



# Designing LLM Chains by Adapting Techniques from Crowdsourcing Workflows

**MADELEINE GRUNDE-MCLAUGHLIN**, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA

**MICHELLE S. LAM**, Computer Science, Stanford University, Stanford, California, USA

**RANJAY KRISHNA**, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA and **PRIOR**, Allen Institute for Artificial Intelligence, Seattle, Washington, USA

**DANIEL S. WELD**, Semantic Scholar, Allen Institute for Artificial Intelligence, Seattle, Washington, USA and **Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA**

**JEFFREY HEER**, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA

---

LLM chains enable complex tasks by decomposing work into a sequence of subtasks. Similarly, the more established techniques of crowdsourcing workflows decompose complex tasks into smaller tasks for human crowdworkers. Chains address LLM errors analogously to the way crowdsourcing workflows address human error. To characterize opportunities for LLM chaining, we survey 107 papers across the crowdsourcing and chaining literature to construct a design space for chain development. The design space covers a designer's *objectives* and the *tactics* used to build workflows. We then surface *strategies* that mediate how workflows use tactics to achieve objectives. To explore how techniques from crowdsourcing may apply to chaining, we adapt crowdsourcing workflows to implement LLM chains across three case studies: creating a taxonomy, shortening text, and writing a short story. From the design space and our case studies, we identify takeaways for effective chain design and raise implications for future research and development.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Computing methodologies** → **Multi-agent systems**; • **Human-centered computing** → **Usability testing**;

Additional Key Words and Phrases: Crowdsourcing workflows, large language model chains, design space, case studies

---

This research was supported in part by the National Science Foundation (awards IIS-1901386 and IIS-2402647), the Office of Naval Research (Grant No. N00014-21-1-2707), and the Allen Institute for Artificial Intelligence (AI2).

Authors' Contact Information: Madeleine Grunde-McLaughlin (corresponding author), Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA; e-mail: mgrunde@cs.washington.edu; Michelle S. Lam, Computer Science, Stanford University, Stanford, California, USA; e-mail: mlam4@cs.stanford.edu; Ranjay Krishna, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA and PRIOR, Allen Institute for Artificial Intelligence, Seattle, Washington, USA; e-mail: ranjay@cs.washington.edu; Daniel S. Weld, Semantic Scholar, Allen Institute for Artificial Intelligence, Seattle, Washington, USA and Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA; e-mail: danw@allenai.org; Jeffrey Heer, Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, Washington, USA; e-mail: jheer@cs.washington.edu.



This work is licensed under Creative Commons Attribution International 4.0.

© 2025 Copyright held by the owner/author(s).

ACM 1557-7325/2025/6-ART27

<https://doi.org/10.1145/3716134>

**ACM Reference format:**

Madeleine Grunde-McLaughlin, Michelle S. Lam, Ranjay Krishna, Daniel S. Weld, and Jeffrey Heer. 2025. Designing LLM Chains by Adapting Techniques from Crowdsourcing Workflows. *ACM Trans. Comput.-Hum. Interact.* 32, 3, Article 27 (June 2025), 57 pages.  
<https://doi.org/10.1145/3716134>

---

## 1 Introduction

People use **Large Language Models (LLMs)** for assistance with both single-step tasks, like suggesting synonyms, and more complex multi-step tasks, like writing stories, editing text, and organizing ideas. Despite their widespread adoption, LLMs suffer from quality deficits like hallucinations [45, 119, 124], brittleness to prompt changes [61, 127, 170], and user interventions limited to prompting [129, 161]. These deficits are exacerbated for complex tasks, which require multiple steps in which errors can arise and propagate [44]. Unfortunately, many common use cases require multiple reasoning steps. For example, the use case of shortening text requires first deciding which portions of the text to edit and then determining how to revise each one [19].

To tackle complex tasks, recent research has turned to LLM chaining techniques. Chaining decomposes a task into multiple calls to an LLM, in which the output of one call affects the input to the next call [162]. For example, when shortening text, an LLM chain could identify verbose sentences, edit each one, and propose outputs of variable lengths by composing multiple edits. Chaining has many benefits: this strategy can enable new abilities like long-form story writing [121, 166], increase accuracy on logical reasoning questions [36, 164, 174], and enable greater transparency and debuggability [65, 135, 159]. LLM chaining is a nascent subfield, with seminal papers emerging in only the past few years [73, 98, 158, 162]. Initial evidence suggests that chaining can mitigate quality deficits, but these solutions continue to face problems such as hallucination [139], bias [121], incorrect responses [123], and expensiveness [175]. There is a vast design space of possible chaining combinations from the suggested operators of this early work, so realizing the potential of chaining is challenging due to the design space’s size and complexity.

The challenge of decomposing complex work, however, is not new. In the more established field of crowdsourcing, researchers have faced a similar problem setup. Groups of people (*requesters*) outsource tasks to other people (*crowdworkers*) on online platforms. Crowdworkers have their own limitations: crowdworkers may output low-quality responses [7, 19, 25, 93], small changes in instructions affect crowdworker responses [80, 120], and requesters have limited modes of interaction with crowdworkers [137]. To compensate for these limitations, crowdsourcing researchers have developed a variety of structured workflows. Crowdsourcing workflows decompose a task into a series of smaller microtasks for crowdworkers to complete. These workflows account for crowdworker limitations by reducing task complexity and leveraging redundancy, aggregation, and explicit validation [9, 19, 29, 169].

Given the parallels between these fields, *we investigate how crowdsourcing techniques may be adapted to design LLM chains* (Figure 1). Selecting relevant strategies from over a decade of crowdsourcing research is challenging, especially as LLMs and crowdworkers have salient differences [161]. For example, crowdworkers may be slow to generate large text amounts, whereas LLMs may lack common sense and factual grounding, making them more sensitive to instructions and likely to “hallucinate” information.

To guide LLM chain development, *we first construct a design space based on a systematic review of the crowdsourcing workflow and LLM chaining literatures*. We analyze 107 papers and perform open coding to identify core design space dimensions. Our design space is organized into the *objectives*

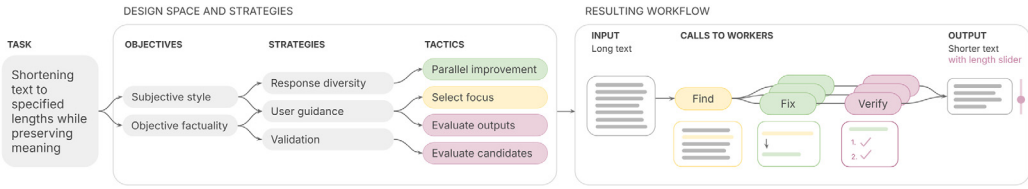


Fig. 1. We contribute (1) a design space, (2) case studies, and (3) a discussion of techniques for LLM chains informed by crowdsourcing workflows. This scaffolding can help designers navigate the large possible space of LLM chains. For example, (Left) given a task of shortening text, as in Soylent [19], our design space aids an LLM chain designer in identifying relevant high-level *objectives*. These objectives incorporate elements of *subjective* and *objective* quality, i.e., creatively shortening input text while verifying its faithfulness to the original. To support these *objectives*, the designer can lean upon concrete *strategies*, such as *validation* and *user guidance*. Strategies are high-level design patterns encapsulating how prior work has manipulated workflow design to achieve *objectives*. These strategies in turn point to lower-level *tactics* that define the chain’s construction, such as a *subtask* to *evaluate* text edit candidates. (Right) The designer can produce LLM chains to support the *objectives* by implementing *strategies* using *tactics*, as outlined in our design space. For example, promoting a high quality *subjective* style can be supported by the *user guidance* strategy, implemented by giving users control to *evaluate* and select the content of the shortened text via a directly manipulable slider.

for which a designer can optimize a workflow (Figure 2) and the *tactics* used to build the workflow (Figure 3).

From this review, we also surface the strategies used to marshal and combine tactics to achieve objectives. For example, tasks in which the ideal output is *subjective* can be supported by strategies that produce *diverse responses*, allow *user guidance*, and utilize *adaptable architectures*. We gather from the literature how tactics have been used in prior work to support strategies. For example, workflows and chains encourage *diverse responses* with *redundant* or *communicative* tactical architectures. Strategies constrain the space and evaluation of tactic choices for both crowds and LLMs (Figure 5).

To better assess the impact of similarities and differences between crowdworkers and LLMs on the implementation of strategies, we implement three case studies adapting crowdsourcing workflows to use LLMs. We chose workflows that enable direct manipulation of the output, an instance of the *user guidance* strategy for quality control that has been explored in crowdsourcing, but is under-explored in chaining. The case studies implement tasks that vary in their degrees of subjectivity, adapting Cascade [32] to control the precision of a generated taxonomy, Soylent [19] to control the length of shortened text, and Mechanical Novel [77] to control the narrative elements of a fictional story.

We evaluate these chains in terms of the quality and controllability of their outputs. For example, compared to zero-shot prompting, chain-generated stories are preferred by crowdsourced raters and the story elements can be more precisely controlled. These case study implementations provide insights into how differences between LLMs and crowdworkers impact chain design to achieve the same objectives. In Mechanical Novel, for instance, we find that the LLM chain requires more explicit support than the crowdsourcing workflow to create creative and varied outputs.

Informed by both our design space and the three case studies, we discuss how differences between LLMs and crowdworkers impact effective tactic designs for LLM chains. For example, because LLM stochasticity differs from human diversity, tactic choices for LLMs must induce *response diversity* more explicitly and account for limitations in LLM *self-validation*. We find that our strategies help focus workflow adaptation and suggest areas of promising future work. Such areas include how the

speed and reduced cost of LLMs relative to crowdworkers presents an opportunity to create higher quality chains, and we recommend improved tools and methods to augment this design process.

## 2 Our Approach within the Crowdsourcing-LLM Intersection

In this article, we integrate a broad survey of literature on LLM chaining and crowdsourcing workflows into a design space. We then use the case studies to explore open questions in this design space. Our broad investigation and resulting design space inform a comprehensive set of recommendations for navigating design decisions and orienting future work. These recommendations span from higher-level guidance, such as adhering to the tradeoff between outcome quality and resource constraints, to lower-level suggestions, such as pursuing architecture-based methods for quality validation. These recommendations expand upon related work by several other groups investigating the intersection of crowdsourcing and LLMs. One such paper explores the impact on requester and crowdworker stakeholders of incorporating LLMs into crowdsourcing workflows [6]. A second considers hybrid crowd-LLM data annotation [99]. Another implements three strategies from crowdsourcing using LLMs for data processing tasks [129]. Closest to our work, the final paper has students implement six crowdsourcing workflows. The article discusses the effectiveness of these workflows and the impact of the difference between LLM and crowdworker abilities [161]. We present a wider scope of recommendations informed by the design space and our case studies.

