generation of alternative ideas inspired users to extend their ideas with details they haven't considered. H8 reflected that through using IDEASYNTH, their research idea *"grew into providing more details or sometimes, an angle that I hadn't thought about."* Participants also reported going through a converging stage during idea exploration. N3 felt like the alternatives generated by IDEASYNTH were *"more out of my scope and out of my expectations. So that's good. Not necessarily means that those ideas are good but still it elicits some new thoughts in me. That's also indirectly helping me to find some good directions."* H2 expressed that they were more confident in their own idea after generating and evaluating alternatives, pruning irrelevant nodes, and converging on the scope of the idea. For participants who ideated on unfamiliar topics (H9, H4), the AI suggestions helped them explore and better understand the information space. For example, H9 began an ideation session on using AI for painting, an unfamiliar topic for them. By searching related works and reviewing literature summaries, they gained insights into how AI-driven approaches influence traditional artistic processes. H9 organized their ideas into two branches, exploring both the benefits and risks of AI-assisted painting. Using insights from the collected papers, they generated "Problem Description and RQ" nodes, evaluated artists' perceptions of AI, and iteratively refined their exploration. This process culminated in a tree-structured canvas, from which H9 selected key facets to craft the final research brief.

Despite the benefits in promoting exploration and ideation, IDEASYNTH has some limitations. Four participants (H1, H12, N5, N6) noted that the system's node-based layout displayed dense information that increased their mental workload. For example, H12 expressed that navigating through numerous suggestions required extra effort to sift through and prioritize relevant ideas and occasionally felt *"a bit overwhelming"*. On the other hand, this sentiment was not reflected in the survey results (IDEASYNTH :  $mean = 4.35, SD = 1.73$ , Baseline :  $mean = 4.55, SD = 1.64$ , Fig.8), suggesting that balancing the costs and benefits of the node-based layout and combined with other features of IDEASYNTH, the overall mental workloads were comparable. Finally, one of the participants (N5, an algorithm researcher), preferred the traditional writing-centric approach, such as using physical pen and paper, which might be specific to research areas like algorithm design where the tactile process of sketching out ideas and deriving mathematical equations felt more intuitive. These observations suggest that, while IDEASYNTH supports ideation effectively, its design could be optimized to balance idea generation with cognitive load, as well as consider researchers' needs in different domains. We further discussed this idea in the Discussion and Future Work section.

**5.3.2 IDEASYNTH helped users expand and iterate on ideas (RQ2).** Almost all participants ( $n = 17$ ) reported that they were



**Figure 7: Key Survey Responses** Participants generated research briefs after 30 minutes of ideation with each of the two systems. Participants self-assessed their research brief and rated their user experience and perceived workload after each task. We pre-selected four key measures and conducted a MANOVA test to correct for multiple comparisons, identifying a statistically significant difference ( $manovaF = 3.30, p < 0.05^*$ ) between IDEASYNTH and the baseline on the combined dependent variables of Impact, Feasibility, Literature Grounding, and Sufficient Exploration. We separately report full survey questions and results in Appendix G.

able to more effectively expand their initial research idea with additional details with IDEASYNTH. This was supported by the survey responses (IDEASYNTH :  $mean = 6.05, SD = 0.89$ , baseline :  $mean = 4.45, SD = 2.04$ , Fig.8). Many participants ( $n = 11$ ) attributed this to being able to break down their ideas into multiple facets using the node representation. Upon connecting facets to compose them, H7 felt more motivated to expand their idea: “once you have that initial idea to start with, that’s when I can see potential for a new node, ... I can imagine each node being a research question and other nodes which are connected to it through edges,...[it’s] motivating [as] that’s when I would actually pay attention to the connections.” One instance that demonstrates IDEASYNTH’s assistance in expanding and iterating on ideas is H8’s exploration of using AI to generate fictional characters’ key story moments based on a general description. By evaluating IDEASYNTH’s node suggestions, H8 iteratively expanded the scope of their research idea to focus on elevating narrative flexibility in creative writing, envisioning an interactive character editor powered by LLM-driven complexity enhancement. Additionally, alternative nodes generated by IDEASYNTH helped H8 identify and integrate the balance between AI creativity and human authorship as a critical “Problem Description and RQ” node—an aspect of the research they had not initially considered but later incorporated into their research brief.

Similarly, N6 leveraged IDEASYNTH to refine their initial idea of reducing cultural bias in LLMs. Building on IDEASYNTH’s suggestions, they connected two “Proposed Design and Solution” facets: one focusing on semantic data augmentation and the other on a cultural bias feedback loop. Through IDEASYNTH, N6 identified the opportunity to quantify the effectiveness of data augmentation by integrating an iterative feedback loop that combined both facets. Inspired by this connection, N6 iterated on the node contents to detail

a human evaluation system for dynamically tuning data parameters based on feedback, adding depth to their research direction. Comparing IDEASYNTH to the baseline system, H3 said that IDEASYNTH “was more targeted in a sense that they [ideas] were separated into the specific things that I had to include in the research brief so I can iterate on one component much easier,” whereas the baseline “just provided general summary that was very general.”

The majority of participants ( $n = 13$ ) also found that the AI suggestions for each idea facet in IDEASYNTH were helpful for iteratively refining their ideas. N2 felt that the suggestions “motivated [them] to think of more details” from their original idea description via clarifying questions. H5 experienced “more control on iteration... after I learn more about this topic, I can add more details... it’s definitely easier to focus on a specific aspect, easier to iterate on that specific node.” H9 also reflected that the node connections were helpful for iterating on an idea as the suggestions are “constantly poking holes and connections, and finding ways to make more connections...while also strengthening my proposal.” In contrast, participants felt that AI suggestions in the baseline system were more like summarization rather than ideation (H8, N4, N8), similar to interacting with tools like ChatGPT (H5, H8, N4, N6), despite using the same underlying LLM and system prompt in both conditions.

**5.3.3 Similar literature search behaviors but IDEASYNTH’s facet-specific summaries and suggestions can surface unexpected connections to prior work (RQ3).** In general, participants ( $n = 12$ ) found the literature review panel in both conditions helpful for exploring related works and gaining an overview of the papers to provide a foundation for research ideation. We did not find a significant difference in users’ confidence in understanding the literature in the post-task survey across two conditions. The behavioral logs also showed that participants spent similar

percentages of their task duration searching and reading paper information (IDEASYNTH = 9.75%, Baseline = 11.7%), and we did not identify a significant difference in the number of papers collected (IDEASYNTH = 5.60, Baseline = 6.05,  $p = 0.63$ ). This is unsurprising as both systems included the same literature search features.

However, some participants ( $n = 7$ ) found IDEASYNTH, with the addition of literature summary and analysis generation at both the canvas-level and node-level, more helpful compared to the baseline. IDEASYNTH summarized the collected papers based on the defined idea facets. H7 found it helpful to seek literature support for ideating specific facets in their ideas: *“I think that’s what I need when I’m brainstorming or ideating. I am confident enough in my methods to not need help...but I think even then this is summarizing motivation...that was pretty helpful.”* N4 commented that *“the option to generate a summary of various paper and...being able to see some connection between the two papers is very interesting and not something that I saw before.”* Similarly, N7 felt that *“[IDEASYNTH] was much better at not just giving you a summary of papers ... it’s... connecting the dots for you.”*

IDEASYNTH additionally presented relevant paper information for individual nodes. H4 commented that *“just being able to explore different types of suggestions was super useful...be able to toggle through and see how it was bringing up some of the literature that I had collected.”* H4 felt that the literature-grounded AI suggestions helped them maintain context of relevant works and prevented idea fixation: *“It’s mostly to offer suggestions and things that wasn’t immediately thinking about. It just helps to prevent myself from pigeonholing because it will point out different papers that I couldn’t just think of referencing immediately that strengthen the research idea...Or talking about related work that I could build upon instead of starting from scratch.”*

Yet, some participants ( $n = 6$ ) shared that the features in IDEASYNTH and the baseline do not replace the literature review process of actually reading the papers. H4 commented that *“I trust the [literature] summaries, but I don’t trust that it’s done an exhaustive search of papers yet.”* Participants found the LLM-generated suggestions based on collected papers can sometimes still be too broad (H6, H10), not entirely trustworthy (N7, H11), and over-reliant on the scope of the collected papers, placing the effort on the users to find a representative initial set of papers (N3, N7, H10). We note that the main focus of IDEASYNTH is not on literature discovery or understanding. We incorporated basic paper search and summary functionalities as literature review plays an important role in research ideation. Future work can explore how to better integrate literature review with research ideation.

**5.3.4 Participants gained increased benefits from AI support in IDEASYNTH and spent less time writing from prior knowledge (RQ4).** To identify the differences in user behavior between the two conditions, we analyzed the interaction log from the user study sessions, where we recorded timestamped user or system actions. We recorded overt actions such as clicking on buttons to trigger system features (e.g. “Get AI Suggestions”, Q&A, search paper), users’ typing traces (e.g. writing research ideas in node content, or in the editor document). We further grouped actions as user-driven actions, such as writing and reviewing literature, or AI-assisted actions, such as generating a response to the user’s prompt,

which is present for both conditions, or generating new nodes in IDEASYNTH. We find that the participants utilized the AI assistance more frequently in IDEASYNTH than in baseline. On average, they spent 36.8% of their time during each task on AI-assisted actions in the treatment condition and 12.9% in the baseline condition. This is expected as IDEASYNTH integrates more LLM-based features to provide suggestions and scaffold idea expansion and exploration throughout the ideation process.

