Recipe 2.0: Information Presentation for AI-Supported Culinary Idea Generation

Angel Hsing-Chi Hwang[1,2] **Samy Badreddine**[1] **Frederick Gifford**[1] **Tarek.Besold**[1]

[1]Sony AI

[2]Cornell University

{Angel.Hwang, Samy.Badreddine, Frederick.Gifford, Tarek.Besold}@sony.com

Abstract

We examine the effective design of information presentation (IP) in AI-supported tools to support culinary creativity. We reviewed the literature to propose a conceptual framework to guide the design of IP, such that it takes into account key constructs and processes of idea generation (i.e., the motivation, subject, and content of information presented to facilitate ideation). In Part 1, we collected feedback from professional chefs regarding concerns and needs around AI tools for culinary innovation. In Part 2, we performed a content analysis on popular culinary content that inspired users to ideate. In Part 3, we designed interactive prototypes based on these insights and conducted a large-scale user study (N = 250). We found that showing the cause-and-effect logic of cooking by demonstrating information in a "what if..." format encouraged new ideas. Novice users were most motivated by understanding the ingredient' constraints and learning cooking practices' rationales. Through this work, we present design implications for AI-supported idea generation and call for more efforts to study how AI can augment human creativity and other open-ended decision-making processes.

Introduction

As artificial intelligence (AI)—e.g., in the form of generative systems—increasingly enters the creative space, we see the need to establish more effective and sustainable forms of human-AI collaboration (Gillies et al. 2016; Kamar 2016). To contribute to this vision, we see ***AI-supported idea generation*** (ASIG) as a promising area for further exploration:

1. Though the majority of work in AI-supported decision-making focuses on building tools that attain more accurate and optimal solutions, abundant decision scenarios are open-ended in nature and do not necessarily entail a single best solution. Understanding how AI can support idea generation can contribute to the latter type of decision-making, which remains under explored.
2. Building on (1.), we see idea generation as an avenue where users can bring their unique strengths in the decision process, while AI offers support that may be hard to attain by humans themselves (e.g., synthesizing and presenting large amount of data). In such dynamics, we see genuine collaboration between the two parties.
3. Compared to obtaining creative end products directly from AI, supporting users at the earlier stage of a work process (i.e., during idea incubation) allows more input and control from users, encouraging a more engaged form of human-AI interaction.

We study ASIG for culinary innovation which allows us to examine the research topic through a number of critical nuances. Both users and AI can bring in unique contributions during the work process, forming authentic collaboration. For instance, AI can summarize and present large amounts of food data (e.g., recipes, nutrition or chemical information of ingredients, as well as food waste and the environmental impact of cooking) which may be difficult for users to collect and digest on their own. Meanwhile, human cooks hold a number of decision factors that may not otherwise be shared by the machine, such as experience of tastes, context for cooking, cultural background, and personal preferences. Taking these human-centered factors into account can inform how AI can support resolving open-ended questions where humans' subjective values – instead of standard metrics – are used for evaluation. One can produce a dish that scores high on nutritional benefit and low on environmental impact, but the success of the cooking idea ultimately depends on whether it tastes good to human diners.

As an application domain, compared to other more specialized practices (e.g., healthcare, legal, finance), cooking covers a wide spectrum of decision-makers, ranging from highly skilled chefs to casual home cooks. Upon reaching an idea, users themselves have to execute the cooking process, meaning this topic requires them to be more involved in the decision process irrespective of expertise.

As we reviewed prior work and theoretical foundations, we also see cooking as a suitable domain to address the various components and cognitive processes of idea generation. Against that backdrop, we conducted a three-part research process to examine the use of AI for culinary idea generation. First, we conducted a formative interview with professional chefs to grasp their opinions toward ASIG tools for culinary innovation. Next, to also seek insights from more diverse communities and general users, we took a computational social science approach to analyze key features in effective online culinary content (e.g., cooking tutorials) that inspire cooking ideas. Based on insights from our content analysis, we designed five different prototypes of interactive tools that support cooking idea generation, and we conducted a large-scale user study to collect user feedback and design implications.

Our work makes practical and theoretical contributions along three dimensions. First, based on a literature review, we synthesize a framework with key constructs of idea generation and show how it can be adopted to inform design of ASIG tools, leading to more fruitful ideation outcomes than those of existing tools. Second, we assemble perspectives from diverse communities, providing concerns that need to be addressed and directions for improving emerging technologies in the culinary domain. Finally, we examine outcomes of taking these directions into actual design, as well as their effect on users' experience levels. We identify the importance of informing users about the logic and constraints in idea generation, and show that these design practices are particularly helpful to inspire new ideas among novice users. In sum, our work contributes to exploring AI's potentials in culinary innovation, while expanding the broader knowledge space of AI's applications to more open-ended, personalized, and creative decision-making.

Background and Related Work

Cognitive Constructs of Idea Generation

While idea generation in each domain applies unique expertise and practices, existing research identifies several universal principles and elements that are crucial to the process of generating and formulating ideas (Girotra, Terwiesch, and Ulrich 2010; Toubia 2006). Most previous work builds on the notion that creative ideas need to be both "new and appropriate" (Sternberg and Lubart 1999; Kaufman and Sternberg 2010). For idea generation, appropriateness needs to take into account the capabilities and constraints of the involved materials and methods (Medeiros et al. 2018; Rietzschel, Slijkhuis, and Van Yperen 2014). Knowing *"what works"* and *"what doesn't work"* become equally important.

Then, how does one progress from ideas that are simply appropriate to those that are truly innovative? Design researchers and practitioners have proposed outlining the governing logic and then encouraging individuals to adapt and apply it to their own ideation problems (Plucker 2004; Policastro and Gardner 1999). In other words, knowing *"how things work"* enables one to apply functional practices across domains and topics, leading to ideas that are out-of-the-box but still possible to execute. Additionally, research in cognitive science found inhibitory control to be a crucial cognitive function for idea generation (Cassotti et al. 2016; Flaherty 2005). Specifically, humans have the tendency to adopt cognitive heuristics (i.e., mental "shortcuts" that allow one to make decisions and take actions quickly). This indicates that ideas that are more common and obvious often come to mind first during idea generation. Therefore, whether a person can suppress these highly accessible, "easy" ideas becomes decisive for innovation.