## 3 Design Space Methodology

We survey papers that involve the design and implementation of crowdsourcing workflows or LLM chains. Although crowdsourcing workflows and LLM chains have many overlapping concepts, the terms used to denote these concepts differ. In this article, we refer to both LLMs and crowdworkers as “workers.” We use the term “workflow” to refer to both crowdsourcing workflows and LLM chains. We refer to intermediate steps of workflows as “subtasks.” We use the term “user” to describe the person receiving workflow outputs (e.g., the “requester” in crowdsourcing terms). The term “designer” in this article refers to a person who defines a workflow to later be used by a “user.” Throughout the discussion of the design space and strategies, we indicate which references are from the crowdsourcing workflow literature [C] and the LLM chaining literature [L].

### 3.1 Paper Collection

We run two paper collection processes, one for the crowdsourcing literature and another for the LLM chaining literature. We use snowball sampling from the relevant citations in and of these papers, including every citation of the initial AI Chains paper [160]. We consider papers available as of August 2023.

*Workflow Literature.* We collect papers through a search of the keywords: *crowd(sourcing) workflow*, *crowd(sourcing) pipeline*, *crowdsourcing complex work*, and *microtask*. We search the ACM CHI Conference on Human Factors in Computing Systems, ACM Conference on Computer-supported Cooperative Work and Social Computing, the ACM Symposium on User Interface Software and Technology, and the AAAI Conference on Human Computation and Crowdsourcing, as well as Semantic Scholar [49] and Google Scholar.

*Chaining Literature.* We collect papers through a search of the keywords: *LLM chain(ing)*, *prompt chain(ing)*, *LLM pipeline*, and *AI chain(s/ing)*. We search the International Conference on Learning Representations, the Conference on Neural Information Processing Systems, the International Conference on Machine Learning, the Association for Computational Linguistics proceedings, the

Association for Computing Machinery proceedings, as well as Semantic Scholar [49] and Google Scholar. We include preprints in our literature search.

*Paper Inclusion Criteria.* We include papers that implement, or support the implementation of, a workflow with at least two steps; i.e., workflows in which the outputs from one worker inform the input to another worker. We include papers that focus on building an interface or domain specific language for other designers to build workflows. We restrict our search to academic papers, rather than software like LangChain [91]. We exclude works that call multiple workers in parallel but never in a chain of steps (e.g., majority voting of labels). This rule limits the crowdsourcing techniques considered to those that reflect the serial and dependent structure of an LLM chain. We also exclude techniques such as **Chain of Thought (CoT)** [155] which issue multiple tasks to a worker in one step instead of incorporating multiple worker calls [30]. We include instances in which workers synchronously work on the same interface, viewing and editing each other's outputs. Finally, the task at hand must have a user working toward some end goal, so we do not consider open-ended collaborative projects like Wikipedia. We make one exception for a paper that explicitly draws from crowdsourcing strategies to complete data aggregation tasks with LLMs [129]. Although they do not implement chains, the connection to crowdsourcing resonates with our design space goals.

*Results.* Our initial keyword search, conducted in August 2023, found 313 papers for workflows and 180 papers for chaining. One author determined whether papers were out of scope given our collection criteria, and any ambiguous cases were discussed among multiple authors. Our final set of papers resulted in 107 papers: 68 for workflows, and 39 for chains.

### 3.2 Coding Strategy

We conduct a thematic analysis [23], following the methodologies of similar design spaces in **Human-Computer Interaction (HCI)** [16, 88, 142]. First, one author extracted (1) the outcomes (*objectives*) and (2) the elements (*tactics*) of the workflows from the paper set. We then employed iterative open coding on the extracted data, in which two authors separately coded each data extraction, sorted resulting codes into sub-categories, and then discussed these categorizations. In three further rounds of iteration, the same authors refined these categorizations. All authors reviewed the resulting design space.

## 4 Design Space

Our design space has two parts: the *objectives* span the space of a designer's goals, and the *tactics* span the space of workflow implementations. These *objectives* are (1) promoting *outcome qualities* while respecting *resource constraints*, (2) specifying an *objective* or *subjective* definition of quality, and (3) defining the *generality* or *specificity* of the set of tasks that the designer wants the workflow to support. For the workflow *tactics*, *actors* complete various *subtasks* that are linked together into the workflow's *architecture*.

We construct the design space from an exhausting survey of existing work, although the conceptual completeness of the design space is limited to this corpus. Defining this space enables understanding, classification, and comparison of existing workflows (Figure 4 and Section 4.3), generation of new workflows (Section 6), and the discovery of design patterns (Section 5) that inform chaining practices and guide directions for future work (Section 8).

### 4.1 Objectives in the Design Space

The *objectives* of a workflow design define the designer's goals within several dimensions. Designers will design for some point along an *outcome quality* vs. *resource constraints* tradeoff, in a range

## Objectives

What the designer wants the workflow to accomplish

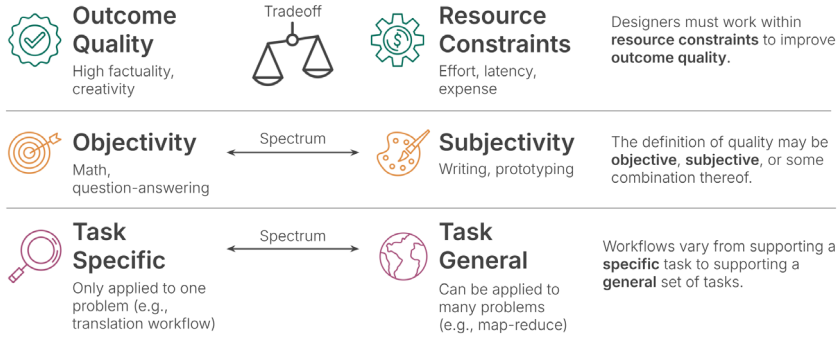


Fig. 2. Objectives. Our design space contains three axes of workflow objectives, shown as rows above. The first regards the impact of *resource constraints* on the ability to improve *outcome quality*. The second axis distinguishes different notions of ‘quality:’ whether the task objective requires *objectivity*, *subjectivity*, or elements of both. The final axis concerns generality: if the workflow is aimed to solve a single, *specific* task or to *generalize* to many tasks.

from *objectivity* to *subjectivity*, and in a range spanning between *task specificity* and *generality*. (Figure 2).

**4.1.1 Outcome Quality vs. Resource Constraints.** When working toward *outcome qualities*, designers must be mindful of *resource constraints*. Workflows can support quality outputs such as by improving accuracy and/or spurring creativity [C[93] L[36, 78, 135, 162]], transparency [L[36, 37, 78, 135, 158, 161]], debuggability [L[159, 162]], and modularity [C[94, 95, 136] L[65, 158, 168, 175]].

Workflows can also address, although imperfectly, quality concerns. Some quality concerns are induced by the workers, such as hallucinations [L[36, 121, 123, 124, 129, 158]] and generative loops [L[121]] in the case of LLMs, and fatigue [C[7, 38]] in the case of crowdworkers. Others are induced by the workflow itself, such as cascading errors [C[53, 82] L[159, 162, 164]] and a lack of coherency [C[19, 60, 77] L[121, 159, 166]]. More concerns are induced by the broader context, such as the need to scale [C[86, 87] L[15]], the workflow’s placement within a larger system [C[29, 34, 131]], unethical use of the workflow [L[130]], and the impact on human labor [L[121, 158]], sense of ownership [L[121]] and sense of agency [L[78, 121]]. However, mechanisms for improving output quality can be more expensive [C[70] L[168]], take longer [C[27, 54]], and require more effort to create [C[82, 94]].

**Expense.** The expense of a workflow refers to paying crowdworkers or paying for an API call to an LLM. Workflows are often more expensive than simpler approaches (e.g., CoT prompting) [C[39, 60, 70, 102, 117, 118] L[43, 51, 123, 153, 168, 175]]. However, the cost is often less than hiring professionals [C[9, 34, 101, 169] L[48]], pretraining models [L[35, 65, 78, 96]], or completing the task yourself [L[154, 174]].

**Latency.** Another constraint is the time for the workflow to complete [C[87, 92, 117, 118]]. Many workflows have lower latency than a single person or a baseline performing the same task, as work can be parallelized across multiple workers [C[11, 31, 54, 56, 87, 92, 95, 102, 133, 147] L[78]].

**Effort.** Workflows can help users complete tasks with less effort [C[87] L[110, 121, 158]]. However, workflows can require high effort to create [C[60, 80, 86]] and oversee [C[94] L[159]].

**4.1.2 Objectivity vs Subjectivity.** One element of output quality of particular importance is if the workflow should support *objective* or *subjective* outputs. Many tasks' definitions of output quality have elements of both objectivity and subjectivity.

**Objectivity.** Objectivity-focused tasks have one correct answer. Some crowdsourcing papers in this space include TurKit [C[105]] and Second Opinion [C[122]]. Examples of these tasks include math problems [L[43, 100, 114, 163, 176]], blurry text recognition [C[104, 105]], and matching images with the same content [C[83, 122]]. For these tasks, workflows can improve performance on benchmarks [C[58, 83, 120] L[35, 47, 97, 139, 163, 174]], enable new tasks [C[54, 57, 70, 105] L[66]], classify examples for which machine learning approaches struggle [C[7, 19, 29, 122]], and reduce hallucinations [L[36, 37, 121, 123, 124, 129, 158]].