In contrast, participants spent a significantly larger portion of time writing in the baseline without utilizing features from the system. On average, they spent 19.2% of the time for each task on writing actions in IDEASYNTH (writing node title, node content, editing generated research brief) and 64.8% of time writing (writing in the document interface, editing AI-revised research brief generated using “AI Assist Writing”) in the baseline condition. Participants’ uninterrupted writing sessions were also longer in the baseline. Analyzing at the action level, an uninterrupted writing action (i.e., no switching to other actions) in baseline lasted on average 124.3 seconds, compared to 51.5 seconds in IDEASYNTH. However, participants engaged in similar total numbers of discrete writing actions (IDEASYNTH = 121, Baseline = 125), making smaller writing edits for each decomposed idea facet in IDEASYNTH, but wrote more extensive passages in the baseline. While spending more time on writing is not necessarily a hindrance to ideation, the document interface in the baseline did not provide users with adequate support for simultaneous divergent exploration of multiple ideas, according to some participants ( $n = 9$ ). H10 remarked that they *“mostly just used the [baseline] to basically prosify the idea that [they] already had, or maybe flesh it out in a little bit more detail, but not necessarily as much to consider alternatives.”* Similarly, H3 recalled that while using the baseline system, they were *“kind of fixated on a design and it’s too much work [to enter] in the whole thing into the assistant [Q&A interface] and ask it for different alternatives. So I just click on the ‘AI assist writing’ to make it to polish the research brief instead of trying to do more ideation.”*

### 5.3.5 Participants felt IDEASYNTH improved task success, but further study is required to evaluate research idea quality.

Participants rated their perceived level of success in accomplishing the task slightly more highly in IDEASYNTH ( $mean = 5.15$ ,  $SD = 0.81$ ) than the baseline ( $mean = 4.55$ ,  $SD = 1.23$ ,  $p < 0.05$  using independent Wilcoxon signed-rank test, Fig.8). In the combined dependent variable, which included Impact and Feasibility as key measures for research brief quality, we identified a significant difference (Fig.7) for IDEASYNTH over baseline. However, measuring participants’ ratings on the five dimensions—Novelty, Impact, Specificity, Feasibility, and Relevance (full definitions in Appendix 5)—independently, we did not identify a significant difference between conditions for any metrics. This might be due to the nature of uncertainty in the research ideation. The comparison of ideas might also be affected by the anchoring effect of the initial insight participants chose to develop. N5 expressed difficulty in evaluating their ideas in both conditions: *“I think the task [of rating idea quality] is honestly a little hard...the questions I answered about feasibility or specificity, it wasn’t about the system that I used, it’s just that particular idea that I came up with.”*

We wanted to capture participants' initial perceptions of the research ideas they developed in a controlled lab setting. However, evaluating research ideas and predicting research outcomes are difficult [49], even when reviewing complete paper submissions [4, 47]. Qualified experts are difficult to recruit and can still face difficulties in judging the quality of a research idea as evaluation criteria can be perceived highly subjectively [7, 49, 50]. To further investigate the naturalistic usage and ecological validity of IDEASYNTH, we conducted a deployment study described in the following section. We further discuss the need for an end-to-end study to predict the ideation outcome in the Discussion.

## 6 Study 2: Field Deployment Study

To further investigate how scholars would use IDEASYNTH to develop real-world research ideas organically, we deployed IDEASYNTH and invited participants from the lab study to continue using the tool on their own computers. Participants were asked to use the tool for at least three days, for at least one hour in total usage, and return for a 15-30 minutes follow-up interview. All seven participants (D1-D7; 3 in the field of HCI and 4 in NLP) completed the deployment study and were compensated \$25 USD for their time. The first author conducted thematic analysis on the interview transcripts[8, 13] (Appendix H). Below we report our findings.

### 6.1 Deployment Study Results

**6.1.1 Participants used IDEASYNTH for research ideas in different phases of ideation.** While we focused the design of IDEASYNTH on supporting idea expansion and refinement, in a field setting, participants tested the system across a wider range of the ideation phase. Predominantly, participants used IDEASYNTH for two primary tasks: research ideation and literature understanding. For research ideation, participants were able to apply IDEASYNTH to their research projects in various stages of development. D1, D3, and D7 explored early insights in IDEASYNTH and expanded them into more concrete structured research ideas. D3 described their process as “[like] I was having a discussion with my supervisor, I think we have a rough idea right now... I was trying to do some research ideation because I don’t really know anything about [topic]...it [IDEASYNTH] gives me a general thinking flow, there’s problem, research question components all the other stuff. So it gives me a complete research ideation process and the proposal.” Similarly, D7 found the tool “convenient because I’m sort of in the midst of a brainstorming phase for a different project. So mostly I just used it to sort of help me come up with ideas to think about different ways of framing, research questions, different methods to use.” Some participants ( $n = 3$ ) externalized their existing ideas using the node-based canvas and examining the connections between idea facets to reflect on their existing ideas more clearly. D4 applied IDEASYNTH on an idea where they already drafted the paper on and “already have a structure of how the research should go. And these texts [on the idea facet nodes] are filled by myself. And this time I focused on using the relationship [node edges] and I feel pretty surprised. I think this one really works pretty good. It gives me some suggestions...and it tells me the strengths of the connection.” Despite further developing a more mature research idea, D4 found benefits using IDEASYNTH: “my idea wasn’t this complex at the very beginning. It really helped

me to make my ideas completed with many branches and I was able to derive something pretty useful.” D2 also utilized the tool for an ongoing project: “I was trying to use the tool to help me with related work and how I might design evaluations.” On the other hand, most participants ( $n = 5$ ) found utility in contextualizing their research ideas using IDEASYNTH’s literature search and summary functionalities. For instance, D6 commented “because I have a paper deadline, I was trying to use the tool to more strengthen, like contextualize my paper in the HCI literature.”

**6.1.2 IDEASYNTH fits into researchers’ ideation workflow and can be used long-term.** We also collected participants’ response on how IDEASYNTH can fit in their research workflow and what some envisioned use cases are. Most participants ( $n = 5$ ) expressed interests to continue using IDEASYNTH for research ideation. D4 felt that the workflow of the tool was “fun to use” and that “one of the biggest contributions of this tool is it brought the barrier in my mind to start a research or start thinking or initiating this process way easier...and this fun or kind of excitement of using this tool to exploring project ideas makes me more like to do it.” D3 and D7 commented that even without the AI features, they would consider using IDEASYNTH’s representation of idea facets to organize and structure their ideas to stimulate thinking. Over the course of the entire research process, D5 suggested they could use the tool to perform “periodic checkpoints in the projects’ lifespan where you force yourself to think at a more high level... and how this fits into the larger narrative.” D1, D2, and D6 also see potential in using IDEASYNTH as a form of collective ideation tool, where researchers can share and iterate on the ideation canvas with collaborators.

**6.1.3 Participants viewed IDEASYNTH as an assistant to support and advocated for human involvement during ideation.** Regarding the role of LLMs in research ideation, most participants ( $n = 4$ ) expressed some level of concerns on the trustworthiness of LLM-generated literature summaries and research ideas. D1 experienced a conflict for adopting the literature findings presented by IDEASYNTH directly and continue with ideation or verifying them by reading the paper, sharing concerns on potential “misinterpretation” of collected papers. D4 felt that in general, “we cannot 100% believe that something that has been generated and click next”, citing worries about hallucination. D2 echoed the concern on generated ideas. Despite using ChatGPT in their research, they “don’t ask LLM to generate ideas for me...I do kind of worry about the novelty variation from the LLM. I didn’t trust that it can actually evaluate the novelty.”

For all participants ( $n = 7$ ), an ideal workflow collaborating with IDEASYNTH or AI systems in general consists of using LLM tools to assist and support the researchers in their ideation process. D6 expressed that while LLM can be useful for exploring ideas, to make decisions and evaluate the ideas, “human need to be more involved rather than just 99% using the tool.” Likewise, D7 reflected that “it’s [IDEASYNTH] most useful just as a way to sort of stir the mental pot. Even if I don’t find myself directly wanting to copy some of the suggestions...I like having that explicit graph structure and sort of explicitly generating alternatives...it sorts of get things bubbling a bit more mentally and forces you to reconsider things.” We further describe the study implications for future human-AI research support tools and the risks of using LLM for research ideation in the Discussion.

## 7 Discussion and Future Work

### 7.1 Tailoring AI Support to User Needs across Different Stages of Research Ideation

Our findings highlight that users desire varying levels of AI support depending on the stage of their research process. Specifically, our participants employed IDEASYNTH in different ways during the prolonged deployment study, emphasizing the need for flexible control, customizability, and adaptability of AI ideation systems.

In the early stages of research ideation, participants appreciated IDEASYNTH's ability to suggest alternative ideas and connections, which helped broaden their thinking. However, the desire for even more control over AI-generated suggestions was frequently mentioned. Users expressed interest in understanding the rationale behind the system's recommendations, particularly when certain papers or ideas were excluded. For instance, H2 questioned why certain relevant papers were not suggested and wanted more insight into the decision-making process behind AI outputs: *"It [system] puts it [generation] through a funnel... but I'd like to potentially just explore all possibilities."* This feedback suggests a need for greater transparency and explainability in AI-driven suggestions to ensure users feel empowered rather than constrained by the system.

As an idea becomes more developed, some participants switched to a more hands-on strategy, taking over more agency by selectively engaging with IDEASYNTH's suggestions to take greater control over their canvas workspace. However, H5 and H6 also reported that while user-driven interactions afforded greater control, they also require higher mental workload, which may hinder creative flow (H5), raising potential needs for additional passive AI support to streamline ideation (H6). For example, H6 remarked, *"I generally want something that passively helps me rather than me having to call into it to help me."* This highlights the tension between user-initiated and passive AI assistance, where users seek a balance between ease of use and meaningful involvement in the ideation process.