Putting these considerations together, a handful of research looked at the process of idea generation, which typically entails the following stages: (1) identifying opportunities and problems, (2) acquiring knowledge and collecting information, (3) generating preliminary ideas, (4) evaluating and further developing ideas, and (5) implementing, revising, and improving ideas (Shneiderman 2002; Treffinger, Isaksen, and Stead-Dorval 2006). Throughout these stages, individuals are said to alternate between two key cognitive approaches: *divergent thinking* and *convergent thinking*. With divergent thinking, one would cast a wide net, trying to collect as much information, identify as many opportunities and gaps, and lay out as many potential ideas as possible (Runco and Acar 2012; Acar and Runco 2019). Divergent thinking is often seen as an indicator of creativity – after all, starting broad would provide a greater number of "candidate" ideas for a person to further develop, which, again, can lead to higher quality of ideation outcomes. However, not all working materials and initial ideas are worth further pursuing and developing, and some may not even be feasible for execution. Convergent thinking narrows down the scope, prioritizes what may be more relevant, and identifies novel problem space (Cropley 2006; Simonton 2015). In particular, convergent thinking supports evaluation, a critical step that allows one to focus on more promising ideas and further develop them, leading to true innovation. Therefore, the use of both divergent and convergent thinking is common in popular ideation techniques, such as brainstorming (Larey 1994), Linkographs (Goldschmidt 2016), and the Double Diamond model (West et al. 2018).

AI-Supported Idea Generation

Existing work in computational creativity has demonstrated AI's potential in automatically and independently generating creative content as well as in augmenting human creativity through the provision of tools or creative collaborators (Davis et al. 2015). Here, we focus on the latter and discuss several cohesive themes from proposed approaches to designing ASIG tools.

AI tools are supposed to ***align well with users' mental and work processes of idea generation***; for instance, (Schleith et al. 2022) proposed to use six key actions (i.e., learn, look up, relate, monitor, extract, and create) as creative prompts to guide users through the ideation process and land on novel ideas. Also, in order to elicit ideas above the ordinary, various studies emphasize the importance of ***creating interactive experiences***. Drawing from work in social robotics, one common approach is to "bounce ideas back and forth" with AI tools, enabling users to take turns and shift initiative between themselves and the tools (Lin et al. 2020; Gero, Liu, and Chilton 2022). Such experiences—especially when AI tools provide unexpected content—were found to spark inspiration and unblock ideation bottlenecks.

Another successful strategy is to ***provide users with more opportunities to collect feedback (and thus evaluate their ideas)***, especially when giving users control to customize the types of feedback to their needs. For instance, (Wu, Terry, and Cai 2022; Wu et al. 2022) created a writing support tool and a music composition tool with large language models while allowing users to create their own interactive experience through prompt chaining. This allowed users not only to better understand how the AI tools worked but also how they could improve their ideation content through more personalized, granular feedback.

Finally, several studies stress AI's capabilities of ***integrating and presenting large amounts of information*** to help users acquire knowledge and enrich sources of inspiration.

Information Presentation (IP) to Support Ideation

We see great promise in the capacity of AI systems to extract information from (often large amounts of) data. To harness this strength, we focus on how AI tools can support idea generation through effective information presentation (IP). How individuals encode information can impact whether they can leverage it for idea generation and creative problem-solving (Mumford et al. 1996; Sawyer 2011). Specifically, presentation that can help users focus on factual information, discount the irrelevant, and connect the dots, can drive higher-quality ideas (Mumford et al. 1996; Mobley, Doares, and Mumford 1992). (Wang and Nickerson 2017) studied tools and systems that support creative work through assembling and presenting information from digital libraries and the web, focusing on task-specific knowledge, and enabling more efficient information search. The review found effective creativity support systems often serve at least one of three functions: structure and organize knowledge hierarchically, synthesize and provide various perspectives to an existing topic, and filter and offer information based on its relevance.

Various approaches to implementing these functions have been examined for their effectiveness on creative problem-solving and ideation. Early work attempted to present information step-by-step depending on users' different stages of design thinking or work processes (Elam and Mead 1990; Marakas and Elam 1997). Alternatively, (Althuizen and Wierenga 2014; Forgionne and Newman 2007) focused on offering concrete case studies and examples, in the hope that users could draw analogies between these references and their own work as inspiration for new ideas. (Wang and Ohsawa 2013; Jenkin et al. 2013) designed tools to extract and offer key notes from large amounts of information, directly highlighting important and novel points for users. Across the board, visualization was found as a particularly helpful means of information consumption (Kohn, Paulus, and Korde 2011).

Beyond idea generation, IP is critical to designing AI tools for decision support in other domains. Besides addressing common challenges about users' trust in AI and its explainability (Goebel et al. 2018), recent work has revealed a lack of actionability as a key drawback of such systems (Yang, Steinfeld, and Zimmerman 2019). To address this issue, we construct a framework to guide designers of ASIG tools to present information while taking into account the key cognitive constructs and processes of idea generation discussed above. Our proposal entails the following three layers:

- Motivation (the "why" problem): Whether the goal of IP is to ***support divergent or convergent thinking***; namely, whether the tool should help users explore a wide range of relevant knowledge, references, and examples (divergent thinking), or focus on just one or a few sample(s) to help users funnel their thoughts to a specific end.
- Subject (the "what" problem): To inform the capabilities and constraints of ***materials and methods*** that one works with for idea generation, ASIG can present information about materials, methods, or a combination of both.
- Content (the "how" problem): To put information into tangible IP, one should address one or several of the following key questions: ***"what works", "what doesn't work", and "how things work"***. For instance, one can provide details for a material or a method that users work on (i.e., ***input of ideation***); one can explain the logic of how things work, inspiring users to generalize and apply them in another domain (i.e., ***rules of ideation***); or one can show possible outcomes and examples as sources of inspiration ***output of ideation***.