**Subjectivity.** Subjective tasks have the widest degree of output freedom. Representative crowdsourcing papers with subjective outputs include Flash Teams [C[136]], VisiBlends [C[33]], and Mechanical Novel [C[77]]. These papers tackle creative writing [C[10, 19, 77] L[121, 166, 168, 173]], idea generation [C[105] L[78, 161]], and prototyping [C[87, 95, 136, 150] L[130]]. Prior work finds that on subjective tasks, workflows can help achieve higher quality [C[10, 77, 95] L[78, 121, 130, 166, 168]] and less constrained outputs [C[42, 60, 95, 107] L[130, 154]]. Workflows may enable novice crowdworkers [C[133]] or users [C[11, 33, 75, 102, 150] L[154, 161]] to perform a subjective task with an otherwise high barrier of entry. Important qualities for subjective tasks include coherence [C[77] L[166, 168]] and aligning to the user's creative vision [C[137]]. On subjective tasks, a workflow is often better suited for creating and lengthening the output, rather than proofreading [C[69, 77, 105] L[121]], so these outputs are sometimes seen as a starting point for editing, rather than a finished product [C[75, 126] L[110, 121, 130]].

**Both Objectivity and Subjectivity.** Many tasks require some degree of objectivity and some degree of subjectivity. In some tasks, there is not one correct answer, but there are restrictions on what outputs are correct. Crowdsourcing papers for these tasks include Cascade [C[32]], Flock [C[29]], Revolt [C[27]], and Storia [C[75]]. Example tasks include categorizing data (categories must be accurate, but there are multiple options and unclear edge cases) [C[11, 22, 27, 29, 31, 32]], summarizing text (summaries must stay faithful, but have df in how to write the summary) [C[19, 75, 131] L[123, 124]], and coding (programs must work, but there are multiple approaches to problem solving) [C[87, 94] L[64, 65, 158]]. The goals of workflows for these tasks are to align with expert outputs [C[32, 76, 107, 137] L[48, 77, 129]] or to enable users to choose among multiple correct possible output options. Using workflows for these tasks can result in an output that is human-understandable [C[29]], that better represents what would be important for people to know [C[138]], and that surfaces the uncertainty among the different possible correct options [C[27]].

**4.1.3 Task Specificity vs Generality.** Some workflows are highly *specialized* to the specifications of the task (e.g., visual blends [C[33] L[154]], masking faces in images [C[7, 74]], and programming [C[94] L[65]]). Others *generalize* to a variety of tasks [C[86] L[36, 37, 124, 158]]. For example, AutoGen [L[158]] uses the same conversation-based workflow for coding as well as chess games, math problems, and retrieval-based question answering. Task-specific workflows support fewer outcome types but can incorporate more tailored support than workflows that generalize.

## 4.2 Tactics in the Design Space

Designers use workflow *tactics* to implement workflows. Tactics need to consider the *actors*, the *subtasks* the actors complete, and the connections of subtasks into workflow *architectures* (Figure 3).

# Tactics

The building blocks of workflows

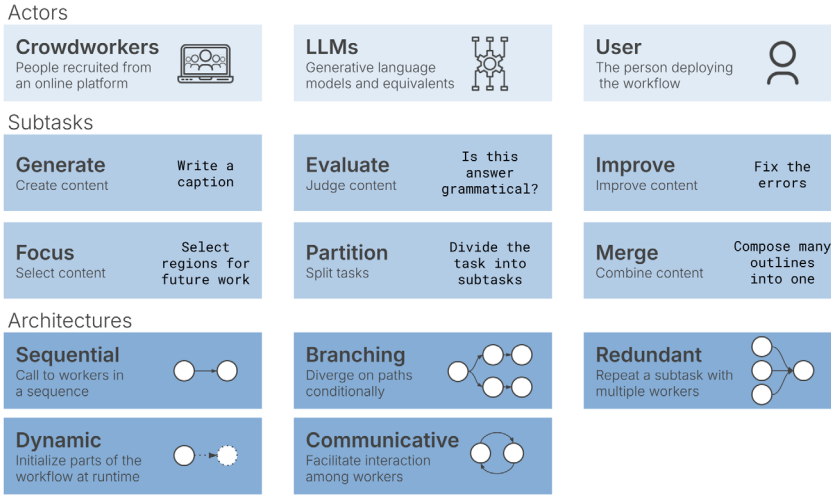


Fig. 3. Tactics. Tactics are the building blocks of workflows. *Actors* complete *subtasks*, which are connected together in *architectures*. There are multiple categories of actors (crowdworkers, LLMs, and the user), subtasks (generate, evaluate, improve, focus, partition and merge), and architectures (sequential, branching, redundant, dynamic, and communicative). Workflows may use multiple tactics in any of these categories.

## 4.2.1 Actors. Actors may be crowdworkers, LLMs, or the user.

**Crowdworkers.** These actors are usually recruited from crowdworking platforms like Amazon Mechanical Turk, or UpWork (formerly known as oDesk). Platforms allow for requests of crowdworker qualifications, although these credentials are sometimes hard to validate [C[4, 87, 150, 169]]. Methods like retainer models can reduce the latency of accessing crowdworkers on these platforms [C[18]]. Workers can also be recruited from the public, including community members [C[115]], conference committees [C[11, 31]], and online deaf communities [C[20]].

**LLMs.** There are LLMs of many sizes and expenses [L[24]]. Models have varying strengths and respond differently to prompts [L[100, 121, 129, 161]]. LLM outputs can be adjusted through parameter changes (e.g., temperature) [L[129, 130]], fine-tuning [L[24, 35–37, 97]], and expressing roles through prompts [L[43, 130]]. Recently, designers often substitute specialized or distilled models, such as embeddings, to achieve LLM functionality at reduced cost [L[97, 121]].

**Comparing Crowdworkers and LLMs.** Some limitations of crowdworkers and LLMs overlap. First, both crowdworkers and LLMs can output low-quality responses. Prior work finds that outcomes from crowdsourcing platforms include low-quality, spam, and off-topic responses [C[7, 19, 25, 67, 93]]. LLMs can provide erroneous outputs: they are known to confidently output incorrect answers [L[43, 129, 154]] and to struggle with self-consistency and logical reasoning [L[36, 37, 100, 129, 164, 168]]. Second, long and complex instructions can lower performance for both crowdworkers and LLMs. Crowdworkers struggle with following long and complex instructions, as people can overly focus on the beginning or end of a given text [C[28, 32, 60, 92]]. LLMs can ignore large fractions of long prompts or focus mostly on recent parts of the given context [L[43, 114, 129, 161]]. Third, both crowdworkers and LLMs can give biased outputs. Crowdworkers can be biased; for example, in one study, crowdworkers performed significantly better in identifying photos of their

own race compared to other races [C[122]]. LLMs also output biased results, a widespread concern [L[64, 110, 130, 158]]. For example, work on screenplay generation found stereotypical outputs, with multiple instances of gender biases [L[121]].

Crowdworkers have some unique limitations. Since crowdworkers are paid when they complete a task, there is an incentive to work quickly even if that means sacrificing [C[19] L[129]]. Additionally, since rejected work can hurt access to future jobs on a platform [C[77]], crowdworkers may overcompensate to a detrimental level in an attempt to make their efforts clear [C[19, 76]]. Crowdworkers may fatigue [C[7, 38]], and they may lack task-relevant knowledge and delete jargon [C[19, 126]]. On synchronous tasks, other limitations become prevalent. Social blocking can make crowdworkers so cautious to change others' work that not enough work gets done [C[60, 92]], and territoriality can motivate edit wars in which crowdworkers' attempts to establish their own work over others degrade the final product quality [C[10, 150]]. Crowdworkers also may be affected by diffusion of responsibility and social loafing in which there is incentive to wait for others to do tasks [C[10, 92]].

LLMs also have unique limitations. On creative tasks, LLMs sometimes produce outputs that are perceived as robotic [L[110]] and that can lack the understanding, nuance, and subtext that people can provide [L[121]]. LLM training on human data increases the risk of accidental plagiarism [L[121]]. Furthermore, models can get stuck in generative loops, degenerating text quality [L[129]]. LLMs also hallucinate more than crowds [L[36, 37, 121, 123, 124, 129, 158]] and may struggle to change outputs through self-reflection [L[100]].

Crowdworkers and LLMs have their comparative strengths as well. Crowdworkers excel at performing high-level structural evaluation and quickly identifying outliers [C[68, 111]]. People are also currently better at selectively choosing among subtasks than LLMs [L[161]]. LLMs may be better at divergent thinking [L[100, 158, 161]] and comparison-based instructions [L[158, 161]] than crowdworkers, who are affected more by anchoring bias. LLMs may also be better than crowdworkers at tasks requiring recollection of large amounts of data [L[129]].

*User.* Finally, users of the workflow can also be involved in revising or guiding the process. Users can be novices using the workflow to adapt to a steep learning curve [C[33, 87, 102, 133] L[154]], or experts using the workflow as a support tool [C[122, 126] L[121]]. However, even with user involvement, prior work finds mixed results on how well tools retain the user's sense of agency and ownership [L[78, 121]].

**4.2.2 Subtasks.** Subtasks are singular tasks in a workflow. Although highly varied, they follow several thematic functions: to *generate*, *evaluate*, *improve*, *focus*, *partition*, and *merge*.

*Generate.* Most workflows involve at least one step in which actors generate content. This generation may be writing words or code [C[131, 156] L[97, 173, 174]] brainstorming and ideating [C[33, 67, 93, 115] L[110, 168]], or answering a question or solving a problem [C[83] L[47, 163, 164]]. Generation tasks may also ask for labels or annotations on given content [C[22, 27, 31, 57, 58]], to translate between languages [C[9, 169]], to contribute to a conversation with other actors [C[115] L[43, 100, 153]], to demonstrate a behavior [C[20, 95]], or to source content from external sources [C[33, 34, 60, 68, 75, 76, 111, 133]]. Some subtasks also require an explanation along with the generated content [C[27, 29, 42, 77] L[48, 51, 139]].

*Evaluate.* One of the most common evaluation tasks is choosing the best among multiple possible options [C[29, 57, 58, 80, 107, 156] L[121, 166, 168]]. Other subtasks may include rating and ranking among multiple options [C[70, 115, 117, 131, 137, 157] L[97, 129, 168]] or a pairwise comparison [C[104, 105, 117, 118, 122] L[24, 129]]. Evaluation subtasks may also determine if another round of generation needs to occur [C[94] L[100, 153, 162]].