During the deployment study, participants further specialized in using IDEASYNTH's features based on their workflow needs. Many found it useful both for literature discovery and ideation but often separated these tasks. D7, for example, preferred to use the node-based system for structuring their research while leveraging the literature review features for understanding and differentiating related work. Similarly, D4 and D5 praised IDEASYNTH's ability to retrieve high-quality papers and quickly summarize them, using it primarily as a tool to strengthen their understanding of the research landscape. On the ideation side, participants like D1 and D3 used IDEASYNTH to explore and refine early-stage ideas, adding and removing nodes as their mental models evolved. Others, like D4, applied the tool to mature projects, finding it helpful for expanding and clarifying connections between research facets. The flexibility to adapt IDEASYNTH based on research idea maturity underscored the importance of customization, as participants tailored the system to their specific stages and tasks, whether brainstorming ideas or fine-tuning concepts. Overall, these findings highlight the need for flexible, user-controlled AI support to align with the dynamic and multifaceted nature of the research process.

### 7.2 Supporting User-Driven Research Processes Instead of Full Automation

There has been increasing interest in employing LLMs to fully automate some or all parts of the research process. Si et al. [49], for example, found evidence that LLM-generated initial research ideas are rated as more novel than ideas generated by human experts.<sup>8</sup> Lu et al. [38] introduced a system that automates the entire AI research workflow, from ideation, to experiment execution, to paper writing. While these works show that *AI alone* can match or even surpass the research task performance of *humans alone*, they do not reflect the more common real-world scenario where humans and AI tackle research problems together in *human-AI teams*.

Our work provides an example of how carefully designed user interfaces and interactions for LLM-powered research tools may be key to enabling researchers to collaborate effectively with AI models. For example, by breaking down ideation into idea facets, researchers could have greater control over how and where AI assistance is employed to further align with their goals during ideation. Participants in our study could easily expand and iterate on variations of a specific idea node within the context of a broader topic (Section 5.3.2). They were able to use the LLM's breadth of knowledge to surface ideas and details that they had not previously considered (Section 5.3.1), as well as uncover new connections between papers (Section 5.3.3). On top of that, our faceted design also allowed participants to better steer the AI when it completed a task in a suboptimal manner (e.g., participants mentioning that LLM outputs were too generic in Section 3.4), which is much more challenging to do in fully automated AI research pipelines.

As we observed in Section 3.1, researchers may not always want or need AI assistance throughout the entire life cycle of a research task, but may appreciate it at specific moments. We demonstrate by strategically guiding AI to assist with tasks where it excels, we can unlock new potential for researchers and AI to co-create higher-quality research outcomes than either humans or AI alone.

### 7.3 Implications of Using LLMs for Research Ideation

**7.3.1 Risks of Using LLMs for Research Ideation.** While utilizing LLMs enables efficient summary and synthesis of information based on the collected papers and user-written idea facets, the model generation could contain inherent biases that negatively impact the outcome of the research ideation. Studies have shown that LLMs could exhibit social biases (e.g. gender, race, religion) stemmed from its training data [32, 40]. Specifically, Dai et al. identified bias and unfairness from the recommendation and results in information retrieval systems integrated with LLMs [14]. In the scientific research context, this could lead to misrepresentation of certain publications, findings, or theories. Moreover, LLMs could potentially skew the research direction by presenting dominant or popular ideas and overlooking more novel or less-explored approaches, which could harm researchers' creativity. Another potential risk is hallucination, where an LLM generates content that misaligns with established world knowledge or previous interaction context [60].

<sup>8</sup>It is worth noting that most expert ideas were broad initial ideas that expert participants came up with on the spot specifically for the study.

This could lead to false ideas or fabricated sources being used during the ideation process.

To mitigate the above risks, IDEASYNTH grounds its generation on the curated paper collection. The generated idea facets and suggestions are often distilled from the original context of the collected papers in an effort to prevent hallucination. In an attempt to prevent misrepresentation of ideas, IDEASYNTH puts the user at the steering wheel when exploring and developing ideas. IDEASYNTH provides on-demand LLM-based support aimed to inspire the user, rather than guiding the user towards certain directions. Future systems can further mitigate these risks by improving model explainability by including the confidence level of generated statement and citing original text from the source.

**7.3.2 Ethical Implications and Intellectual Ownership.** An ethical concern that arises with human-AI collaborative ideation is the controversy around the intellectual ownership of the work and ideas produced [51]. While LLMs can assist researchers in synthesizing information and generating connections between ideas, their role in the creative process raises questions: if an LLM significantly influences the framing or direction of a research endeavor, to what extent can the resulting ideas be considered original contributions of the researcher? This also poses broader ethical questions about attributing co-authorship to machines: should LLMs be acknowledged as contributors in academic works? While current norms do not recognize machines as authors, there is a need for clearer guidelines on how to appropriately credit the influence of AI systems in research ideation without undermining the originality of human contributions. In addition, the LLM outputs, fundamentally derived from pre-existing data, might not accurately attribute the presented ideas to the original source document, which could lead to concerns regarding academic integrity and plagiarism.

There is also the risk of hindering human creativity due to over-reliance on LLMs [9, 31, 55]. If researchers begin to lean heavily on machine-generated suggestions, they may become less engaged in the iterative, reflective processes essential for developing innovative and impactful research. Such reliance could result in a superficial exploration of research ideas, with researchers focusing only on what is directly surfaced by the LLM while potentially overlooking subtler, more nuanced insights that require human intuition. From our user study, participants remained cautious employing LLM-generated ideas and generally examined the idea's validity before incorporating them onto the canvas. However, non-CS researchers might be less familiar with the potential pitfalls of relying on LLMs. The system should be explicit about the limitations in its generation. Ensuring that researchers maintain agency and actively critically evaluate generated outputs is crucial to avoiding this pitfall.

To address these concerns, IDEASYNTH does not aim to automate the ideation process or replace the human researcher. Instead, our system prioritizes providing users with controls and presenting relevant information grounded in literature to assist in user-driven ideation. Future work could further explore strategies for integrating LLM outputs with researcher-led reflection, ensuring that the machine-generated suggestions are contextualized and critically evaluated. This would not only safeguard the depth and originality of the research but also preserve the intellectual agency of human

researchers, fostering a truly collaborative rather than substitutive dynamic between humans and AI systems.

## 7.4 Limitations and Future Work

**7.4.1 Generalizability and Domain Customization.** Our studies involved participants from specific research areas in HCI and NLP, limiting the generalizability of our system and findings. We also observed that participants from each domains might have different workflows. For example, some participants from more technical domains, such as algorithm research, preferred writing equations by hand, highlighting the need for future systems to support diverse research methods. While our participants with diverse research topics within HCI and NLP domains benefited from our general approach, conducting further studies to cater to different areas of researchers could examine the system's ecological validity across a broader range of disciplines.

One of our key design feature is supporting users in structuring their research ideas using idea nodes. Participants found values in using the structure as a scaffold to further develop their ideas by exploring alternatives for different facets. However, while the current predefined structure is generalizable across multiple research domains and scenarios, some participants desired even more structures in the form of finer-grained facets or by defining their own customized facets (H10, N3). This observation points to interesting future directions of exploring customizable idea facets. Broadly, we envisioned two ways of supporting this. Firstly, we can define a more comprehensive taxonomy of research idea facets, and allow users to dynamically create hierarchies for their ideas. For example, under the current *proposed design and solution*, user could additionally create children nodes for *system architecture*, *interaction design*, etc. Alternatively, we could also allow users to define their own facets as they ideate in IDEASYNTH. Helping user maintain a useful and manageable structures and provide facet-specific AI supports remains exciting open challenges for future research.

**7.4.2 Evaluation of Research Ideas.** Estimating the quality of a research idea is challenging due to the uncertain nature of research progression, and even fully executed projects can be hard to predict for continuous impact [18, 28, 34, 49, 53]. Personal research interests can also influence how ideas are judged [21]. Therefore, we scope this project on the early idea development process and take a user-centered perspective, focusing on how IDEASYNTH affected researchers' ideation process based on participants' self-perception. The short timeframe of the lab study limits our ability to assess the long-term impact of IDEASYNTH on research outcomes. While our longer field deployment uncovered additional strategies participants employed to enhance real-world idea development, assessing the long-term quality, impact, or success of ideas generated through system usage remains a challenge. Future work could explore recruiting expert researchers who are experienced in evaluating early research ideas (e.g. senior professors) to assess the early ideation artifacts and predict the idea's impacts. However, to fully examine the outcome of LLM-assisted research ideation, future works should conduct end-to-end longitudinal studies incorporating public deployments or diary studies, focusing on tracking the lifecycle of ideas from inception to implementation and their

eventual outcomes in real-world contexts. For example, Si et al. conducted a large scale study where NLP experts evaluated LLM- and human-generated ideas, and they proposed a follow-up end-to-end study where researchers are recruited to execute the ideas into full projects [49]. Such efforts would provide deeper insights into the ecological validity of the system and its sustained utility over time.

**7.4.3 Paper Discovery.** Since research ideas rarely emerge in isolation, participants started ideation with both their initial ideas and any relevant papers they found. While IDEASYNTH used basic tools like search and a recommendation system to expand the paper set, the process remained largely manual. Users were responsible for the paper curation as an attempt to mitigate LLM hallucination on source information and ground generations in real paper contents. However, some participants felt the scope of recommended papers was too limited to their existing collection, making it harder to find additional relevant work. This suggests that while focused idea development is important, enabling broader exploration in the knowledge space might also be valuable. Future work should balance broader paper discovery with focused, specific feedback to avoid idea drift while supporting further development.