Technology for Culinary Innovation

Prior human-computer interaction work has explored various approaches to augment culinary innovation, including enhancing social engagement in cooking experiences (Isaku and Iba 2015; Svensson, Höök, and Cöster 2005) or generating 3D artifacts of food to serve as sources of inspiration (Sauvé, Bakker, and Houben 2020; Naritomi and Yanai 2021; Punpongsanon et al. 2022). Other work looked at supporting users to collect helpful information (e.g., nutrient data, cooking techniques, cookware, and recipes) in order to come up with novel cooking ideas (Baurley et al. 2020; Yoneda and Nadamoto 2018; Kato and Hasegawa 2013). Here, one of the key challenges lies in the divergent and highly subjective notions of what users consider as useful and meaningful data (depending, e.g., on taste preferences, dining experiences, and cultural background). Correspondingly, efforts have been made to build more personalized recommendation systems (Chen et al. 2021) as well as IP and visualization tools that effectively reveal insights and spark inspiration for cooks (Chang et al. 2018). Another approach to leveraging insights from food-related data is through directly generating new cooking ideas, e.g., by providing food pairing recommendations (Gim et al. 2022), by suggesting how food ingredients and cooking methods pair well together (Baurley et al. 2020; Yoneda and Nadamoto 2018; Kato and Hasegawa 2013), or by generating entire new recipes from food tutorial clips (Fujii et al. 2020). Finally, across the various approaches to supporting culinary innovation, there is a growing trend to adopt and present information from multiple domains, e.g., simultaneously providing information about food and its environmental impact (Kuznetsov, Rodriguez Vega, and Long 2022; Sauvé, Bakker, and Houben 2020).

Several of these IP approaches have been implemented and productized; for instance, recipe recommendations and food pairing functions are shown in various commercial applications (e.g., BigOven, PlantJammer, FoodPairing). However, these applications face the challenge to fit well and embed into users' existing cooking practices. In particular, it is often required that users possess clear cooking goals prior to starting information collection. This counteracts with more common approaches to creative ideation, which more often start from broader, divergent scopes, and later converge to more specific paths (Urban Davis et al. 2021). Moreover, while existing literature has explored a variety of materials and formats targeting specific cooking elements (e.g., ingredient, cookware, time, temperature) to support culinary innovation, more principled guidelines to inform this design space are lacking.

Regarding the general study of ASIG, creativity-support tools for cooking offer an interesting application domain in that the previously described theoretical constructs of idea generation can be practically operated and assessed:

- Motivation (the "why" problem): Divergent approaches to present culinary information would demonstrate the various dishes that an ingredient or technique can be applied to, while convergent reasoning would highlight how information can be applied to a specific dish.
- Subject (the "what" problem): Under the context of cooking, subjects that users work on to generate new ideas include ingredients (materials) and cooking techniques (methods).
- Content (the "how" problem): For cooking, one can focus on presenting general good practices to cooking (i.e., rules), informing specifics of an ingredient or technique (i.e., input), or showing examples of cooking outcomes (i.e., output).

In the following we examine how different approaches to IP can support generation of creative cooking ideas. We begin with understanding users' perspectives on existing AI tools for culinary ideation to identify strengths and pain points of these tools and possible means for improvement. While the majority of these creativity-support tools target culinary professionals, we further ask whether their current advantages and disadvantages are applicable to general users and how adaptation can be achieved. Taking these general directions for improvement and/or adaptation, we then study how to execute them into actual content and design. Here we take a computational social science approach to observe and pursue inspirations on the public discourse. Specifically, we perform content analysis on popular cooking tutorials and their audiences' responses in order to find content and design strategies that generally work well and elicit creative cooking practices. Finally, informed by the content analysis, we execute effective content design into five interactive prototypes and launch a large-scale user study to examine users' behavior, experience, and feedback.

Part 1: Perspectives on Existing Tools

A wide range of technologies for culinary innovation have been designed and targeted at expert users (i.e., professional chefs). As well as a general interview series with eighteen chefs, and in-depth interviews with six chefs and two industry advisors, we conducted a semi-structured, formative interview with 16 professional chefs from Europe, North America, South America, and Asia to capture their feedback on the advantages and drawbacks of some existing tools, specifically comparing effective recommendation and food pairing functions. Here, we summarize insights that are relevant to the current work.

First, experts favored the specificity of information; for them, it is the details that matter and serve as effective inspiration. For instance, when providing information on an ingredient, it is important to specify its origin and processing techniques. Also, presenting information about *either* ingredient *or* method was seen as an ineffective approach to professional cooks. In general, chefs favoured a convergent approach to IP; they showed little interest in seeing a large amount of general information or its summary, as most of it usually seemed already familiar and less inspiring to them. Instead, they preferred focusing on specific, commonly unseen ingredients or techniques as sources of inspiration. Experts' opinions on IP can be summarized as follows:

- Motivation (the "why" problem): Adopt a ***convergent approach*** and present mainly selective, highly relevant, and previously unseen information.
- Subject (the "what" problem): Present information about both ***materials and methods*** side by side.
- Content (the "how" problem): Present detailed information about the ***input*** for idea generation.

Are these insights applicable to designing ASIG tools for general users? Generally, layman users often lack specific goals when initiating idea generation processes (Sawyer 2011). Taking a convergent rather than divergent approach to design ASIG tools may therefore not be as effective for non-experts. Furthermore, jointly presenting different types of information (i.e., combining ingredient- and method-related information) may cause information overload for non-experts, as they are less apt at connecting the dots across abundant information than their expert counterparts (Mumford et al. 1996; Mobley, Doares, and Mumford 1992). Finally, providing highly specific, detailed information to general users may also reduce the flexibility for innovation (Kletke et al. 2001). Unlike professional chefs, who know how to substitute one ingredient with another, amateur cooks may not know the alternatives to specialized items. Therefore, in the following, we gathered insights from a broader, more general user base.

Part 2: Effective Content to Support Ideation

Method

To get insights from general users and to see how IP can be executed in effective content design to inspire ideation, we took a computational social science approach to review popular culinary content on YouTube and corresponding audience responses. We chose this medium for content analysis as cooking tutorials on YouTube are one of the most common resources where the general population seek cooking information (Benkhelifa and Laallam 2018). We focus on understanding what types of content—as well as their approach to IP—can elicit more creative ideas with lay users.

Cooking videos for content analysis We examined the top 10 most popular channels under YouTube's food and cooking category, and used `BeautifulSoup` from Python to crawl the links and metrics of the top 10 most viewed cooking videos in each channel and the top 100 most engaged comments for each video. In total, we obtained 100 cooking videos and 10,000 comments for the analysis.