Workflows can decide what moves forward through majority vote [C[39, 58, 69, 76, 92, 107]], other aggregations of voting tasks (e.g., expectation maximization [C[42] L[129]] or weighted mean [C[122]]), returning the highest ranked or rated items [C[70, 131]], or automated methods. These automated methods include measuring agreement from another step [C[18, 27, 60, 68, 70, 107] L[124]], weak supervision [L[15]], and using non-AI checks like semantic embeddings and rule-based heuristics [C[9] L[129, 154]]. There are also subtasks that seek critiques and reviews from the user [C[126] L[139]] or a worker [C[34, 77, 87, 115, 150] L[100, 114, 129, 139]]. These subtasks may ask for feedback or verify if errors or other qualities exist [C[11, 19, 22, 25, 33, 34, 68, 69, 76, 82, 86, 156] L[51, 65, 100, 161, 166]].

*Improve.* Other subtasks improve content generated by previous steps, sometimes by incorporating feedback from evaluation steps. Improving on previously generated content can come in many forms: updating the format of the content [L[15, 47, 65, 153, 162]], pruning bad results [C[107] L[51]], adding missing results [C[107, 156] L[65]], correcting errors [C[42, 54, 83] L[65, 166]], and performing direct improvements [C[19, 34, 39, 58, 60, 69, 77, 95, 101, 105, 115, 120, 126] L[114, 139]].

*Focus.* One set of subtasks looks to select the attention of the upcoming tasks within a broader context. For example, this step may find patches of the input that need editing, extract information from context, or focus on common keywords [C[18, 70, 77] L[36, 37, 51, 124]].

*Partition.* Some subtasks transform a task by breaking it up into smaller pieces [C[82, 86, 94, 150] L[66, 161, 162, 176]]. For example, the Divide step in the Turkomatic workflow asks crowdworkers to divide a given subtask into simpler subtasks. The workflow calls this step recursively until the subtask complexity is valued (by a separate crowdworker) at a specified cost [C[86]].

*Merge.* Subtasks can merge multiple streams of incoming content. This merge could be algorithmic [C[4, 7, 8, 33, 82, 147]] or accomplished by an actor [C[9, 80, 82, 86, 94, 95, 111, 133, 147] L[161, 162]].

**4.2.3 Workflow Architecture.** Workflows combine subtasks into larger architectures that may be *sequential*, *branching*, *redundant*, *dynamic*, or *communicative*. One workflow may employ multiple architectural patterns.

*Sequential* architectures order subtasks in a chain in which the output of one actor passes forward to another actor [C[34, 69, 70, 120] L[47, 51, 64, 65, 110, 123, 135, 154, 161, 162, 176]].

*Branching* workflows apply conditionals to choose different execution paths based on the outputs of prior steps [C[67, 80] L[65, 159, 175]].

*Redundant* architectures require different actors to complete the same task in parallel or iteratively. In parallel generation, multiple actors perform the same subtask separately [C[8, 25, 58, 69, 102, 156] L[35]]. This subtask is often, but not always, followed by an aggregation or evaluation subtask to merge or select among the different parallel generations. Unlike parallelization that runs a similar subtask on different items at the same time (e.g., five crowdworkers label one image each for five labeled images total), parallel redundant architectures run the same subtask on the same item (e.g., five crowdworkers independently label the same image for one image total). Iterative architectures run iteratively over a changing input. For example, iterative improvement takes the output of a subtask and iteratively gives it to a series of actors to improve, sometimes with interspersed decision tasks choosing between the current and prior iterations [C[10, 38, 53, 92, 94, 111, 157]]. The number of iterations can be determined to be a fixed number of steps [C[70, 77, 105, 169] L[161]], by including a subtask to check for sufficiency [C[56, 87, 137] L[36, 37, 100, 114, 126, 153, 163]], until the user stops calling iterations [C[29, 33, 87, 157] L[135]], until the outputs converge [C[101]

L[43, 51, 139]], until all quality checks are passed [C[22, 32, 77, 117, 118, 147] L[166]], through increasing complexity levels [C[7, 8]], and through decision-theoretic approaches [C[39] L[130]].

*Dynamic* architectures support changing the workflow architecture as it is executed. Map-reduce is a dynamic architecture in which actors break down a task recursively until the subtasks are an appropriate size [C[74, 75, 82, 86, 101, 150] L[66, 161]]. The literature defines such workflows as a hierarchy [C[82]]. Other dynamic workflow architectures include doing beam search across multiple options [L[36, 164, 168]] and allowing for actors to flexibly create tasks on a shared to-do list before choosing which of those to complete [C[33, 82, 94, 115, 126, 171]]. Some other architecture methods may have dynamic elements, such as parallelizing an input into chunks and running the workflow for each chunk [C[11, 54, 74, 76, 131, 169]] or dynamically determining the number of iterative steps [C[28, 31] L[114, 124, 176]].

*Communicative* architectures can facilitate debate and conversation among actors [C[27, 28, 42, 83, 138, 147] L[43, 100, 158]] and allow workers to return feedback to the user [C[86, 92, 126]]. Some workflows create roles for different actors, often based on roles held in real-life organizations (e.g., journalism teams, software development teams) [C[3, 136, 150] L[130]]. Some papers integrate these roles into a hierarchy of communication, with managers, midlevel managers, and different forms of communication between different levels of the hierarchy [C[3, 136]]. Communication may be restrained through subtasks [C[27, 28]], free-form through e-mail [C[3, 150]], or enabled through a shared interface [C[87, 92, 93, 95, 115, 126]]. Communication among actors may be synchronous or asynchronous. Synchronous workflows help enable collaboration [C[10, 31, 33, 138]] and can reduce latency [C[18]]. In asynchronous workflows or workflows with limited communication architectures, a working memory can help actors access progress reports, what other actors think are good next steps, and the overarching goals [C[60, 93, 94, 150, 171] L[124, 158, 166]].

### 4.3 Comparing the Utilization of Techniques

This section discusses the distribution of the crowdsourcing workflow and LLM chaining literature across the design space. We fit each paper in the corpus into its design space categories. We detail the labeling procedure in Appendix A and the outcomes in Tables Table 1 and 2 as well as Supplementary materials. The labeling procedure and outcomes are detailed in Supplementary materials and in Tables 1 and 2. Here we provide a snapshot of the LLM chaining sub-field as it expands, but these distributions will change over time with more published papers.

Comparing these distributions first shows how both literatures cover all design space categories, with the natural exception of the *LLM* and *crowdworker* categories (Figure 4). The methods for creating the design space did not enforce overlap of the literatures, so the high degree of overlap between them provides a sanity check that both fields can share a design space.

However, their distribution differs across some categories. For instance, a greater proportion of crowdsourcing papers promote the decrease of *cost* compared to experts, while LLM chains' costs are usually compared to a baseline of one call to the LLM, inherently increasing the *costs*. Due to the high range of tasks and sparsity of reporting, we cannot compare the scale of the relative costs of workflows and chains. Furthermore, differences in execution details hinder comparison across the literatures even for similar tasks.

Across the objectivity-subjectivity spectrum, we find that crowdsourcing workflows more often implement a task with *subjectivity*, while LLM chains more often have purely *objective* tasks such as math. Finally, more chains in the LLM literature are designed to be *generalizable* to many tasks, while a greater proportion of workflows in the crowdsourcing literature are designed for a *specific* task function.

Table 1. Crowdsourcing Workflow Corpus Categorization

Reference	Title	Venue	Year	Quality Increase	Quality Decrease	Cost Increase	Cost Decrease	Objectivity	Subjectivity	Task General	Task Specific	Sequential	Branching	Redundant	Dynamic	Communicative	Generate	Evaluate	Improve	Focus	Partition	Merge	Crowdworkers	LLMs	User
[19]	Soylent: A Word Processor with...	UIST	2010																						
[39]	Decision-Theoretic Control of...	AAAI	2010																						
[104]	Exploring Iterative and Parallel...	CKDD-HCOMP	2010																						
[105]	TurKit: Human Computation...	UIST	2010																						
[82]	CrowdForge: Crowdsourcing Complex...	UIST	2011																						
[4]	The Jabberwocky Programming...	UIST	2011																						
[118]	Demonstration of Quirk: A Query...	SIGMOD	2011																						
[117]	Human-Powered Sorts and...	VLDB	2011																						
[18]	Crowds in Two Seconds: Enabling...	UIST	2011																						
[101]	An Iterative Dual Pathway...	AAAI	2011																						
[169]	Crowdsourcing Translation...	ACL	2011																						
[9]	Collaborative Workflow for...	CSCW	2012																						
[86]	Collaboratively Crowdsourcing...	CSCW	2012																						
[80]	CrowdWeaver: Visually Managing...	CSCW	2012																						
[157]	Strategies for Crowdsourcing...	CHI	2012																						
[171]	Human Computation Tasks with...	CHI	2012																						
[11]	Community Clustering: Leveraging...	HCOMP	2013																						
[93]	Chorus: A Crowd-Powered...	UIST	2013																						
[133]	CrowdBand: An Automated...	HCOMP	2013																						
[22]	Crowdsourcing Multi-Label...	HCOMP	2013																						
[32]	Cascade: Crowdsourcing Taxonomy...	CHI	2013																						
[38]	POMDP-Based Control of Workflows...	Elsiver	2013																						
[76]	Crowdsourcing Step-by-Step...	CHI	2014																						
[102]	Crowdsourced Explanations for...	HCOMP	2014																						
[131]	Video Digests: A Browseable...	UIST	2014																						
[136]	Expert Crowdsourcing with Flash...	UIST	2014																						
[34]	STEP: A Scalable Testing and...	HCOMP	2014																						
[10]	Effects of Simultaneous and...	CHI	2014																						
[87]	Wish: Amplifying Creative Ability...	HCOMP	2014																						
[94]	Microtask Programming: Building...	UIST	2014																						
[31]	Frenzy: Collaborative Data...	CHI	2014																						
[156]	Learnersourcing Subgoal Labels...	CSCW	2015																						
[3]	Crowdsourcing in the Field: A...	HCOMP	2015																						
[68]	Guardian: A Crowd-Powered Spoken...	HCOMP	2015																						
[29]	Flock: Hybrid Crowd-Machine...	CSCW	2015																						
[92]	Apparition: Crowdsourced User...	CHI	2015																						
[111]	Crowdlines: Supporting Synthesis...	HCOMP	2015																						
[42]	MicroTalk: Using Argumentation...	HCOMP	2016																						
[60]	The Knowledge Accelerator: Big...	CHI	2016																						
[75]	Storia: Summarizing Social Media...	CSCW	2016																						
[120]	Why Is That Relevant?...	HCOMP	2016																						
[126]	WearWrite: Crowd-Assisted Writing...	CHI	2016																						
[53]	Understanding Crowdsourcing...	HCOMP	2016																						
[56]	VizLens: A Robust and Interactive...	UIST	2016																						
[147]	Supporting Collaborative Writing...	CHI	2016																						
[25]	More Than a Feeling: The MiFace...	UIST	2017																						
[27]	Revolt: Collaborative Crowdsourcing...	CHI	2017																						
[57]	Facade: Auto-Generating Tactile...	CHI	2017																						
[69]	Supporting ESL Writing by...	HCOMP	2017																						
[95]	SketchExpress: Remixing Animations...	UIST	2017																						
[138]	Toward Scalable Social Alt Text...	HCOMP	2017																						
[150]	Flash Organizations: Crowdsourcing...	CHI	2017																						
[74]	CrowdMask: Using Crowds to...	HCOMP	2017																						
[77]	Mechanical Novel: Crowdsourcing...	CSCW	2017																						
[137]	Communicating Context to the...	CSCW	2017																						
[107]	ConceptScape: Collaborative...	CHI	2018																						
[67]	Evorus: A Crowd-Powered...	CHI	2018																						
[54]	EURECA: Enhanced Understanding...	HCOMP	2018																						
[115]	CommunityCrit: Inviting the...	CHI	2018																						
[83]	An Empirical Study on Short...	HCOMP	2018																						
[58]	StateLens: A Reverse Engineering...	UIST	2019																						
[122]	Second Opinion: Supporting...	HCOMP	2019																						
[33]	VisiBlends: A Flexible Workflow...	CHI	2019																						
[28]	Cicero: Multi-Turn, Contextual...	CHI	2019																						
[7]	Privacy-Preserving Face Redaction...	HCOMP	2020																						
[70]	Crowdsourced Detection of...	CHI	2020																						
[20]	ASL Sea Battle: Gamifying Sign...	CHI	2021																						
[8]	Pterodactyl: Two-Step Redaction...	HCOMP	2021																						