## 8 Conclusion

While many existing research ideation tools focused on supporting broad initial ideation, this paper focused on supporting the iterative expansion and deeper refinement stages of research ideation. To bridge this gap, we introduced IDEASYNTH, a research idea development system that provides iterative support for idea refinement and evaluation. IDEASYNTH delivers literature-based feedback grounded by user-collected papers, helping researchers articulate idea facets, including research problems, solutions, evaluations, and contributions. Our system's node-based canvas enables users to visualize variation of idea facets, iteratively refine them, and eventually connect different nodes to compose a coherent research brief. Our lab study demonstrated that participants using IDEASYNTH explored more idea variations and expanded their initial ideas with greater detail compared to a strong baseline. The deployment study further reinforced the system's utility in more realistic scenarios, with participants reported benefits using IDEASYNTH across various stages of their real-world research projects. These findings underscore the importance of tailoring AI support for different ideation stages to flexibly iterate on and develop research ideas.

## References

- [1] 2008–2024. GROBID. <https://github.com/kermitt2/grobid.swh:1.dir:dab86b296e3c3216e2241968f0d63b68e8209d3c>
- [2] Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th Conference on Creativity & Cognition*. 413–425.
- [3] Chittaranjan Andrade. 2011. How to write a good abstract for a scientific paper or conference presentation. *Indian Journal of Psychiatry* 53 (2011), 172 – 175. <https://api.semanticscholar.org/CorpusID:28535477>
- [4] Ines Arous, Jie Yang, Mourad Khayati, and Philippe Cudré-Mauroux. 2021. Peer Grading the Peer Reviews: A Dual-Role Approach for Lightening the Scholarly Paper Review Process. *Proceedings of the Web Conference 2021* (2021). <https://api.semanticscholar.org/CorpusID:233334418>
- [5] Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2024. ResearchAgent: Iterative Research Idea Generation over Scientific Literature with Large Language Models. doi:10.48550/arXiv.2404.07738 arXiv:2404.07738 [cs].
- [6] Judith Bell and Stephen Waters. 1993. Doing Your Research Project: A Guide for First-time Researchers. <https://api.semanticscholar.org/CorpusID:60359883>
- [7] Alina Beygelzimer, Yann Dauphin, Percy Liang, and Jenn Wortman Vaughan. 2023. Has the Machine Learning Review Process Become More Arbitrary as the Field Has Grown? The NeurIPS 2021 Consistency Experiment. *ArXiv abs/2306.03262* (2023).
- [8] Richard E. Boyatzis. 1998. Transforming Qualitative Information: Thematic Analysis and Code Development.
- [9] Zana Bućinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 188 (April 2021), 21 pages. doi:10.1145/3449287
- [10] Clara E. Busse and Ella T. August. 2017. How to Write and Publish a Research Paper for a Peer-Reviewed Journal. *Journal of Cancer Education* 36 (2017), 909 – 913. <https://api.semanticscholar.org/CorpusID:13253899>
- [11] Eric Chamoun, Michael Schlichtrull, and Andreas Vlachos. 2024. Automated Focused Feedback Generation for Scientific Writing Assistance. *arXiv preprint arXiv:2405.20477* (2024).
- [12] Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel S. Weld. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. *ArXiv abs/2004.07180* (2020). <https://api.semanticscholar.org/CorpusID:215768677>
- [13] Lynne M. Connelly. 2013. Grounded theory. *Medsurg nursing : official journal of the Academy of Medical-Surgical Nurses* 22 2 (2013), 124, 127.
- [14] Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. Bias and Unfairness in Information Retrieval Systems: New Challenges in the LLM Era. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Barcelona, Spain) (KDD '24). Association for Computing Machinery, New York, NY, USA, 6437–6447. doi:10.1145/3637528.3671458
- [15] Mike D'Arcy, Tom Hope, Larry Birnbaum, and Doug Downey. 2024. MARG: Multi-Agent Review Generation for Scientific Papers. doi:10.48550/arXiv.2401.04259 arXiv:2401.04259 [cs].
- [16] Annie S. Ditta and Benjamin C. Storm. 2016. That's a good idea, but let's keep thinking! Can we prevent our initial ideas from being forgotten as a consequence of thinking of new ideas? *Psychological Research* 81 (2016), 678 – 689. <https://api.semanticscholar.org/CorpusID:40612122>
- [17] LD Douglas, MH Jillian, LR Thomas, and LS Eric. 2006. Identifying quality, novel, and creative ideas: constructs and scales for idea evaluation. *Journal of the Association for Information Systems* 7, 10 (2006), 646.
- [18] M. Ferioli, Elies Dekoninck, Stephen J. Culley, Benoit Roussel, and Jean Renaud. 2010. Understanding the rapid evaluation of innovative ideas in the early stages of design. *International Journal of Product Development* 12 (2010), 67–83. <https://api.semanticscholar.org/CorpusID:111197215>
- [19] Allen Foster. 2005. A non-linear model of information seeking behaviour. *Inf. Res.* 10 (2005). <https://api.semanticscholar.org/CorpusID:33942217>
- [20] Karan Girotra, Lennart Meincke, Christian Terwiesch, and Karl T Ulrich. 2023. Ideas are dime a dozen: Large language models for idea generation in innovation. *Available at SSRN 4526071* (2023).
- [21] Xuemei Gu and Mario Krenn. 2024. Generation and human-expert evaluation of interesting research ideas using knowledge graphs and large language models. doi:10.48550/arXiv.2405.17044 arXiv:2405.17044 [cs].
- [22] Sandra G Hart. 2006. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, Vol. 50. Sage publications Sage CA: Los Angeles, CA, 904–908.
- [23] Jessica He, Stephanie Houde, Gabriel E. Gonzalez, Dario Andrés Silva Moran, Steven I. Ross, Michael Muller, and Justin D. Weisz. 2024. AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (June 2024), 1–14. doi:10.1145/3663384.3663398 Conference Name: CHIWORK 2024: Annual Symposium on Human-Computer Interaction for Work ISBN: 9798400710179 Place: Newcastle upon Tyne United Kingdom Publisher: ACM.
- [24] Tom Hope, Doug Downey, Daniel S Weld, Oren Etzioni, and Eric Horvitz. 2023. A computational inflection for scientific discovery. *Commun. ACM* 66, 8 (2023), 62–73.
- [25] Nanna Inie, Jonas Frich, and Peter Dalsgaard. 2022. How Researchers Manage Ideas. In *Proceedings of the 14th Conference on Creativity and Cognition (C&C '22)*. Association for Computing Machinery, New York, NY, USA, 83–96. doi:10.1145/3527927.3532813
- [26] Richard H. Kallet. 2004. How to write the methods section of a research paper. *Respiratory care* 49 10 (2004), 1229–32. <https://api.semanticscholar.org/CorpusID:12280435>
- [27] Hyeonsu B. Kang, Sheshera Mysore, Kevin Huang, Haw-Shiuan Chang, Thorben Prein, Andrew McCallum, Aniket Kittur, and Elsa Olivetti. 2022. Augmenting Scientific Creativity with Retrieval across Knowledge Domains. <http://arxiv.org/abs/2206.01328> arXiv:2206.01328 [cs].
- [28] Hyeonsu B. Kang, Xin Qian, Tom Hope, Dafna Shahaf, Joel Chan, and Aniket Kittur. 2022. Augmenting Scientific Creativity with an Analogical Search Engine. *ACM Transactions on Computer-Human Interaction* 29, 6 (Nov. 2022), 57:1–57:36. doi:10.1145/3530013