Coding approach To analyze the video content and understand users responses to the culinary information presented in these videos, we first viewed the videos and

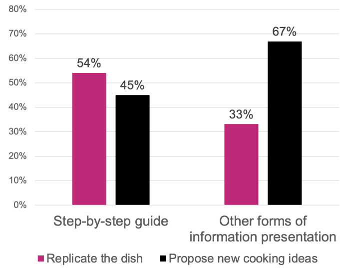

Figure 1: Viewers' responses to popular cooking videos

remarked how information was presented in each video through our proposed framework. We then examined the audience feedback to each video through two specific lenses: whether viewers mention cooking ideas and/or actions taken upon viewing the video, which helps us investigate if the video was effective in triggering idea generation; and how viewers described the content and what they found beneficial from each video, which allows us to distill features for effective content design for culinary innovation.

Results

Among the examined culinary videos, a considerable portion (28%) simply followed a conventional, step-by-step tutorial style to show viewers how to cook a dish in a specific way from the start to the end. The remaining videos demonstrated cooking information through various formats, such as doing experiments to show how switching ingredients and applying different cooking techniques may influence cooking outcomes. The majority of these cooking tutorials applied convergent approaches (76%), featuring the production of only one dish in each tutorial, and covered both input (specified information about ingredients: 69% of all videos; specified information about cooking techniques: 88% of all videos) and output of cooking (showed results of cooking: 70%). Furthermore, as many as 72% of videos applied examples to present culinary knowledge. In the contrary, IP emphasized "what worked" for each dish (79% of all videos), while the logic of "how things worked" (9%) and "what didn't work" (32%) were less mentioned.

Examining feedback to the cooking videos, we noted different patterns in viewers' comments depending on how culinary information was presented. Whether or not a video used step-by-step cooking instructions (i.e., presenting one way to "do the dish right") has a salient effect on viewers' ideas and intended actions. After watching a video showing how a dish was made from the start to the end, viewers were more likely to express interest in making the dish (mentioned among 41% of comments) with more than half (54%) intending to replicate what they learned from the tutorial step-by-step. Two-thirds (67%) of those who watched videos presenting culinary information in more diverse ways would instead propose their new ideas for cooking.

We examined IP formats that inspired viewers to generate their own cooking ideas. Viewers responded particularly positively to three types of content. First, the audience was interested in ***understanding the "cause-and-effect" rationales*** behind cooking. Participants proposed more new cooking ideas after watching videos that asked numerous "what if..." questions and demonstrated how changing one component in cooking (e.g., varying a cooking technique, ingredient, or cookware) would influence the outcome of the dish (e.g., flavor profile, texture). For instance, in one of the videos, the chef varied the time and temperature to sear a piece of steak and examined its tenderness. Experimenting and showing how different methods led to distinct outcomes was seen as particularly informative and motivated viewers to come up with their own plans to create steak dishes that suited their taste preferences. Secondly, viewers prefer information that ***advises how they can adapt what they learned*** from a cooking video to their own kitchen. Understanding how a cooking technique can be generalized to handle different ingredients and produce various dishes was, thus, found especially helpful and triggered more proposals of new cooking ideas. Third, ***presenting constraints*** of an ingredient or method is valuable. As mentioned in viewers' comments, this information helped them to understand what may have gone wrong in past cooking attempts and to come up with new ideas avoiding those mistakes. Likewise, seeing chefs' trial-and-error processes was perceived as helpful to comprehend what would not work and to come up with ideas that could make a cooking plan successful. Moreover, these three approaches to presenting culinary information not only encouraged idea generation, but they were also reported to enhance the positive effect—such as fun and enjoyment—of cooking processes, which is yet another motivating factor of innovation (Sawyer 2011).

Summarizing our findings from studying general audience comments on publicly available cooking videos:

- Motivation (the "why" problem): Public comments showed no particular preference between presenting information in either convergent or divergent approaches.
- Subject (the "what" problem): Public comments showed no particular preference between understanding information about materials and methods.
- Content (the "how" problem): Viewers expressed particular interest in understanding the cause-and-effect (i.e., explaining rules and "how things work"), adaptability (i.e., explaining rules, "how things work", and "what works"), and common mistakes in cooking (i.e., information related to constraints of input and "what doesn't work").

Part 3: Designing AI Tools to Support Culinary Ideation

Method

Based on insights from the previous sections, we designed five different prototypes, each applying a unique IP strategy to support culinary innovation. In Table 1, we summarize how each design of the five conditions corresponds to key constructs of idea generation. We then conducted a between-subject user study to examine whether and how users' performance differs when adopting different idea generation tools. Participants read an informed-consent form and completed a pre-survey on their cooking practices and habits. As a key part of the study asks participants to gen-

Table 1: Approaches to information presentation (IP) and their corresponding outcomes

Insights (Part 1 & 2) or Condition (Part 3)	Motivation of IP (The "Why" Problem)	Subject of IP (The "What" Problem)	Content of IP (The "How" Problem)	Show "what works"	Show "what doesn't work"	Explain "how things work"	Show examples	Ideation Outcomes	Affected Users
Feedback from experts (Part 1)	Support **convergent thinking**	Material + Method	Input	✓				(Not applicable)	(Not applicable)
Popular content on public discourse (Part 2)	Support convergent thinking: 76% Support divergent thinking: 24%	Material: 69% Method: 88%	Input: 95% Output: 70% Rule: 32%	79%	9%	29%	72%	(Not applicable)	(Not applicable)
Feedback from audiences' comments (Part 2)	(Not available)	No noticeable preference	Input + Rule (the cause-and-effect, adaptability, and common mistakes)	✓	✓	✓		(Not applicable)	(Not applicable)
"Baseline" condition (Part 3)	Support **convergent thinking**	Material + Method	Input	✓				(+) for revising ideas (-) for generating new ideas	Encourage experts to generate new ideas but lead novice users to copy
"Pairing" condition (Part 3)	Support **divergent thinking**	Material	Input	✓				(+) for revising ideas (-) for generating new ideas	Encourage novice users to revise existing ideas
"Generalizable Method" condition (Part 3)	Support **divergent thinking**	Method	Rules + Output	✓		✓	✓	Encourage new ideas but also replication	No particular affected users
"Constraint" condition (Part 3)	Support **divergent thinking**	Material	Rules		✓	✓		Encourage new ideas but also replication	Encourage novice users to generate new ideas and revise existing ones
"What if" condition (Part 3)	Support **divergent thinking**	Material + Method	Rules + Output	✓		✓	✓	Enhance generation of new ideas	Encourage most users, especially those with more experiences, to come up with new ideas

erate a cooking idea, we used the pre-survey to screen and exclude participants who had no cooking experience.