Each row references a crowdsourcing workflow paper from our corpus, and each column represents a category of the design space. Colored cells indicate that the paper’s workflow contains the category element. Details on the categorization process can be found in Appendix A.

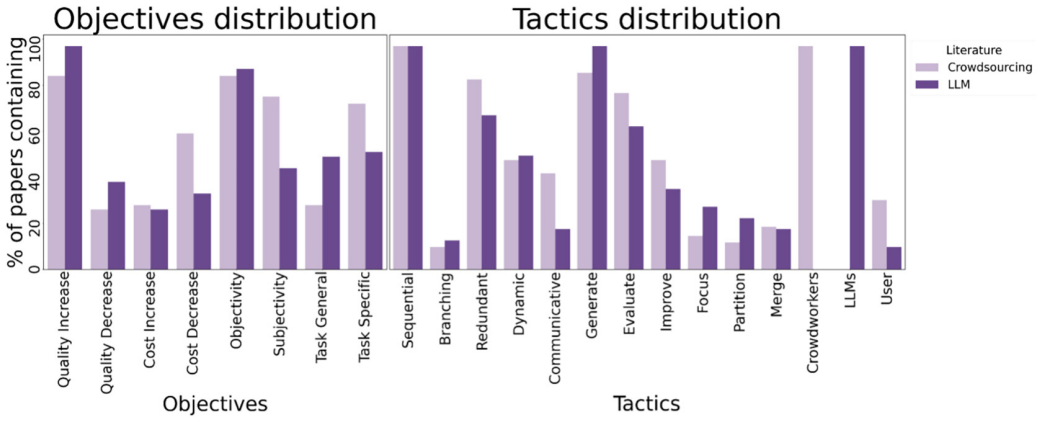


Fig. 4. The distribution of the corpus in the design space. We categorize the papers in our corpus across different categories of the design space. Each set of bars is the categorization, split by the crowdsourcing and LLM literatures. There are fewer LLM papers in our corpus (39) than crowdsourcing papers (68). The length of the bar is the proportion of all papers in each literature that incorporate the design space element. We provide the categorization of each article in the [Supplementary materials](#), and we provide the criteria for our categorizations in Appendix A. Through this comparison, we see how the literatures similarly span the entirety of the design space, with the exception of using LLMs or crowdworkers. We surface some differences as well, such as how crowdsourcing papers are more often *task-specific*, incorporate *subjectivity*, and use *communicative* architectures.

The crowdsourcing workflow literature uses *redundant* and especially *communicative* architectures more frequently. More chaining papers use only a *sequential* architecture than crowdsourcing papers (9 papers vs. 1 paper). There are slight differences in subtask use, in which the crowdsourcing workflow literature uses *evaluate* and *improve* subtasks more often, and the LLM chaining literature uses *generate*, *focus*, and *partition* subtasks more often. Finally, more papers in the crowdsourcing workflow literature explicitly incorporate the *user* than those in the LLM chaining literature.

Surfacing the differences between these distributions focuses attention on potentially informative sections of the design space. These differences could result for a multitude of reasons. We speculate that the lower rate of *subjective* tasks in chaining reflects a relative preference for benchmark evaluations in the LLM chaining community. The incorporation of the *user* and the focus on *task specificity* could reflect that crowdsourcing work was primarily housed in the HCI field. The variations in these distributions could also reflect inherent differences between LLMs and crowdworkers or indicate areas of potential future work.

Comparing the two distributions confirms the many similarities between crowdsourcing and chaining methods while focusing attention on the explanatory and actionable potential of the differences. We explore these differences in further detail in the case studies (Section 7) and the discussion (Section 8).

## 5 Strategies Integrating the Design Space

Choosing *tactic* designs to support *objectives* is a non-trivial task, as combining different tactic types and implementations supports a huge range of possible workflows. However, through the thematic coding process we surfaced strategies by which tactics are used to support objectives (Figure 5). Specifically, these strategies (5.1) encourage diverse responses, (5.2) validate outputs from subtasks, (5.3) enable user guidance, (5.4) utilize adaptable architectures, (5.5) improve actor

Table 2. LLM Chain Corpus Categorization

Reference	Title	Venue	Year	Quality Increase	Quality Decrease	Cost Increase	Cost Decrease	Objectivity	Subjectivity	Task General	Task Specific	Sequential	Branching	Redundant	Dynamic	Communicative	Generate	Evaluate	Improve	Focus	Partition	Merge	Crowdworkers	LLMs	User
[35]	Training Verifiers to Solve Math...	arXiv	2021																						
[159]	PromptChainer: Chaining Large...	CHI	2022																						
[24]	Learning to Perform Complex...	EMNLP	2022																						
[47]	Honest Students from Untrusted...	arXiv	2022																						
[162]	AI Chains: Transparent and...	CHI	2022																						
[174]	Automatic Chain of Thought...	ICLR	2022																						
[36]	Faithful Reasoning Using Large...	arXiv	2022																						
[37]	Selection-Inference: Exploiting...	arXiv	2022																						
[96]	Standing on the Shoulders of...	arXiv	2022																						
[139]	PEER: A Collaborative Language...	arXiv	2022																						
[153]	Iteratively Prompt Pre-Trained...	EMNLP	2022																						
[130]	Social Simulacra: Creating...	UIST	2022																						
[166]	Re3: Generating Longer Stories...	EMNLP	2022																						
[135]	Iterated Decomposition: Improving...	arXiv	2023																						
[15]	Ask Me Anything: A Simple Strategy...	ICLR	2023																						
[66]	AI Chain on Large Language Model...	arXiv	2023																						
[129]	Revisiting Prompt Engineering...	arXiv	2023																						
[154]	PopBlends: Strategies for...	CHI	2023																						
[164]	Decomposition Enhances Reasoning...	arXiv	2023																						
[168]	Tree of Thoughts: Deliberate...	Neurips	2023																						
[175]	DeepPavlov Dream: Platform for...	ACL	2023																						
[48]	Assessing the Potential of...	medRxiv	2023																						
[64]	Adaptive Intellect Unleashed...	CORR	2023																						
[97]	Teach LLMs to Personalize – An...	arXiv	2023																						
[110]	Tweetorial Hooks: Generative AI...	ICCC	2023																						
[123]	Generating Medically-Accurate...	ACL	2023																						
[158]	AutoGen: Enabling Next-Gen LLM...	arXiv	2023																						
[176]	Least-to-Most Prompting Enables...	ICLR	2023																						
[65]	PCR-Chain: Partial Code Reuse...	ICSE	2023																						
[51]	Self-Verification Improves...	arXiv	2023																						
[78]	Metaphorian: Leveraging Large...	DIS	2023																						
[100]	Encouraging Divergent Thinking...	arXiv	2023																						
[114]	Self-Refine: Iterative Refinement...	Neurips	2023																						
[121]	Co-Writing Screenplays and...	CHI	2023																						
[173]	VISAR: A Human-AI Argumentative...	UIST	2023																						
[43]	Improving Factuality and Reasoning...	arXiv	2023																						
[124]	DERA: Enhancing Large Language...	arXiv	2023																						
[161]	LLMs as Workers in Human-Computational...	arXiv	2023																						
[163]	An Empirical Study on Challenging...	ICLR Workshop	2024																						

The color cells indicate when a paper from the LLM chaining corpus (rows) includes an element of the design space (columns).

selection, (5.6) align a subtask’s required skills with an actor’s capabilities, (5.7) promote global context, and (5.8) ensure outputs meet quality thresholds.

We outline the strategies that surfaced from the literature in a manner adapted from *design patterns* [5, 50] used in architecture and software engineering. For each, we define the strategy, explain the effects on objectives that motivate using the strategy, and how to implement the strategy with tactics. We then discuss secondary benefits, tradeoffs, examples from the literature, and connections to other strategies. While we note some differences between the crowdsourcing and chaining literatures throughout this section, we leave the bulk of such discussion to Section 8.

5.1 Diverse Responses

To promote objective or subjective quality, source multiple responses to the same subtask with parallel, iterative, or communicative workflows (Figure 6).