- [29] Hyeonsu B. Kang, Sherry Tongshuang Wu, Joseph Chee Chang, and Aniket Kittur. 2023. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–19. doi:10.1145/3586183.3606759 arXiv:2308.07517 [cs].
- [30] Jeffrey W. Knopf. 2006. Doing a Literature Review. *PS: Political Science & #38; Politics* 39, 1 (2006), 127–132. doi:10.1017/S1049096506060264
- [31] Harsh Kumar, Jonathan Vincentius, Ewan Jordan, and Ashton Anderson. 2024. Human Creativity in the Age of LLMs: Randomized Experiments on Divergent and Convergent Thinking. arXiv:2410.03703 [cs.HC] <https://arxiv.org/abs/2410.03703>
- [32] Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards Understanding and Mitigating Social Biases in Language Models. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 6565–6576. <https://proceedings.mlr.press/v139/liang21a.html>
- [33] Weixin Liang, Yuhui Zhang, Hancheng Cao, Binglu Wang, Daisy Ding, Xinyu Yang, Kailas Vodrahalli, Siyu He, Daniel Smith, Yian Yin, et al. 2023. Can large language models provide useful feedback on research papers? A large-scale empirical analysis. *arXiv preprint arXiv:2310.01783* (2023).
- [34] Brian Licuanan, Lesley R. Dailey, and Michael D. Mumford. 2007. Idea evaluation: Error in evaluating highly original ideas. *Journal of Creative Behavior* 41 (2007), 1–27. <https://api.semanticscholar.org/CorpusID:144243454>
- [35] Weng Marc Lim, Satish Kumar, and Faizan Ali. 2022. Advancing knowledge through literature reviews: ‘what’, ‘why’, and ‘how to contribute’. *The Service Industries Journal* 42, 7–8 (2022), 481–513.
- [36] Yiren Liu, Si Chen, Haocong Cheng, Mengxia Yu, Xiao Ran, Andrew Mo, Yiliu Tang, and Yun Huang. 2024. How AI Processing Delays Foster Creativity: Exploring Research Question Co-Creation with an LLM-based Agent. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 17, 25 pages. doi:10.1145/3613904.3642698
- [37] Yiren Liu, Mengxia Yu, Meng Jiang, and Yun Huang. 2023. Creative Research Question Generation for Human-Computer Interaction Research.. In *IUI Workshops*. 58–66.
- [38] Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. The AI scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292* (2024).
- [39] Chris A. Mack. 2012. How to write a good scientific paper: title, abstract, and keywords. *Journal of Micro-nanolithography Mems and Moems* 11 (2012), 020101–020101. <https://api.semanticscholar.org/CorpusID:12481302>
- [40] Roberto Navigli, Simone Conia, and Björn Ross. 2023. Biases in Large Language Models: Origins, Inventory, and Discussion. *J. Data and Information Quality* 15, 2, Article 10 (June 2023), 21 pages. doi:10.1145/3597307
- [41] Harshit Nigam, Manasi S. Patwardhan, Lovekesh Vig, and Gautam M. Shroff. 2024. Acceleron: A Tool to Accelerate Research Ideation. *ArXiv abs/2403.04382* (2024). <https://api.semanticscholar.org/CorpusID:268264612>
- [42] Paul B Paulus and Huei-Chuan Yang. 2000. Idea generation in groups: A basis for creativity in organizations. *Organizational behavior and human decision processes* 82, 1 (2000), 76–87.
- [43] Biqing Qi, Kaiyan Zhang, Haoxiang Li, Kai Tian, Sihang Zeng, Zhang-Ren Chen, and Bowen Zhou. 2023. Large language models are zero shot hypothesis proposers. *arXiv preprint arXiv:2311.05965* (2023).
- [44] Marissa Radensky. 2024. Mixed-Initiative Methods for Co-Creation in Scientific Research. *Creativity and Cognition* (June 2024), 1–7. doi:10.1145/3635636.3664627 C&C '24: Creativity and Cognition.
- [45] Marissa Radensky, Simra Shahid, Raymond Fok, Pao Siangliulue, Tom Hope, and Daniel S Weld. 2024. Scideator: Human-LLM Scientific Idea Generation Grounded in Research-Paper Facet Recombination. *arXiv preprint arXiv:2409.14634* (2024).
- [46] Orit Shaer, Angelora Cooper, Osnat Mokryn, Andrew L Kun, and Hagit Ben Shoshan. 2024. AI-Augmented Brainwriting: Investigating the use of LLMs in group ideation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 1050, 17 pages. doi:10.1145/3613904.3642414
- [47] Nihar B. Shah. 2022. An Overview of Challenges, Experiments, and Computational Solutions in Peer Review (Extended Version). <https://api.semanticscholar.org/CorpusID:255751797>
- [48] Hua Shen, Chieh-Yang Huang, Tongshuang Wu, and Ting-Hao Kenneth Huang. 2023. ConvXAI : Delivering Heterogeneous AI Explanations via Conversations to Support Human-AI Scientific Writing. In *Computer Supported Cooperative Work and Social Computing*. ACM, Minneapolis MN USA, 384–387. doi:10.1145/3584931.3607492
- [49] Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2024. Can LLMs Generate Novel Research Ideas? A Large-Scale Human Study with 100+ NLP Researchers. *arXiv preprint arXiv:2409.04109* (2024). <https://api.semanticscholar.org/CorpusID:272463952>
- [50] Müge Simsek, Mathijs de Vaan, and Arnout van de Rijt. 2024. Do grant proposal texts matter for funding decisions? A field experiment. *Scientometrics* 129 (2024), 2521–2532. <https://api.semanticscholar.org/CorpusID:269924192>
- [51] Jan Smits and Tijn Borghuis. 2022. *Generative AI and Intellectual Property Rights*. T.M.C. Asser Press, The Hague, 323–344. doi:10.1007/978-94-6265-523-2\_17
- [52] K M Sudheesh, Devika Rani Duggappa, and S. S. Nethra. 2016. How to write a research proposal? *Indian Journal of Anaesthesia* 60 (2016), 631 – 634. <https://api.semanticscholar.org/CorpusID:10708507>
- [53] Shiro Takagi, Ryutaro Yamauchi, and Wataru Kumagai. 2023. Towards Autonomous Hypothesis Verification via Language Models with Minimal Guidance. *arXiv preprint arXiv:2311.09706* (2023).
- [54] Jan P Vandenbroucke and Neil Pearce. 2018. From ideas to studies: how to get ideas and sharpen them into research questions. *Clinical epidemiology* (2018), 253–264.
- [55] Helena Vasconcelos, Matthew Jörke, Madeleine Grunde-McLaughlin, Tobias Gerstenberg, Michael Bernstein, and Ranjay Krishna. 2022. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 7 (2022), 1 – 38. <https://api.semanticscholar.org/CorpusID:254591809>
- [56] Samangi Wadinambiarachchi, Ryan M Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18.
- [57] Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. 2024. SciMON: Scientific Inspiration Machines Optimized for Novelty. doi:10.48550/arXiv.2305.14259 arXiv:2305.14259 [cs].
- [58] Qingyun Wang, Lifu Huang, Zhiying Jiang, Kevin Knight, Heng Ji, Mohit Bansal, and Yi Luan. 2019. PaperRobot: Incremental Draft Generation of Scientific Ideas. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 1980–1991. doi:10.18653/v1/P19-1191 arXiv:1905.07870 [cs].
- [59] Zonglin Yang, Xinya Du, Junxian Li, Jie Zheng, Soujanya Poria, and Erik Cambria. 2023. Large language models for automated open-domain scientific hypotheses discovery. *arXiv preprint arXiv:2309.02726* (2023).
- [60] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. Siren’s Song in the AI Ocean: A Survey on Hallucination in Large Language Models. arXiv:2309.01219 [cs.CL] <https://arxiv.org/abs/2309.01219>



A Demographic Details of Formative Study Participants

P#	Gender	Age Range	Country	Research YoE	Research Area	AI Use Frequency
P1	Man	25–34	U.S.	6–10	Human-AI interaction	A couple times
P2	Man	25–34	U.S.	6–10	LLM evaluation	Occasionally
P3	Woman	18–24	U.S.	2–5	LLMs + society	Never
P4	Man	25–34	Canada	2–5	Multilingual NLP	Occasionally
P5	Woman	25–34	U.S.	6–10	Human-AI interaction	Occasionally
P6	Man	25–34	U.S.	2–5	Multimodal AI & HCI	Occasionally
P7	Woman	25–34	South Korea	2–5	LLM retrieval	Occasionally
P8	Man	25–34	U.S.	6–10	NLP & HCI	Occasionally

**Table 1: Participants from our formative study. Country refers to the country in which the participant primarily conducts research at the time of the study. Research YoE refers to the years of experience conducting academic research. AI Use Frequency refers to how frequently the participants uses AI tools to ideate and/or iterate on research ideas. Participants selected “Occasionally” based on the description “I don’t rely on AI but sometimes tinker with it.”**

B Formative Study Wizard-of-Oz Prototype LLM Prompt

You will be given a brief description of a researcher's research interests and attachments with their most recent publications.  
Generate 3 novel research ideas based on the information. The format of each idea should be as follows:  
  
{idea title}  
{a super concise and specific bulleted list describing the idea. Include no more than 3 bullet points.}  
  
Here is the brief description of research interests:  
{user\_research\_interest\_from\_survey}

C Formative Study Qualitative Analysis Themes

Theme	Description
Ideation process	How participants formulated initial ideas, reviewed and synthesized literature, backtracked and changed direction of their ideas, evaluated and converged on an idea, and developed their own approaches to ideation.
Workflow challenges	Challenges encountered by participants during ideation, related to areas such as managing ideas and information, evaluating ideas, scoping ideas, effectively communicating/gathering feedback on ideas, and reviewing literature.
Artifacts generated	Characteristics of artifacts participants generated throughout their ideation and research process, both in their own work outside of the study as well as during the study.
Use of AI	How participants used AI and envisioned to use AI, both in their own work outside of the study as well as during the study. Limitations and avoidance of AI were also included.
Feedback	Participants’ feedback on the interactive capabilities of the probe and responses provided by the AI system, as well as new features and interactions they wanted to see going forward.

**Table 2: Themes derived from a qualitative analysis of our formative study transcripts and recordings.**

## D Demographic Details of Comparative Lab Study Participants

P#	Gender	Age Range	Country	Research YoE	Research Area	AI Use Frequency
H1	Woman	25–34	Canada	6–10	HCI + AI	Occasionally
H2	Man	25–34	Canada	6–10	Mixed Reality	Never
H3	Woman	25–34	U.S.	2–5	HCI + AI	Occasionally
H4	Man	18–24	U.S.	2–5	Fabrication	A couple times
H5	Man	18–24	U.S.	6–10	Ubiquitous Computing	Always
H6	Man	25–34	Canada	2–5	Interaction Techniques	A couple times
H7	Woman	25–34	U.S.	6–10	HCI + Health	A couple times
H8	Man	25–34	U.S.	2–5	HCI + Generative AI	Always
H9	Man	25–34	U.S.	2–5	HCI + AI	Always
H10	Man	25–34	U.S.	6–10	Computational Social Science	A couple times
H11	Non-binary	25–34	U.S.	2–5	HCI + Culture	Never
H12	Woman	25–34	Canada	6–10	Haptics + Robotics	Frequently
N1	Woman	18–24	U.S.	2–5	Language Modeling	Frequently
N2	Man	25–34	Canada	6–10	LLM Interpretability	Always
N3	Man	25–34	U.S.	2–5	Conversational Agents	Always
N4	Man	18–24	U.S.	2–5	NLP Fairness + Law	A couple times
N5	Woman	25–34	U.S.	2–5	LLM Alignment + Security	Never
N6	Man	25–34	U.S.	2–5	Multimodal LLMs	A couple times
N7	Woman	25–34	U.S.	2–5	LLM Interpretability	Never
N8	Woman	18–24	U.S.	0–1	Educational Data Science	Frequently

**Table 3: Participants from our laboratory study. P# is the participant ID. Country refers to the country in which the participant primarily conducts research at the time of the study. Research YoE refers to the years of experience conducting academic research. AI Use Frequency refers to how frequently the participants uses AI tools to ideate and/or iterate on research ideas. Participants selected “Occasionally” based on the description “I don’t rely on AI but sometimes tinker with it.”**