Participants were randomly assigned to one of the five prototypes, to explore, use and come up with cooking ideas. While exploring, they composed a recipe plan, which entails describing a dish they would like to create, the ingredients needed, and the steps they would take to cook the dish. Participants were asked to pull up and view their assigned prototype alongside their recipe planning screen. They did not have to memorize and could instead refer to information in the prototype when generating ideas. As last step, participants filled out a short exit survey to reflect their user experience and reported their demographic data. The entire study took around 30 minutes to complete, and participants received $10 to compensate their time and participation. The study was reviewed and approved by the Research Ethics Board at the authors' affiliation.

Interactive prototypes We used Figma to create interactive prototypes for the five conditions of our user study. The designs of the conditions were informed by the findings we reported in Part 1. To create equivalent initial states, all conditions started with showing users a home page with four dishes as sources of inspiration. Participants could click to explore each dish. From there, each prototype applied a unique strategy to present information and inspire idea generation. The five conditions include: (1) **baseline condition**: a classic step-by-step recipe; (2) **ingredient pairing**: a recipe showing the molecular and recipe fit of ingredients used; (3) **constraint condition**: a recipe showing the constraint of each ingredient used; (4) **generalizable method condition**: a recipe showing the cooking techniques used and other cuisines that can be made applying the same methods; (5) **what-if condition**: a recipe showing possible outcomes as one switches the ingredients and cooking techniques. Each condition had a similar amount of information for exploration with each dish having 6 pages of content to click through and each condition having in total 24 pages of information to consume. To check if participants had explored the content on their assigned prototypes, we included a page code on each piece of the content, and participants were asked to record and report the code of pages they had viewed in the exit survey. None of our participants failed this validity check.

Measures of the user study We collected three main categories of data, including users' existing cooking practices and experience (frequency, expertise, and years of cooking), the recipe idea they planned out, and their user experience during idea generation (measured through the usability scale (Bangor, Kortum, and Miller 2008) and the self-efficacy scale (Sherer et al. 1982)). With participants' cooking plans, we coded each idea into one of three types: (1) *copy* indicates a participant was simply copying the idea from one of the four source dishes; (2) *revision* indicates participants adopted one of the four source dishes but made a twist to its original recipe (e.g., swap ingredients or replace a cooking technique); (3) *new* indicates a participant came up with their own cooking idea that are distinct from the four sample dishes presented in the prototype.

Participants We recruited 250 participants on Amazon Mechanical Turk (AMT) through the following screening criteria: participants were located in the United States and have completed more than 1000 HITs with a HIT approval rate greater than 98% at the time when the study was conducted. Average age of participants was 36.30 ($S.D. = 11.17$). The majority of participants was Caucasian (70%), while 54% identified as female and 40% as male. As stated in our recruitment message and research consent form, participants also needed to have at least one month of cooking experience at the time they participated in the study. Overall, participants have, on average, 16.29 years of cooking experience ($S.D. = 11.63$). The final sample size was predetermined by conducting a pilot study and performing a power analysis based on the pilot data.

Results

We saw that **the types of recipe ideas participants came up with differed significantly by the prototype conditions** ($\chi^2 = 23.898$, $p = 0.002$). Specifically, those who explored the *what if* condition generated the most new recipes (76.32%), followed by those who explored *generalizable methods* (55.81%) and *constraints* (56.10%) of ingredients. Participants who viewed the *baseline* (47.73%) and the *ingredient pairing* conditions (40.48%) came up with the fewest new ideas. Still, although not inspiring idea generation, the last two prototypes seemed effective in providing information that can be useful for adaptation. Correspondingly, we saw the largest portion of participants who revised recipes from the four source dishes (baseline: 25.00%; ingredient pairing: 26.19%). Also in their descriptions of recipe plans we saw the highest percentages of mentions of how participants were able to apply learned information for the *baseline* and *pairing* prototypes.

We also ascertained the degree of complexity of participants' recipe ideas by examining the number of ingredients and the number of planned out steps. The number of ingredients used in different conditions differed marginally by conditions ($F = 2.043$, $p = 0.089$). Participants who viewed the *baseline* condition ($M = 9.43$, $S.D. = 4.61$) applied the most ingredients, followed by those in the *what if* condition ($M = 9.03$, $S.D. = 3.80$), and those in the *ingredient pairing* condition ($M = 8.20$, $S.D. = 3.80$). There was no significant difference either in the number of steps planned out in participants' cooking ideas ($F = 0.502$, $p = 0.735$) or in participants' self-reported user experience and ease of use in the different prototypes explored ($F = 1.672$, $p = 0.158$). This rules out the alternative explanation that participants were more likely to come up with their own cooking ideas simply because they couldn't acquire or comprehend the sources of inspiration in a prototype.

We used the number of years participants spent cooking multiplied by their cooking frequency as a proxy to assess their cooking experience levels. Overall, we saw a marginal effect of participants' cooking experience on their idea generation outcomes ($\beta = 0.02$, $S.E. = 0.01$, $t = 1.68$, $p = 0.095$). Moreover, **participants' existing cooking experience moderates their ideation outcomes**, as we found an interaction effect between participants' experience levels and the recipe condition they explored ($\beta = -0.03$, $S.E. = 0.02$, $t = -1.79$, $p = 0.075$). To be specific, for the *what if* or the *baseline* prototypes, more experienced participants came up with more new ideas; conversely, novice participants were triggered to generate more new ideas when they explored the other conditions.

Discussion

We reviewed prior literature and synthesized a theoretical framework to guide the design of IP in ASIG tools. Specifically, we proposed that effective IP should cover three key constructs of idea generation, responding to the "why" problem (i.e., whether the motivation is to support divergent or convergent thinking), the "what" problem (i.e., whether the subject entails the material or method of idea generation), and the "how" problem (i.e., whether the content addresses "what works," "what doesn't work", and "how things work" through providing information about the rules, input, or output of idea generation). We first obtained perspectives from experts regarding the benefits and shortcomings of existing tools, sought insights from the public discourse on effective content that can inspire ideation, and conducted a user study to examine the effectiveness of different interactive prototypes. We now elaborate on the various theoretical and practical implications we gather from the results.