Motivation. Individual actors can give low-quality outputs, affecting the quality of the overall output. Designers may want to improve robustness to low-quality outputs, outliers, and spammers

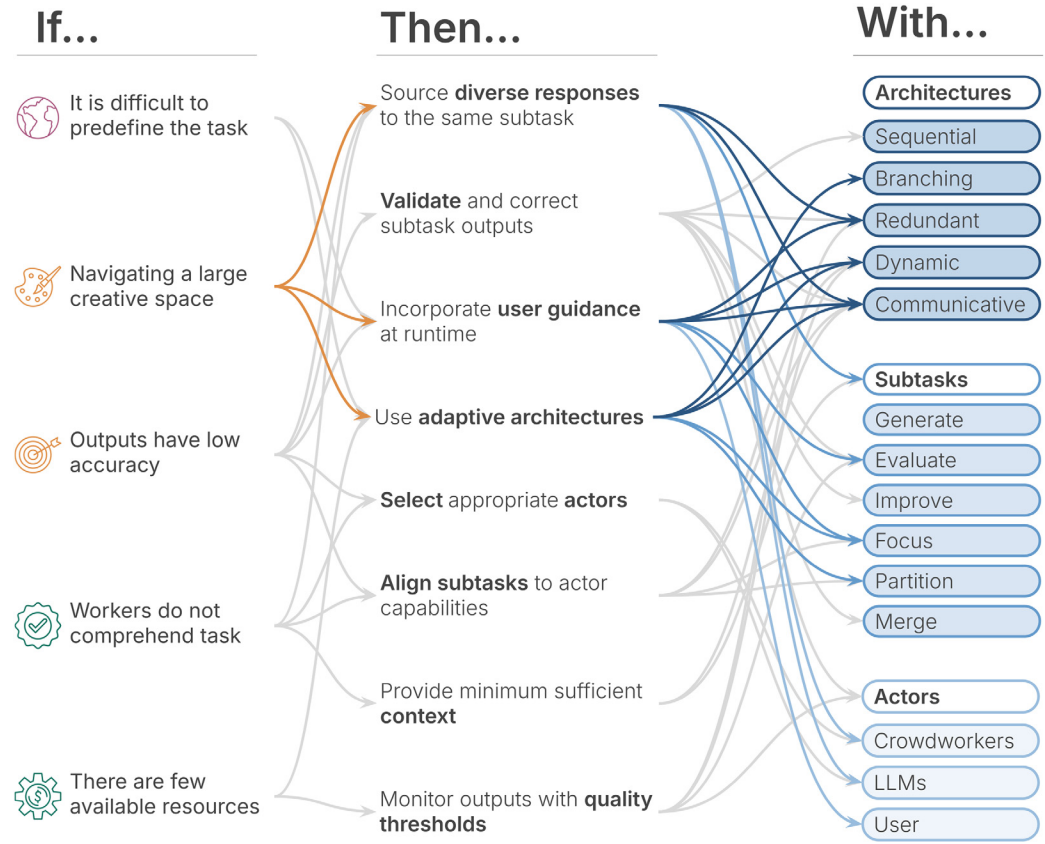


Fig. 5. Design space connections. In working toward objectives, challenges can arise. This figure captures a subset of possible issues in achieving objectives. Based on our design space, these issues may be addressed by utilizing different strategies. In turn, these strategies are supported by various tactics. For example, if the objective requires navigating a large creative space of *subjective* outcomes, a challenging proposition, then designers can incorporate the *adaptive architectures*, *diverse responses*, and *user guidance* strategies. Each of these strategies has been implemented in prior work by a variety of architectures and other tactics.

to support the accuracy of objective outcomes [C[9, 60, 67, 70, 93, 101, 169] L[43, 135]]. For *subjective* tasks, they may want to encourage divergent and diverse thinking [C[69, 92, 111, 147] L[100, 158]].

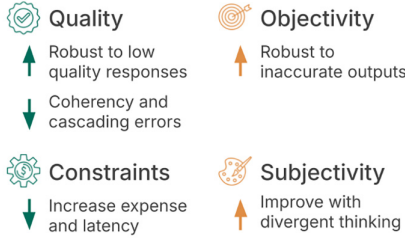
**Implementation.** Getting responses from multiple actors on a subtask can help support these goals. Crowdsourcing workflows and LLM chaining source diverse responses differently. Crowdsourcing does so by querying different people. Chaining leverages the stochastic nature of LLMs (e.g., with a higher temperature parameter setting) or uses different models or prompt variants (e.g., role-based prompts) to source a variety of responses [L[15, 130, 161]].

*Redundant* architectures are the most direct way of sourcing diverse responses, but *communicative* architectures, such as those that structure debates, can source diverse responses as well. Some studies in crowdsourcing systematically study the impact of different architectures (e.g., parallel vs. iterative) and architecture parameters (e.g., the number of iterations) on the diversity and *quality*

## Diverse responses

To promote objective or subjective quality source multiple responses to the same subtask with parallel, iterative, or communicative workflows.

### EFFECTS ON OBJECTIVES



### TACTICS USED

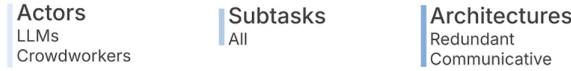


Fig. 6. An overview of the tactics used for and the effects of the diverse responses strategy.

of the outcomes [C[10, 53, 82, 104]]. However, it is unknown if these findings transfer to similar tactic choices with LLMs.

*Secondary Benefits.* There are some benefits of diverse responses that are solely discussed in the crowdsourcing literature, as they are dependent on the variations among people. Crowdsourcing elicits the benefit of having different viewpoints [C[9, 28, 93]], filling in other crowdworkers’ deficits to create a more rounded final product [C[75, 171]], and sourcing perspectives from expert crowdworkers [C[31]] or from crowdworkers with various skill levels [C[115]].

*Consequences.* Sourcing diverse responses is inherently more *costly* because multiple actors complete the same subtask. Merging multiple responses can induce *outcome* quality errors: coherency for parallel architectures [L[159]] and cascading errors for iterative architectures [C[53]].

*Examples.* Examples of the diverse responses strategy include: Apparition [C[92]], chains for multi-agent debate [L[100]], and an exploration of “human computation processes” [C[104]].

*Related Strategies.* Agreement among diverse responses can serve as *validation*. Setting *quality thresholds* can impact the number of diverse responses requested.

## 5.2 Validation

*To improve quality of the overall output, validate and correct subtask outputs with deterministic, architectural, or actor-based checks (Figure 7).*

*Motivation.* Individual actors can give low-quality responses. Unaddressed, these can lead to cascading errors [C[19, 53, 104, 105] L[159]]. Therefore, higher quality subtask outcomes throughout the workflow can have myriad benefits on the overall *outcome quality* for both *objectivity* and *subjectivity* focused tasks, increasing robustness to low-quality actor outputs [C[9, 19]].

*Implementation.* The validation process is often *sequential*, in which content is first *generated* before being validated [L[51]]. In parallel and iterative *redundant* workflows, errors from prior passes can be found and fixed or rejected through *improve* subtasks [C[60, 169]] and differing opinions can be merged [C[82, 111]] through *merge* and *evaluate* subtasks [C[9, 86] L[139, 158, 161, 175]].

## Validation

To improve quality of the overall output, validate and correct subtask outputs with deterministic, architectural, or actor-based checks.



Fig. 7. An overview of the tactics used for and the effects of the validation strategy.

Users can break ties [C[133]] and intervene if there are cascading errors [C[86]]. Non-worker methods such as taking the majority vote of the outputs of redundant *evaluate* subtasks can also be effective. Automated methods of flagging inattentive crowdworkers, gold-standard validation steps, or requiring a certain agreement threshold among actors within a *redundant* architecture can also provide validation [C[70, 92, 101] L[129]].

Debate-based *communicative* workflows can promote reflection on disagreements [C[28, 83] L[43, 100]]. Using different models and roles in debate shows promise [L[43]], as an LLM judge in a debate was found to more likely agree with its same underlying model [L[100]]. Other workflows use *adaptable architectures* to make all production always editable so actors can fix errors in an adaptable manner [C[94, 107]].

*Consequences.* Costs increase when validation involves more calls to actors. The literature shows mixed results for an LLM’s ability to validate its own responses in *evaluate* subtasks. Some studies suggest that adding a self-validation step improves outcomes [L[51, 114, 164, 168]], while others suggest that LLMs are likely to agree with their previous outputs, even when incorrect, and propose debate-based *communicative* workflows as an alternative [L[43, 100]].

*Examples.* Examples of the validation strategy include: the “Verify” step of SoyLent [C[19]], a Dual-Pathways structure for transcription [C[101]], and self-evaluation guided reasoning [L[164]].

*Related Strategies.* Validation often involves considering the agreement among *diverse responses*.

### 5.3 User Guidance

*To help the user explore the design space and align subjective outcomes to their vision, incorporate their input into the workflow at runtime (Figure 8).*

*Motivation.* Users may want to use the workflow to reduce the *effort* of exploring the design space, to align the workflow’s output to their creative vision, or some combination thereof [C[87, 126, 157]]

## User guidance

To help the user explore the design space and align subjective outcomes to their vision, incorporate their input into the workflow at runtime.

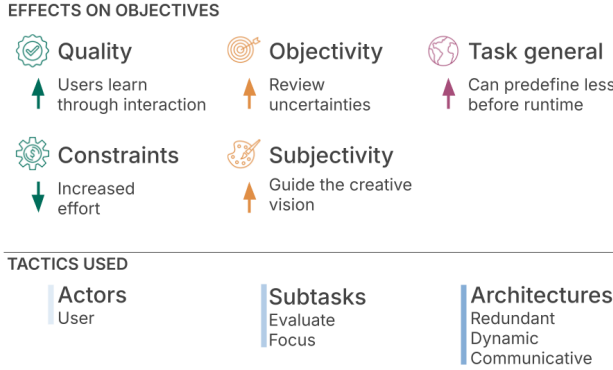


Fig. 8. An overview of the tactics used for and the effects of the user guidance strategy.

L[130]]. If a process includes data that is inherently uncertain it could benefit from more explicit guidance by the user [C[27, 29, 157]].