## E Comparative Lab Study User Task Instruction

<p>User Task Sheet</p> <p>Welcome to the Research Ideation Session!</p> <p>You will be working on research tasks within the domain of [Task Theme]. Below, you'll find the details of your assigned tasks, along with important criteria to keep in mind as you develop your ideas.</p> <p>Task Overview</p> <p>You will be assigned one of the following broad topics:</p> <p>Task 1: [Task Topic 1]</p> <p>Task 2: [Task Topic 2]</p> <p>Your Goal:</p> <p>You will have 30 minutes to complete each task. The focus of the study is on the research ideation process. For each task, explore research ideas and literature to draft a research brief of around 12 sentences.</p> <p>Your brief should include the following sections:</p> <p>Problem Description and RQ</p> <p>Proposed Design and Solution</p> <p>Evaluation Method</p> <p>Contribution and Impact</p> <p>Evaluation Criteria:</p> <p>As you work on your tasks, consider the following criteria for your research brief:</p> <p>Novelty</p> <p>Definition: The research brief justifies the research idea's originality by clearly differentiating it from prior works.</p> <p>Impact</p> <p>Definition: The research brief describes ideas that make a nontrivial contribution to advance knowledge in the field.</p> <p>Specificity</p> <p>Definition: The research brief is clearly elaborated and sufficiently detailed with a complete description of the research idea(s).</p> <p>Feasibility</p> <p>Definition: The research brief describes the implementation of the ideas in a reasonable way that does not violate known constraints.</p>
---

**Relevance**

**Definition:** The research brief addresses the task topic at hand and can be reasonably expected to apply to the research problem.

Please ensure that your research brief reflects these attributes.

Please click on the System URL and the User Survey to start.

Good luck, and enjoy the ideation process!

## F Comparative Lab Study Qualitative Analysis Themes

Theme	Description
Comparison of system and workflows	How participants compared our IDEASYNTH system with the baseline system, as well as comparisons of their workflow in IDEASYNTH with their existing or past workflows without IDEASYNTH.
Use cases	Comments on the specific use cases of IDEASYNTH. Examples include idea generation, literature review, sharing research ideas, creating mind maps, onboarding new researchers, and more.
General system benefits	Participants' perceived benefits of IDEASYNTH. Examples include improving literature understanding, more thoroughly expanding idea facets, refining ideas through iteration, exploring ideas through alternatives, preventing idea fixation, and more.
Usage of node-based canvas interface	How participants made use of IDEASYNTH's node-based canvas interface. Includes comments about the use of individual nodes, node edges, selecting a group of nodes and generating a research brief, and AI suggestions stemming from nodes.
System limitations	Participants' comments about IDEASYNTH's limitations and drawbacks, most notably drawbacks of the node-based canvas interface, broadness of LLM suggestions, misalignment between LLM output and participant's idea, and lack of trust towards AI suggestions/summaries.
Future directions	Participants' suggestions for directions to improve IDEASYNTH as well as addressing tradeoffs and ethical concerns of using the system.

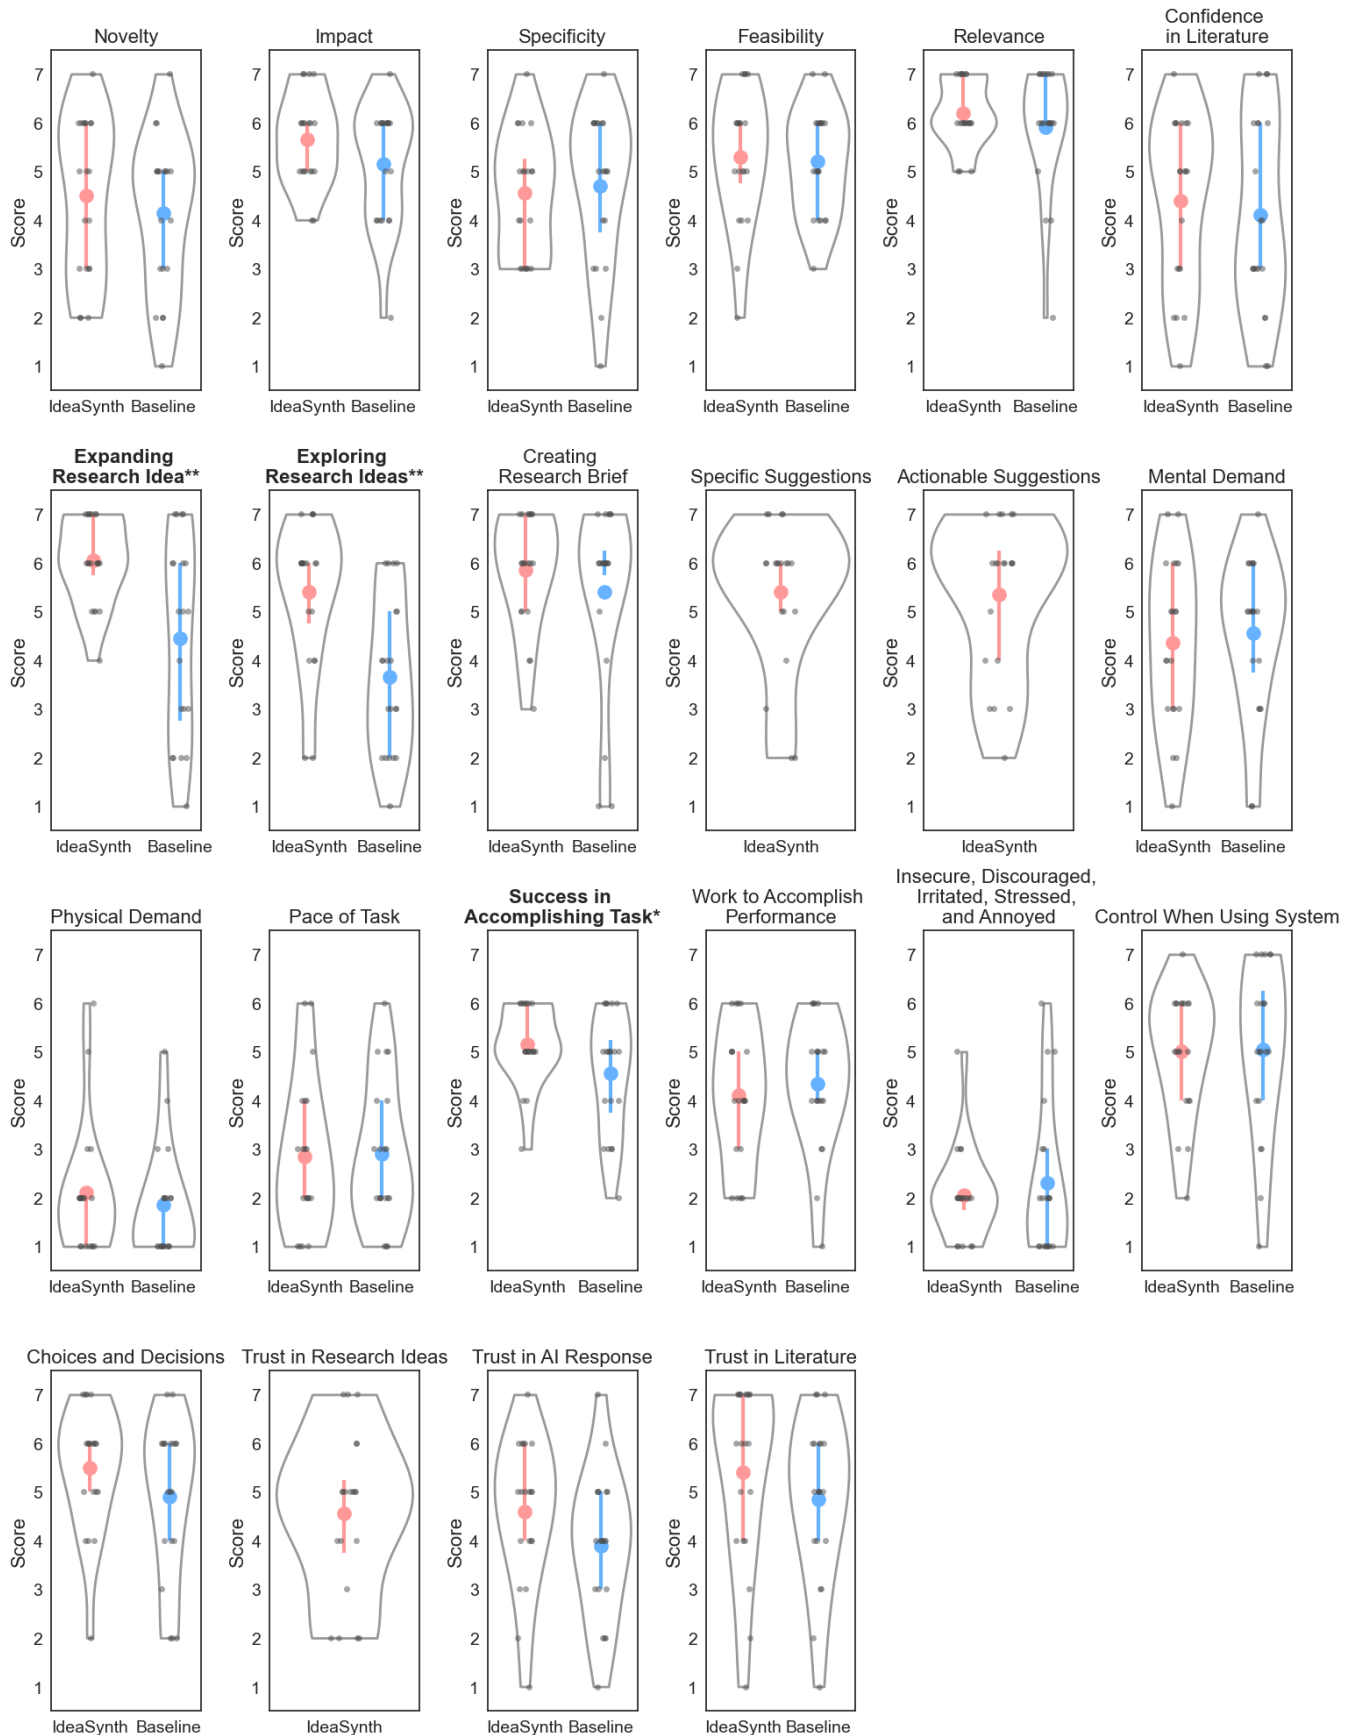
**Table 4: Themes derived from a qualitative analysis of our lab study transcripts and recordings.**