Design Implications

The different patterns in viewers' responses to differing styles in popular culinary content reconfirmed the important roles of IP in inspiring idea generation and creative problem-solving (Mumford et al. 1996; Mobley, Doares, and Mumford 1992). Specifically, upon viewing more of a constrained, conventional style, presenting step-by-step guides to "do things right in one specific way", the audience were more likely to replicate the same course of actions instead of coming up with their own cooking approaches and ideas. As regards the design of IP that can effectively inspire ideation, we found that all three of the newly proposed interactive prototypes (*generalizable method* condition, *constraint* condition, and *what if* condition) led to more new ideas generated than what was the case for approaches seen in existing tools (i.e., *baseline* and *pairing* conditions).

In view of similar results in the *baseline* and *pairing* conditions, we realized that the shift from adopting convergent to divergent approaches per se did not necessarily lead to different idea generation outcomes. What seemed to matter more is what information was covered and how IP was designed in the actual content. Regarding IP, we first compare the outcomes of the *generalizable method* condition and *what if* condition, as these two conditions differed only in the subjects presented, with all other design factors following similar strategies. Having more new ideas generated in the *what if* condition suggests covering both the material and method as subjects of IP has a additive effect on idea generation. However, while covering two subjects at once may be informative for more experienced users, we saw that those with fewer cooking experiences in these cases ended up copying more. We also saw a similar pattern in the *baseline* condition, which also jointly covers ingredients and cooking techniques. Indeed, as seen in Part 1, professional chefs particularly requested to see detailed information about ingredients and techniques side by side. Still for lay users this design may introduce too much information at once therefore becoming less helpful.

The general public especially expressed interest in understanding the logic behind "how things work" and identifying common mistakes. We found that less experienced users performed particularly well in generating new ideas when they worked with the *constraint* condition. This highlights the effectiveness of explaining applicable rules, reasoning, and limitations of an ideation subject in IP, while, according to the data from Part 2, these design strategies were less often applied in existing content. Judging from our literature review, they are also less emphasized in prior work on

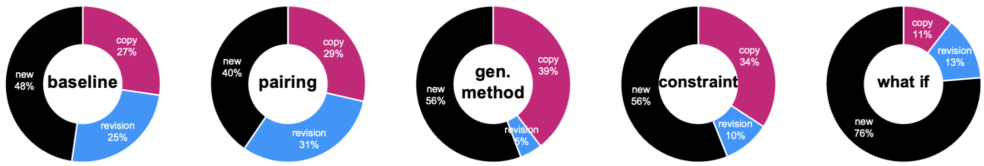

Figure 2: Breakdown of recipe ideas generated in each prototype condition

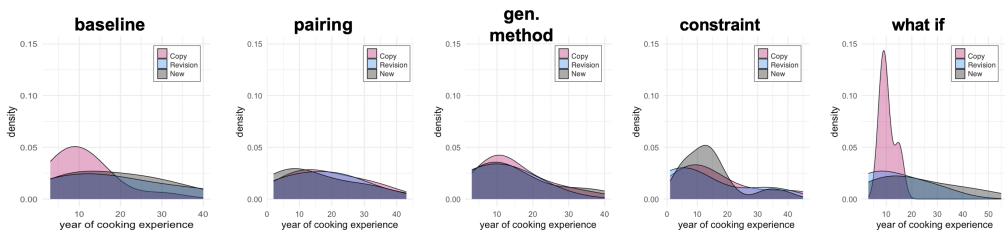

Figure 3: Density distribution of idea generation outcomes by participants' years of cooking experience in each condition

designing ASIG tools. It is also worth noting that—while the literature often emphasizes the usefulness of presenting examples as sources of inspiration—our findings from the *constraint* condition suggest that the absence of examples at least did not hinder users from coming up with new ideas.

Limitations and Future Work

We acknowledge several limitations of our current work. First, to compare how different designs of IP affects users with higher versus lower levels of experience, we used participants' years of cooking experience as a proxy to their experience level. This may not be rigorous enough and can be subject to other confounding factors. For instance, a user may have had formal culinary training and be more skilled than a home cook who has spent many more years cooking. Ideally, we would want to bring the interactive prototypes to professional chefs and conduct user studies with them as well. We encourage future work to perform more rigorous comparison between experts' and novices' responses.

Second, while we examine whether participants came up with new ideas as a key measure, we acknowledge that the nature of idea generation is much more complex and should be further explored through multiple dimensions. For instance, although participants in the *baseline* and *pairing* conditions did not generate as many new ideas, they tended to work well on revising ideas. This may respond to the general public's interest in adapting what they learned to their own cooking environment. At the same time it also suggests these forms of IP may facilitate users to learn and absorb knowledge. Because a positive relationship between learning and creativity has been found in the long run, we might observe different effects of the two conditions if we extended the study period over a longer time span. This is another potential direction which we encourage future research to pursue.

We asked each participant to plan out just one recipe idea. We adopted this approach to focus on investigating the quality instead of the quantity of participants' idea generation outcomes. Still, participants thus did not have the opportunity to compare, evaluate, and select the best idea out of a pool of candidate ideas they generated. While evaluation serves as a critical component in the full process of idea generation, we see the need for additional work to examine effective IP to support users' evaluation and selection of ideas.

More broadly our current work contributes to understanding how humans work with AI to resolve open-ended problems that rely on users' subjective, personal experiences in decision-making processes. Building on our findings, an important next step will be to compare how effective design of IP to support open-ended decision-making differs from information presented for close-ended decision-making.

In summary, our present work is an initial attempt to propose more systematic approaches to designing IP in ASIG tools. We use culinary innovation as a domain to operationalize our theoretical framework and conduct user studies, while we encourage future work to examine the topic across (and adopt our proposed framework to) other specialty areas. Additionally, collecting information is just one of the various steps in idea generation; designing AI tools to support other parts of the ideation process still remains an underexplored area offering itself for further research.

References

Acar, S., and Runco, M. A. 2019. Divergent thinking: New methods, recent research, and ext. theory. *Psychology of Aesthetics, Creativity, and the Arts* 13(2):153.

Althuizen, N., and Wierenga, B. 2014. Supporting creative problem solving with a case-based reasoning system. *Journal of Management Information Systems* 31(1):309–340.

Bangor, A.; Kortum, P. T.; and Miller, J. T. 2008. An empirical evaluation of the system usability scale. *Intl. Journal of Human–Computer Interaction* 24(6):574–594.