**Implementation.** This strategy incorporates the user into a critical role within the workflow. Users can guide outcomes by dynamically answering questions and providing feedback [C[137]] or monitor *dynamic* workflows to catch cascading errors [C[86]] with *evaluate* subtasks. Users can be called upon when necessary with a “panic button” that other workers can invoke [C[86]], or automatically when ties occur [C[133]]. With a shared *communicative* interface between users and all workers, users can intervene when needed and make real-time suggestions for prototyping directions [C[92, 126]]. Users may review at intermediate steps to edit or regenerate sections of the workflow in *evaluate* or *focus* subtasks [C[87]] L[78, 130, 173]]. A *redundant* iterative architecture may also include multiple iterations with user revisions between each one [C[29, 87, 157]] L[78, 130]]. Sometimes disagreements raised with *redundant* architectures may not need to be resolved within the workflow and can instead support an explorable interface for users to directly manipulate [C[27, 29, 122, 157]]. They can sort and filter the returned data [C[29, 122, 157]], use a slider to control output characteristics like text length [C[19]], and consider category labels at multiple levels of precision [C[27]].

Crowdsourcing workflows integrate the user in more varied ways than LLM chaining, including providing results without aggregation to sort and filter [C[102, 122]], user guidance for prototyping [C[29, 87, 136]], and actor-initiated user revision [C[137]].

**Secondary Benefits.** Incorporating the user into the process at runtime can reduce the amount of pre-definition of workflow steps needed, making workflows more *generally* applicable [C[87, 126]].

**Consequences.** These benefits come at the cost of more *effort* for the user than an entirely automatic workflow.

**Examples.** Examples of the user guidance strategy include: Second Opinion [C[122]], Social Simulacra [L[130]], and WearWrite [C[126]].

**Related Strategies.** Incorporating user guidance can aid with *validation*.

## Adaptable architectures

To build workflows that require less effort to design and that generalize to many inputs, use adaptive architectures that build themselves at runtime.

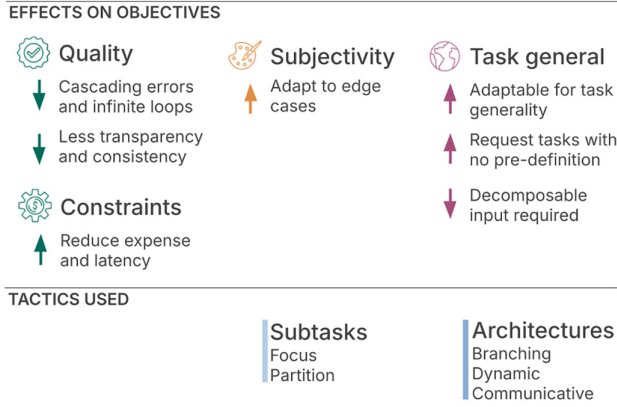


Fig. 9. An overview of the tactics used for and the effects of the adaptable architectures strategy.

### 5.4 Adaptable Architectures

To build workflows that require less effort to design and that generalize to many inputs, use adaptive architectures that build themselves at runtime (Figure 9).

*Motivation.* It is challenging and requires high effort to design a rigid and pre-defined workflow that can effectively generalize to many tasks or that can accommodate for many edge cases, especially in subjective tasks [C[75, 82, 86, 150]]. Instead, designers may want workflows that are adaptable and responsive [C[28, 75, 87, 150] L[36, 124, 158]] and to support a variety of inputs without the effort of explicit pre-definition [C[82, 86] L[36, 37]].

*Implementation.* Adaptable workflows are not rigid in their progression of subtasks but instead are flexible and generate parts of the workflow architecture at runtime [C[28, 75, 82, 87, 150] L[65, 66]]. Workflows can incorporate subtasks defined by the crowd with *partition* and *focus* subtasks in *dynamic* and *communicative* architectures [C[3, 4, 136, 171] L[135]]. Changes in these adaptable workflows may be made by the worker [C[94, 136, 150]], user [C[86, 126, 136, 150]], or an algorithm [C[22, 28, 38, 53]]. Although *branching* architectures are pre-defined, they also allow for adaptability in which different inputs take different paths.

*Consequences.* However, there are tradeoffs. Adaptable workflows can lead to poor quality worker-created subtasks and instructions, causing cascading errors and infinite loops [C[86]]. They also are restricted to tasks with an easily decomposable input [C[82]].


Chaining papers highly emphasize the importance of transparency [L[36, 37, 51, 135, 158, 159, 162]], increasing the risk of using adaptable workflows which may cause unforeseeable errors as each step is not pre-defined. The degree to which adaptable workflows degrade this transparency is yet to be explicitly explored.


Some effects of adaptable workflows on worker experiences matter for crowdsourcing, but do not matter for chaining. Crowdsourcing work has allowed crowdworkers to choose their own tasks as an incentive to improve task performance [C[33, 82, 94, 115, 126, 171]], but the need to incentivize

## Actor selection

To improve outcome quality at each step, curate a high quality actor pool by filtering actors by their skills and improving actor capabilities over time.

### EFFECTS ON OBJECTIVES

 **Quality**  
↑ Higher quality  
subtask responses

 **Constraints**  
↓ Increase cost and  
effort of filtering

### TACTICS USED

**Actors**  
LLMs  
Crowdworkers

Fig. 10. An overview of the tactics used for and the effects of the actor selection strategy.

does not apply to chaining. One crowdsourcing study found users were frustrated with too much adaptability, an emotion that does not apply to LLMs [C[150]].

*Examples.* Examples of adaptable architectures include: map-reduce [C[82]], Flash Teams [C[136]], and microtask programming [C[94]].

*Related Strategies.* Adaptable architectures must support the dynamic creation of *aligned subtasks* that have well-defined *context* and worker instructions. Adaptable architectures can also dynamically change course once *quality thresholds* are met.

### 5.5 Actor Selection

*To improve outcome quality at each step, curate a high quality actor pool by filtering actors by their skills and improving actor capabilities over time (Figure 10).*

*Motivation.* Actors that provide higher-quality responses support *outcome quality* [C[4, 42, 68]].

*Implementation.* In crowdsourcing, selecting high-quality *actors* can mean filtering crowdworkers by pre-specified qualifications or using platforms with more specialized crowdworkers such as Upwork [C[150]]. With support, crowdworkers can improve performance over time [C[28, 42, 83, 156]]. For LLM chains, model choice, specialization, and fine-tuning impact abilities [L[24, 35–37]].

*Consequences.* Higher quality actors may be more *expensive* [C[34, 68] L[24]], although smaller specialized models may reduce *expense* [L[24]]. There is also *effort* and *expense* involved in filtering crowdworkers for skill [C[87, 150, 169]] or fine-tuning models to improve specific skills [L[24]].

*Examples.* Examples that pay particular attention to actor selection include Flash Organizations [C[150]], and chains that fine-tune models [L[24]] or use multiple models [L[100]].

*Related Strategies.* Actor selection works hand in hand with *aligning subtasks* to actor capabilities. Actor selection must also be considered when using *adaptive architectures* to define what kind of workers to hire during runtime.

### 5.6 Subtask Alignment

*To improve outcome quality at each step, align subtasks to actor capabilities (Figure 11).*

## Subtask alignment

To improve outcome quality at each step, align subtasks to the actors' capabilities.

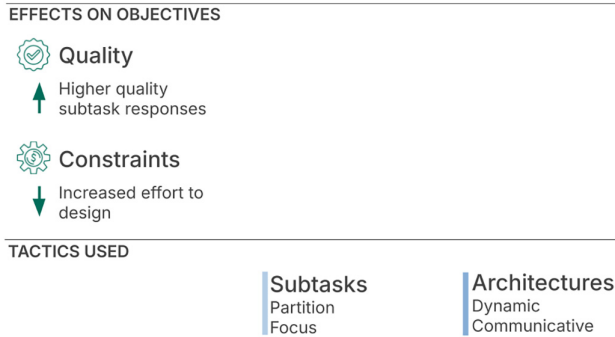


Fig. 11. An overview of the tactics used for and the effects of the subtask alignment strategy.

*Motivation.* One critical strategy is ensuring that *subtasks* match the *actors'* capabilities appropriately. Simplifying tasks can increase *outcome quality* [C[32, 93, 133] L[162]].

*Implementation.* Compared to entire tasks, smaller subtasks are often easier to complete, validate, and correct [C[9] L[36, 37, 78, 139, 162, 166]]. Designers can also adapt the subtask content to fit the actor's capabilities [C[27, 122, 134] L[15]]. For example, researchers found that crowdworkers were better at generating predictive features than at estimating if a feature is predictive; so they adapted the workflow accordingly [C[29]].

When the appropriate subtasks are unknown, *dynamic* architectures can use *partition* or *focus* subtasks to use actors' input to break down tasks into an appropriate size at runtime [C[77, 82, 94] L[66]]. *Communicative* architectures can create appropriate subtasks for actors when they have a running "todo" list from which actors can choose their preferred task of many [C[92, 115, 171]].

*Consequences.* The capabilities of LLMs and appropriate subtask content are far from fully understood. Furthermore, it is unknown the degree to which LLMs can judge the appropriateness of subtasks for other LLMs to complete, bringing into question the effectiveness of some *partition* subtask and *communicative* architecture approaches taken when the actors are *crowdworkers*. Therefore, testing out these different options to find appropriate subtasks is *resource* intensive in the design process.

*Examples.* Examples that pay particular attention to subtask alignment include: Revolt [C[27]], CrowdBand [C[133]], and Ask Me Anything [L[15]].

*Related Strategies.* Subtasks must align to the capabilities of *selected actors* and requirements must be communicated with sufficient *context*. Subtask alignment is especially challenging when subtasks are created during runtime with *adaptable architectures*.

### 5.7 Context

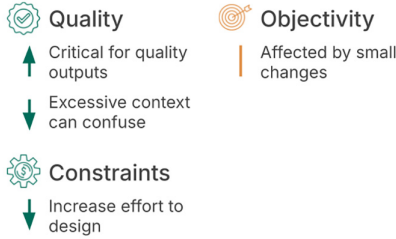
*To promote coherence and improve subtask outcomes, provide the minimum sufficient global context while optimizing for clarity of actor instructions (Figure 12).*

*Motivation.* Workflows need to have high-quality subtask outputs that promote cohesion among disparate parts of the workflow.

## Context

To promote coherence and improve subtask outcomes, provide the minimum sufficient global context while optimizing for clarity of actor instructions.

### EFFECTS ON OBJECTIVES



### TACTICS USED



Fig. 12. An overview of the tactics used for and the effects of the context strategy.