## G Comparative Lab Study Survey Response

On a scale of 1-Strongly disagree, 2-Disagree, 3-Slightly disagree, 4-Neutral, 5-Slightly agree, 6-Agree, 7-Strongly agree, how much do you agree with the following statements:	
The research brief justifies the research idea's originality by clearly differentiating from prior works.	Novelty
The research brief describes ideas that make a nontrivial contribution to advance the knowledge in the field.	Impact
The research brief is clearly elaborated and sufficiently detailed with the complete description of the research idea(s).	Specificity
The research brief describes the implementation of the ideas in a reasonable way that does not violate known constraints.	Feasibility
The research brief addresses the task topic at hand and can be reasonably expected to apply to the research problem.	Relevance
I felt confident in my understanding of the existing literature related to my research idea.	Confidence in Literature
I was able to effectively expand on my initial research idea with more details.	Expanding Research Idea
I was able to sufficiently explore and evaluate alternative versions of my research ideas.	Exploring Research Ideas
I was able to create a well-structured and coherent research brief.	Creating Research Brief
The system provides suggestions that are specific to my research ideas and the related literature.	Specific Suggestions
The system provides suggestions that are actionable for me to advance my research ideation process.	Actionable Suggestions
How mentally demanding was the task? 1-Very low, 7 Very high	Mental Demand
How physically demanding was the task? 1-Very low, 7 Very high	Physical Demand
How hurried or rushed was the pace of the task? 1-Very low, 7 Very high	Pace of Task
How successful were you in accomplishing what you were asked to do? 1-Very low, 7 Very high	Success in Accomplishing Task

How hard did you have to work to accomplish your level of performance? 1-Very low, 7 Very high	Work to Accomplish Performance
How insecure, discouraged, irritated, stressed, and annoyed were you? 1-Very low, 7 Very high	Insecure, Discouraged, Irritated, Stressed, and Annoyed
I felt that I was in control when I used the system.	Control When Using System
I felt like the system allowed me to make choices and decisions while developing the research idea.	Choices and Decisions
I trust the system's generated research ideas to be used in a real research scenario.	Trust in Research Ideas
I trust the system's AI response to my prompt.	Trust in AI Response
I trust the system's reference and summary for the related literature.	Trust in Literature

**Table 5: Survey questions and their corresponding labels**



**Figure 8: Survey Response.** Bolded plots show statistically significant differences ( $p < 0.05^*$ ,  $p < 0.01^{**}$ ) between IDEA SYNTH and baseline, using Wilcoxon signed-rank test if each survey question is treated independently. However, to avoid multiple comparison bias, we only report pre-selected measures (Fig.7) before analyzing the survey results.

## H Deployment Study Qualitative Analysis Themes

Theme	Description
Task and Tool Usage	How participants utilized IDEASYNTH in their real-life projects, the stages of ideation they used the tool for, and the outcome of the tool usage.
Helpful Features	Significant system functionalities that assisted participants, and the specific benefits resulted from those features.
Future Use Case and Workflow	Envisioned use case of IDEASYNTH longitudinally, how the tool can be applied in real-life research setting, and how the tool can be fit into participants' workflows.
Concerns with LLM in Research	Trust towards LLM generated literature summary, concern about evaluating LLM-generated research ideas, and tension on user control.
System Suggestions	Expanding literature scope to suggest more papers, customizability to define idea facets based on the research domain, and adaptive assistance for different parts of the research process.

**Table 6: Themes derived from a qualitative analysis of our deployment study transcripts and recordings.**

## I IDEASYNTH LLM Prompts

### I.1 System Prompt

You are a helpful research assistant. Your task is to assist a researcher in developing their initial research idea into a comprehensive research brief during a research ideation session.  
The research brief should include the following facets:  
'Problem Description and RQ', 'Proposed Design and Solution', 'Evaluation Method', 'Contribution and Impact'.  
These facets are represented as nodes, and are interconnected, with each connection representing the logical relationship between them.  
The entire path of research facets and their connections will be used to draft the final research brief.

### I.2 Paper Processing Prompt

Answer a question based on the following excerpts from the full text of the paper.

{full\_text}

Given the information above, please summarize the following aspects of the paper each in a paragraph:  
Problem Description and RQ, Proposed Design and Solution, Evaluation Method, Contribution and Impact, Limitation and Future Work.

Respond in proper JSON format like this {"key": "value"}.

If the question cannot be answered from the provided context, reply {"No answer": ""}

### I.3 Node Suggestion Prompt

Idea Facet: {idea\_facet}

Research Idea:

Title: {node\_title}

Content: {node\_content}

Here are related literatures, ONLY use these provided papers:

{papers\_data}

Your main tasks are:

1. Regenerate Current Idea Facet: Ask clarifying questions to help users reflect on their description. Provide evidence to support or challenge the user's claims.
2. Generate Alternatives: If the user wishes to expand on the current facet, generate a few alternative versions based on their feedback or selected context.
3. Generate New Idea Facets: Help generate new facets that are specific and relate to the existing facets and idea context when the user wants to expand their idea.

Provide your response in the following format:

```
{
  "ai_suggestion": [
    {
      "idea_facet": [Idea Facet to be act on],
      "action": [Action that the user can take or the system should perform, from the list above, based on the input data],
    }
  ]
}
```

```

        "suggestion": [Proposed Suggestion, whether it is questions for the user, supporting information, or
        suggested edits, if action is targeted to a specific idea facet, the suggestion should name that facet. If
        the suggestion references a paper, add a citation tag in the format of @ref[corpusId].],
    },
    ...
]
}}

```

Here are two example outputs:

```

{{
  "ai_suggestion": [
    {{
      "idea_facet": "Problem Description and Research Question",
      "action": "Regenerate Current Idea Facet",
      "suggestion": "Could you elaborate on the specific features of your system that are hypothesized to enhance
      [the task context]? How do these features align with the findings from studies like 'Paper2'@ref[corpusId]
      and 'Paper1'@ref[corpusId]?"
    }},
    {{
      "idea_facet": "Proposed Design/Solution",
      "action": "Generate New Idea Facets",
      "suggestion": "Maybe we can start to think about the next facet, Evaluation Method. How can we design the
      metrics and experimental setup to measure the effectiveness of the proposed system's features?"
    }}
  ]
}}

```

You may include multiple suggestions, like in the example, if necessary, but each suggestion should only be 1-2 sentences and easy to read. Always respond in a clear and concise manner, focusing on the specific idea facet or task at hand. Do not summarize or reiterate what the user has said. Begin your response by identifying which idea facet the user is focusing on, and then proceed with the appropriate task based on their input. Reference the related literature to provide evidence or context for your suggestions.

## I.4 Node Generation Prompt

You will be generating a new research idea facet based on the user's existing idea facet and context. Your goal is to create a specific and relevant facets that contributes to the overall research brief.

Here is the current idea facet:

```
{current_node_data}
```

And here is the context for the new idea facet:

```
{node_context}
```

Generate new idea facet(s) based on the following steps:

1. Review the user's current idea facet and the context of all connected nodes.
2. Identify key themes, objectives, or gaps that can be addressed by the new facet.
3. Generate a new idea facet that is specific, relevant, and logically connected to the current idea facet.
4. Ensure that the new facet contributes to the overall coherence and completeness of the research brief.

Here are the instructions:

```
{action_to_take}
```

Use the suggestion to guide the generation of the new idea facet(s).

If the instruction for Action is Regenerate Current Idea Facet or Regenerate Current Node, use the suggestion from existing facet and the user's feedback to generate ONLY 1 refined version of the existing facet.

If the instruction for Action is Generate Alternatives, generate up to 3 alternatives of the same type as the original facet to ensure that the new facets offers distinct perspectives or focuses.

If the instruction for Action is one of 'Problem Description and RQ' or 'Proposed Design and Solution' or 'Evaluation Method' or 'Contribution and Impact', generate up to 3 new nodes of ONLY that specific facet type and clearly establish the logical connection to the original facet.

Provide your response in the following format:

```

{{
  "new_nodes": [
    {{
      "ideaFacet": [New Idea Facet Type],
      "title": [Title of the new idea facet in less than 10 words],
      "content": [Content of the new idea facet in 1-3 short sentences, use the (author, year) format to refer to
      paper without the corpusId],
    }},
    ...
  ]
}}

```

ideaFacet should be one of the following: 'Problem Description and RQ', 'Proposed Design and Solution', 'Evaluation Method', 'Contribution and Impact'.

Ensure that each new idea facet is clearly defined and contributes to the overall research brief. Your response should be concise, focused, and directly relevant to the user's input.

## I.5 Node Edge Generation Prompt

You will be establishing connections between two different idea facets provided by the user. Your goal is to create logical relationships between the facets that form a coherent research brief.

Here are the two idea facets you need to connect:

```
{source_node_data}
```

```
{target_node_data}
```

To establish a connection between the two facets, follow these steps:

Analyze facets:

- Identify main ideas and key elements.
- Determine strengths and limitations.