Baurley, S.; Petreca, B.; Selinas, P.; Selby, M.; and

Flintham, M. 2020. Modalities of expression: Capturing embodied knowledge in cooking. In *Proc. of the 14th Int. Conf. on Tangible, Embedded, and Embodied Interaction*, TEI '20, 785–797. New York, NY, USA: ACM.

Benkhelifa, R., and Laallam, F. Z. 2018. Opinion extraction and classification of real-time youtube cooking recipes comments. In Hassanien, A. E.; Tolba, M. F.; Elhoseny, M.; and Mostafa, M., eds., *The Int. Conf. on Advanced Machine Learning Technologies and Applications (AMLTA2018)*, volume 723, 395–404. Cham: Springer International Publishing.

Cassotti, M.; Agogué, M.; Camarda, A.; Houdé, O.; and Borst, G. 2016. Inhibitory control as a core process of creative problem solving and idea generation from childhood to adulthood. *New directions for child and adolescent development* 2016(151):61–72.

Chang, M.; Guillain, L. V.; Jung, H.; Hare, V. M.; Kim, J.; and Agrawala, M. 2018. Recipescape: An interactive tool for analyzing cooking instructions at scale. In *Proc. of the 2018 CHI Conf. on Human Factors in Computing Systems*, CHI '18, 1–12. New York, NY, USA: ACM.

Chen, Y.; Subburathinam, A.; Chen, C.-H.; and Zaki, M. J. 2021. Personalized food recommendation as constrained question answering over a large-scale food knowledge graph. In *Proc. of the 14th ACM Int. Conf. on Web Search and Data Mining*, WSDM '21, 544–552. New York, NY, USA: ACM.

Cropley, A. 2006. In praise of convergent thinking. *Creativity research journal* 18(3):391–404.

Davis, N.; Hsiao, C.-P.; Popova, Y.; and Magerko, B. 2015. An enactive model of creativity for computational collaboration and co-creation. In Zagalo, N., and Branco, P., eds., *Creativity in the Digital Age*. London: Springer. 109–133.

Elam, J. J., and Mead, M. 1990. Can software influence creativity? *Information Systems Research* 1(1):1–22.

Flaherty, A. W. 2005. Frontotemporal and dopaminergic control of idea generation and creative drive. *Journal of Comparative Neurology* 493(1):147–153.

Forgionne, G., and Newman, J. 2007. An experiment on the effectiveness of creativity enhancing decision-making support systems. *Decision Support Systems* 42(4):2126–2136.

Fujii, T.; Orihara, R.; Sei, Y.; Tahara, Y.; and Ohsuga, A. 2020. Generating cooking recipes from cooking videos using deep learning considering previous process with video encoding. In *Proc. of the 3rd Int. Conf. on Applications of Intelligent Systems*, APPIS 2020. New York, NY, USA: ACM.

Gero, K. I.; Liu, V.; and Chilton, L. 2022. Sparks: Inspiration for science writing using language models. In *Prof. of the 2022 ACM Designing Interactive Systems Conference*, DIS '22, 1002–1019. New York, NY, USA: ACM.

Gillies, M.; Fiebrink, R.; Tanaka, A.; Garcia, J.; Bevilacqua, F.; Heloir, A.; Nunnari, F.; Mackay, W.; Amershi, S.; Lee, B.; et al. 2016. Human-centred machine learning. In *Ext. Abstr. of the 2016 CHI Conf. on Human Factors in Computing Systems*, 3558–3565.

Gim, M.; Choi, D.; Maruyama, K.; Choi, J.; Kim, H.; Park, D.; and Kang, J. 2022. Recipemind: Guiding ingredient choices from food pairing to recipe completion using cascaded set transformer. In *Proc. of the 31st ACM Int. Conf. on Information & Knowledge Management*, CIKM '22, 3092–3102. New York, NY, USA: ACM.

Girotra, K.; Terwiesch, C.; and Ulrich, K. T. 2010. Idea generation and the quality of the best idea. *Management science* 56(4):591–605.

Goebel, R.; Chander, A.; Holzinger, K.; Lecue, F.; Akata, Z.; Stumpf, S.; Kieseberg, P.; and Holzinger, A. 2018. Explainable ai: the new 42? In *Machine Learning and Knowledge Extraction: Second IFIP TC 5, TC 8/WG 8.4, 8.9, TC 12/WG 12.9 Int. Cross-Domain Conf., CD-MAKE 2018, Proceedings 2*, 295–303. Springer.

Goldschmidt, G. 2016. Linkographic evidence for concurrent divergent and convergent thinking in creative design. *Creativity research journal* 28(2):115–122.

Isaku, T., and Iba, T. 2015. Creative cocooking patterns: A pattern language for creative collaborative cooking. In *Proc. of the 20th European Conf. on Pattern Languages of Programs*, EuroPLoP '15, 1–17. New York, NY, USA: ACM.

Jenkin, T. A.; Chan, Y. E.; Skillicorn, D. B.; and Rogers, K. W. 2013. Individual exploration, sensemaking, and innovation: A design for the discovery of novel information: A design for the discovery of novel information. *Decision Sciences* 44(6):1021–1057.

Kamar, E. 2016. Directions in hybrid intelligence: Complementing ai systems with human intelligence. In *Proc. of the 25th Int. Joint Conf. on Artificial Intelligence*, 4070–4073.

Kato, F., and Hasegawa, S. 2013. Interactive cooking simulator: Showing food ingredients appearance changes in frying pan cooking -. In *Proc. of the 5th International Workshop on Multimedia for Cooking & Eating Activities*, CEA '13, 33–38. New York, NY, USA: ACM.

Kaufman, J. C., and Sternberg, R. J. 2010. *The Cambridge handbook of creativity*. Cambridge University Press.

Kletke, M. G.; Mackay, J. M.; Barr, S. H.; and Jones, B. 2001. Creativity in the organization: the role of individual creative problem solving and computer support. *International Journal of Human-Computer Studies* 55(3):217–237.

Kohn, N. W.; Paulus, P. B.; and Korde, R. M. 2011. Conceptual combinations and subsequent creativity. *Creativity Research Journal* 23(3):203–210.

Kuznetsov, S.; Rodriguez Vega, A.; and Long, E. 2022. A study of solar cooking: Exploring climate-resilient food preparation and opportunities for hci. In *Proc. of the 2022 CHI Conf. on Human Factors in Computing Systems*, CHI '22. New York, NY, USA: ACM.