Identify Connections:

- Find common themes or complementary aspects.
- Consider how one facet could enhance or support the other in the research process.

Suggest Enhancements:

- Provide specific ways to refine the connection, with examples or references.

Format your response as follows, replacing the placeholders with the relevant information:

```
{{
  "connectionStrength": [a number value from 0 to 1 indicating the coherence of the connection, 1 being the
strongest],
  "suggestion": [Start with a statement, such as "The connection is weak/average/strong". Then use one sentence to
offer suggestion for enhancing the connection. Make your point concise and clear],
}}
```

Here are two example responses:

```
{{
  "connectionStrength": 0.8,
  "suggestion": "The connection is relatively strong. Consider making the measureable outcomes, such as task
completion time, explicit in the evaluation methods to directly link back to metrics that can be extracted from the
proposed design."
}}
{{
  "connectionStrength": 0.3,
  "suggestion": "The connection is weak. The dataset in your proposed method might not offer the user interaction
granularity that is required to answer the research questions. Consider framing the questionis more broadly, like
in paper @ref[corpusId]."
}}
```

Your response should be concise, clear, and focused on the logical relationship between the two facets. Aim to provide actionable suggestions that help the user improve the coherence and flow of their research brief.

connectionStrength should be a number between 0 to 1 inclusive, be fair and critical about the coherence between the facets connected, if the facets are not specific enough to evaluate the connection, rate it with a low number and suggest further clarification. Be strict and use the whole range of values so the differences between strong and weak connections are distinct. Do not use 0.5 as that is the default neutral value.

If papers are referenced, use the citation tag in the format of @ref[corpusId].

## I.6 Research Brief Generation Prompt

You will create a one page academic research brief based on the provided idea facet nodes, their connecting edges, and literature references to collected papers. Your goal is to create a coherent and structured research brief that integrates the facets logically and effectively.

Here are the idea facet nodes and their connections:

```
{research_ideas}
```

Here are the collected papers for literature references, ONLY use these provided papers:

```
{papers_data}
```

Follow these steps to generate the research brief:

1. Review the idea facet nodes and their connections to understand the logical flow of the research brief.
2. Identify literature references from the collected papers that support or relate to each idea facet.
3. Develop a one page academic research brief that includes the following sections in this order, each section should be around 3 sentences:
  - Title: A concise and descriptive title for the research brief.
  - Problem Description and Research Question: Connect and expand on the content covered in Problem Description and RQ nodes.
  - Proposed Design and Solution: Connect and expand on the content covered in Proposed Design and Solution nodes.
  - Evaluation Method: Connect and expand on the content covered in Evaluation Method nodes.
  - Contribution and Impact: Connect and expand on the content covered in Contribution and Impact nodes.



Include in-text citations to the literature references (e.g. [1, 3, 5] refer to the first, third and fifth publication in the literatureReferences) where relevant throughout different sections. Ensure that the research brief is coherent, structured, and effectively integrates the idea facets with the literature.

Keep as much of the text in the node title and content as possible. Only make minor additions or modifications to connect and expand on existing nodes.

Provide your response strictly in the following format:

```
{
  "researchBrief": {
    "title": [Title of the research brief],
    "problemDescription": [Problem description and research question section of the research brief],
    "proposedDesign": [Proposed design and solution section of the research brief],
    "evaluationMethod": [Evaluation method section of the research brief],
    "contributionImpact": [Contribution and impact section of the research brief],
  },
  "literatureReferences": [
    {
      "citation_id": [id of the citation used to refer to the paper in the research brief],
      "paper_title": [Full Paper Title],
    },
    ...
  ]
}
```

Ensure that the research brief is coherent, structured, and effectively integrates the idea facets with relevant literature references.

Only respond with the JSON object as specified above, do not include any additional information or summaries.

## I.7 AI Chat Q&A Response Prompt

You will be generating a response based on the user's prompt, the collected paper information and summary, and the user's research ideas. Your goal is to provide a concise and informative response that addresses the user's prompt and references the relevant literature.

Here is the user's prompt:

```
{user_prompt}
```

Here are the papers data, ONLY use these provided papers:

```
{papers_data}
```

Here is the user's research idea:

```
{research_ideas}
```

Based on the prompt, the papers, and the research ideas, generate a concise response in the following format:

```
{
  "litReviewResponse": "Response to the user's prompt if provided. If the response references a paper, add a citation tag in the format of @ref[corpusId]"
}
```

Only respond with the JSON object as specified above, do not include any additional information or summaries.

## I.8 Literature Review Summary Prompt

You will be generating a concise one-paragraph literature review based on provided list of corpus\_id:

```
{corpusIds}
```

Here are the papers data, ONLY use these provided papers:

```
{papers_data}
```

Refer to the papers by their short-form titles with citation format (AuthorLastName et al., year) and focus on the key method, findings, limitations, and future works of the papers.

Output in the following format:

```
{
  "litReviewSummary": "One paragraph summary of the literature review based on the provided papers and prompt. If the summary references a paper, add a citation tag in the format of @ref[corpusId]"
  "corpusIds": [List of corpusIds of the mentioned papers to the user's prompt.]
}
```

Only respond with the JSON object as specified above, do not include any additional information or summaries.

## I.9 Literature Review Analysis Prompt

You will be generating a literature analysis based on provided paper summaries and current research ideas. First, here are the summaries of relevant papers, ONLY use these provided papers:

```
{paper_summaries}
```

```
Next, here are the current research ideas:
{research_ideas}
```

Analyze the paper summaries and research ideas:

1. Identify the main themes and concepts present in both the papers and research ideas.
2. Look for connections, similarities, and differences between the existing literature and the proposed research ideas.
3. Consider how the existing literature supports or challenges the research ideas.

Generate a concise literature review with the following 2 sections:

1. Highlight key sections from the papers that are most relevant to the research ideas, whether it's research problem statement, method, evaluation, or contribution and impact. Discuss how the section relates to the proposed research ideas, noting any knowledge gaps in existing works or areas where the new ideas extend current knowledge.
2. Conclude next steps for the researcher to address based on the literature review and the proposed research ideas, provide actionable steps like pivot the research question, refine existing methods, or develop evaluation.

Format your response as follows, remove the placeholders and replace them with the relevant information:

```
{
  "analysis": [
    {
      "section_title": "[One-line summary of the relevance between the paper section and the research idea.]",
      "paper_title": "[Full Paper Title. Add a citation tag in the format of @ref[corpusId]]",
      "corpus_id": "[Corpus ID 1]",
      "key_section": "[3-5 sentences description of the relevant content in the key section, don't just describe a framework/method/result was proposed, but what it is about.]",
      "connection_to_ideas": "[2-3 sentences description of how this section relates to the proposed research ideas, name the specific idea being discussed.]",
      "next_steps": [
        "[Actionable steps for the researcher to address based on the literature review and the proposed research ideas.]",
        ...
      ],
    },
    ...
  ],
}
```

Only output the above JSON object. Ensure that your literature analysis is concise, coherent, and effectively synthesizes the information from both the paper summaries and research ideas. Refer to the papers by their short-form titles with citation format (AuthorLastName et al., year) in key\_section and connection\_to\_ideas.

## I.10 Node Literature Review Prompt

You will be analyzing the user's research idea based on related literature. Your goal is to provide a comprehensive literature-based analysis of the particular research idea, identify the most relevant papers and sections, and suggest refinements based on the existing literature.

Here's the user's initial research idea:

```
Idea Facet: {idea_facet}
Title: {title}
Content: {content}
```

Now, review the following list of related papers, including their titles, TLDRs, abstracts, and section summaries, ONLY use these provided papers:

```
{paper_summaries}
```

Consider the following summary of all related works in relevance to the user's ideas:

```
{lit_review_summary}
```

To analyze the research idea based on the literature, follow these steps:

1. Reading the user's research idea, related papers, and summaries.
2. Identifying key concepts and themes in the idea and comparing these with the related papers.
3. Determining the most relevant papers based on research questions, methodologies, and findings.
4. Highlighting specific sections in these papers that directly relate to the idea.
5. Suggesting refinements to the research idea by addressing gaps, incorporating methods, adjusting scope, or considering alternative approaches.

Your output has three facets:

**most\_relevant\_sections:** List the 1-3 most relevant sections from papers, explaining why each is particularly important to the research idea. Include the paper title and a brief explanation of its relevance. For each of the most relevant papers, identify the key section that are particularly important to the research idea. Explain why these sections are significant and how they relate to the proposed research.

**suggestions:** Offer 1-2 specific suggestions for refining or improving the research idea based on your analysis of the literature. Explain the rationale behind each suggestion and how it could strengthen the research.

Remember to support your analysis with specific references to the papers and their contents. Be objective in your assessment and provide constructive suggestions for improvement.

Present your analysis in the following format in JSON:

```
{
  "most_relevant_sections": [
    {
      "section_title": "[One-line summary of the relevance between the paper section and the research idea.]",
      "paper_title": "[Full Paper Title, add a citation tag in the format of @ref[corpusId]]",
      "key_section": "[3-5 sentences description of the relevant content in the key section, don't just describe a framework/method/result was proposed, but what it is about.]",
      "connection_to_ideas": "[2-3 sentences description of how this section relates to the proposed research ideas, name the specific idea being discussed.]"
    },
    ...
  ],
  "suggestions": [
    "[Actionable steps in 2-3 sentences for the researcher to address based on the literature review and the proposed research ideas.]",
    ...
  ]
}
```

Only output the above JSON object and NOTHING ELSE. Ensure that your literature analysis is concise, coherent, and effectively synthesizes the information from both the paper summaries and research ideas. Refer to the papers by their short-form titles with citation format (AuthorLastName et al., year) in key\_section and connection\_to\_ideas.