Larey, T. S. 1994. *Convergent and divergent thinking, group composition, and creativity in brainstorming groups*. The University of Texas at Arlington.

Lin, Y.; Guo, J.; Chen, Y.; Yao, C.; and Ying, F. 2020. It is your turn: Collaborative ideation with a co-creative robot through sketch. In *Proc. of the 2020 CHI Conf. on Human*

Factors in Computing Systems, CHI '20, 1–14. New York, NY, USA: ACM.

Marakas, G. M., and Elam, J. J. 1997. Creativity enhancement in problem solving: Through software or process? *Management Science* 43(8):1136–1146.

Medeiros, K. E.; Steele, L. M.; Watts, L. L.; and Mumford, M. D. 2018. Timing is everything: Examining the role of constraints throughout the creative process. *Psychology of Aesthetics, Creativity, and the Arts* 12(4):471–488.

Mobley, M. I.; Doares, L. M.; and Mumford, M. D. 1992. Process analytic models of creative capacities: Evidence for the combination and reorganization process. *Creativity Research Journal* 5(2):125–155.

Mumford, M. D.; Baughman, W. A.; Supinski, E. P.; and Maher, M. A. 1996. Process-based measures of creative problem-solving skills: Ii. information encoding. *Creativity Research Journal* 9(1):77–88.

Naritomi, S., and Yanai, K. 2021. 3d mesh reconstruction of foods from a single image. In *Proc. of the 3rd Workshop on AIxFood*, AI & Food'21, 7–11. New York, NY, USA: ACM.

Plucker, J. A. 2004. Generalization of creativity across domains: Examination of the method effect hypothesis. *The Journal of Creative Behavior* 38(1):1–12.

Policastro, E., and Gardner, H. 1999. From case studies to robust generalizations: An approach to the study of creativity. In Sternberg, J., ed., *Handbook of creativity*. Cambridge University Press. 213–225.

Punpongsanon, P.; Miyatake, Y.; Iwai, D.; and Sato, K. 2022. Demonstration of interiqr: Unobtrusive edible tags using food 3d printing. In *Adjunct Proc. of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22 Adjunct. New York, NY, USA: ACM.

Rietzschel, E. F.; Slijkhuis, J. M.; and Van Yperen, N. W. 2014. Task structure, need for structure, and creativity: Task structure and creativity. *European Journal of Social Psychology* 44(4):386–399.

Runco, M. A., and Acar, S. 2012. Divergent thinking as an indicator of creative potential. *Creativity research journal* 24(1):66–75.

Sauvé, K.; Bakker, S.; and Houben, S. 2020. Econundrum: Visualizing the climate impact of dietary choice through a shared data sculpture. In *Proc. of the 2020 ACM Designing Interactive Systems Conference*, DIS '20, 1287–1300. New York, NY, USA: ACM.

Sawyer, R. K. 2011. *Explaining creativity: The science of human innovation*. Oxford university press.

Schleith, J.; Norkute, M.; Mikhail, M.; and Tsar, D. 2022. Cognitive strategy prompts: Creativity triggers for human centered ai opportunity detection. In *Creativity and Cognition*, C&C '22, 29–37. New York, NY, USA: ACM.

Sherer, M.; Maddux, J. E.; Mercandante, B.; Prentice-Dunn, S.; Jacobs, B.; and Rogers, R. W. 1982. The self-efficacy scale: Construction and validation. *Psychological reports* 51(2):663–671.

Shneiderman, B. 2002. Creativity support tools. *Commun. ACM* 45(10):116–120.

Simonton, D. K. 2015. On praising convergent thinking: Creativity as blind variation and selective retention. *Creativity Research Journal* 27(3):262–270.

Sternberg, R. J., and Lubart, T. I. 1999. The concept of creativity: Prospects and paradigms. *Handbook of creativity* 1(3-15).

Svensson, M.; Höök, K.; and Cöster, R. 2005. Designing and evaluating kalas: A social navigation system for food recipes. *ACM Trans. Comput.-Hum. Interact.* 12(3):374–400.

Toubia, O. 2006. Idea generation, creativity, and incentives. *Marketing science* 25(5):411–425.

Treffinger, D. J.; Isaksen, S. G.; and Stead-Dorval, K. B. 2006. *Creative problem solving: An introduction*. Prufrock Press Inc.

Urban Davis, J.; Anderson, F.; Stroetzel, M.; Grossman, T.; and Fitzmaurice, G. 2021. Designing co-creative ai for virtual environments. In *Creativity and Cognition*, C&C '21. New York, NY, USA: ACM.

Wang, K., and Nickerson, J. V. 2017. A literature review on individual creativity support systems. *Computers in Human Behavior* 74:139–151.

Wang, H., and Ohsawa, Y. 2013. Idea discovery: A scenario-based systematic approach for decision making in market innovation. *Expert Systems with Applications* 40(2):429–438.

West, J.; Fusari, G.; Raby, E.; Alwani, R.; Meldaikyte, G.; Wojdecka, A.; and Matthews, E. 2018. Developing the double diamond process for implementation. In Barron, D., and Seemann, K., eds., *Proc. of the 4th Int. Conf. on Design4Health*, volume 1, 310–312.

Wu, T.; Jiang, E.; Donsbach, A.; Gray, J.; Molina, A.; Terry, M.; and Cai, C. J. 2022. Promptchainer: Chaining large language model prompts through visual programming. In *Ext. Abstr. of the 2022 CHI Conf. on Human Factors in Computing Systems*, CHI EA '22. New York, NY, USA: ACM.

Wu, T.; Terry, M.; and Cai, C. J. 2022. Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In *Proc. of the 2022 CHI Conf. on Human Factors in Computing Systems*, CHI '22. New York, NY, USA: ACM.

Yang, Q.; Steinfeld, A.; and Zimmerman, J. 2019. Unremarkable ai: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proc. of the 2019 CHI Conf. on Human Factors in Computing Systems*, CHI '19, 1–11. New York, NY, USA: ACM.

Yoneda, Y., and Nadamoto, A. 2018. Knack for cooking extraction from user generated recipe sites. In *Proc. of the 20th Int. Conf. on Information Integration and Web-Based Applications & Services*, iiWAS2018, 134–137. New York, NY, USA: ACM.