*Implementation.* Communicating context to workers can help with improving the quality of the output and creating more coherence among disparate subtasks [C[9, 19, 33, 60, 86, 111, 147] L[159]]. There are many forms of contextual communication, from input text [C[9, 19, 86]], to workflow goal tracking [C[77] L[121, 166]], to explicit communication between users and workers [C[137]].

Instruction wording is a critical part of all *subtask* designs to support high *outcome quality* [L[129]]. A key difficulty is communicating the user's intent through language. One study reports that crowdworkers tend to need less explicitly defined instruction than LLMs but can still struggle with vague or ambiguous instructions [L[161]]. Both crowdworkers and LLMs benefit from examples, namely screening and tutorial tasks for crowdworkers [C[19, 28, 42, 92, 93, 133, 147]] and few-shot examples for LLMs [L[51, 123]]. Interfaces for crowdworkers can structure their output, while this structure needs to be communicated to LLMs in natural language [C[76] L[36, 161]].

*Consequences.* Although context is generally important, overcomplicating instructions can lead to confusion [C[28, 92] L[114, 161]], so giving the right amount of context is important [C[82]]. One in-depth crowdsourcing study found that providing context from the user was helpful early on in the workflow if content is poor but becomes detrimental if communicated later or with high-quality content [C[137]]. Minimizing context seen can also help with preserving privacy [C[20, 74]]. For LLMs, context may be included as part of the prompt, although LLMs have the additional barrier of an input token limit [L[121, 130]]. Therefore, workflow designs should aim for a minimum amount of sufficient context, to provide coherence while avoiding confusion.

For both crowdworkers and LLMs, small changes in instructions or prompts can lead to dramatically different outcomes, impacting the accuracy of the final result [C[80, 120] L[15, 51, 66, 121, 129, 159, 161, 162]]. Additionally, different LLMs may respond differently to the same input prompt [L[129]]. Therefore, figuring out the clearest instructions requires non-trivial *effort* and testing.

*Examples.* Examples that pay particular attention to context include Turkomatic [C[86]], collaborative writing with microtasks [C[147]], and a study on communicating context [C[137]].

*Related Strategies.* When using *adaptable architectures*, dynamically made tasks can struggle to provide quality worker instructions and context. Additionally, the effectiveness of this strategy is limited if there is improper *subtask alignment*.

## Quality thresholds

To more precisely balance the quality-cost tradeoff, measure quality repeatedly and stop the workflow once a quality threshold is reached.

### EFFECTS ON OBJECTIVES



Quality



Constraints



Improves control  
of tradeoff

### TACTICS USED

Actors

LLMs  
Crowdworkers  
User

Subtasks

Evaluate

Architectures

Redundant  
Dynamic

Fig. 13. An overview of the tactics used for and the effects of the quality thresholds strategy.

## 5.8 Quality Thresholds

To more precisely balance the quality-cost tradeoff, measure quality repeatedly and stop the workflow once a quality threshold is reached (Figure 13).

*Motivation.* Sometimes, designers may not know prior to running the workflow how much work needs to be done to get a high-quality output. In those situations, it would be helpful to have the workflow decide how many resources are necessary to use on an input-specific level.

*Implementation.* Some workflows give the *actor* control over the number of *redundant* generations using *evaluate* subtasks [C[42, 70] L[129, 168]]. Others allow for tuning *redundant* architectures giving flexible goal precision, recall, and agreement thresholds [C[70] L[168]], use decision-theoretic choices of the number of responses to obtain in *redundant* workflows [C[8, 22, 38, 39]], or *dynamically* choose between workflows by their maximum expected value of utility [C[53]].

*Consequences.* Deciding if a threshold is met can require higher costs in terms of *effort* and *expense*.

*Examples.* Examples of this strategy include POMDP-based control [C[38]], Pterodactyl [C[8]], and a workflow for detecting emotionally manipulative language [C[70]].

*Related Strategies.* Quality thresholds may be determined by leaning on the *adaptive architectures* and *user guidance* strategies.

## 6 Using the Design Space and Strategies

In this section, we show the generative value of the design space through two vignettes. In contrast to the paper's case studies (Section 7), which originate with an existing crowdsourcing workflow, here we use the design space to generate the designs of novel workflows from scratch.

After defining our example task of **Retrieval-Augmented Generation (RAG)** (Section 6.1), we present two vignettes that implement this task in different contexts (Sections 6.2 and 6.3, Figure 14). These vignettes exemplify how our design space informs which tactic to employ in response to shifting task contexts. Section 6.4 discusses the function of the design space within the broader design process. While we have neither implemented nor evaluated these chains—and so make no claims that they perform better than existing approaches reported in the literature—they are (to our knowledge) novel and offer the potential for high performance.

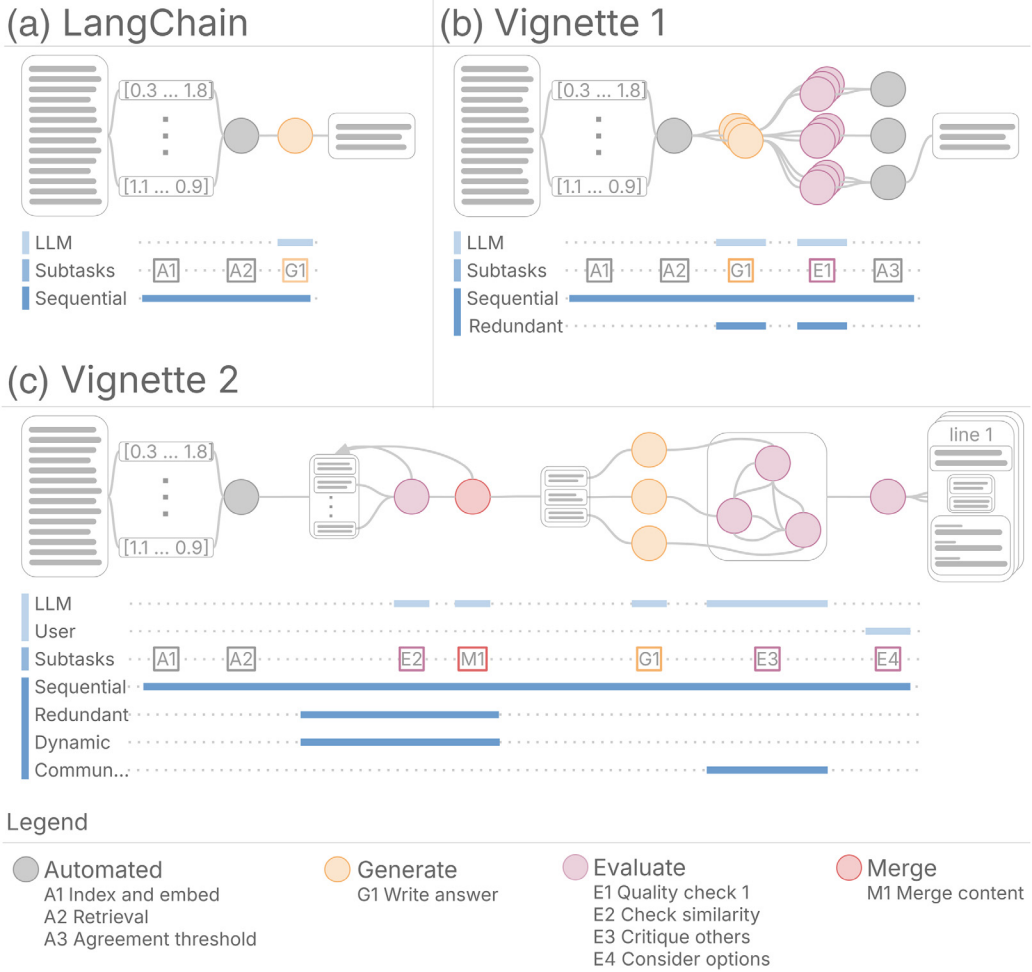


Fig. 14. Vignette chains. We show through two vignettes—both variations of RAG—how our design space can be used to generate novel chain designs. It is common practice for a designer to default to implementing a generic RAG chain, such as (a) the one recommended by the popular chaining software LangChain [91]. While simple, such a design may not fit the specialized needs of a given application. By using our design space a person could be led to consider alternative designs, such as (b) a chain design focused on achieving *objective* answers and a *specific* task with *constrained resources*, or (c) one focused on incorporating *subjectivity* for *general* subjects with high *outcome quality*.

## 6.1 Task: RAG

We consider the task of RAG, in which an LLM generates responses with the help of a larger corpus from which it can retrieve information. RAG systems are usually implemented with a multi-step chaining process due to context size [141]. Specifically, we consider a RAG question-answering task in which the system must answer questions using information from the corpus.

Conceptually, we can break the RAG QA task into retrieval and generation steps. The retrieval step inputs the question and the corpus and outputs relevant information from the corpus. The generation step inputs the question and said relevant information, then outputs an answer. The

recommendations provided by the LLM chaining software LangChain reflect the standard approach for RAG [1]. They recommend chunking and indexing documents in the corpus, using the query to retrieve documents by semantic similarity, then placing the retrieved chunks into an LLM context and generating an answer (Figure 14).

## 6.2 Vignette 1: Customer Service

For this vignette, we consider implementing RAG as a part of a customer service system, intended to answer questions about company policies and terms of service.

*Prioritized Objective: Objectivity.* For customer queries, we will make the assumption that there is just one answer as all information in the corpus should be self-consistent. The *diverse responses* and *validation* strategies are especially useful for improving accuracy on *objective* tasks. *Evaluate* and *improve* subtasks support *validation*. Additionally, requiring a high threshold of agreement across *redundant* outputs, with a plan for reconciliation if outputs do not reach that standard, uses the *diverse responses* strategy to *validate* by measuring consistency.

*Prioritized Objective: Task Specific.* This application of RAG is specific to one corpus about the company policies. The most helpful strategy for this objective is precise *validation*, since the specificity of the task allows for pre-definition of qualities to be validated. The tactical instantiation of this strategy is crafting highly specific *evaluate* subtasks.

*Prioritized Objective: Resource Constraints.* *Latency* and user *effort* lead to customer frustration so are important to minimize in a customer service system. *Expense* and *effort* to design are also desirable to minimize, but less critical in this task context. To reduce the strain on the latency of the *diverse responses* and *validation* strategies, the architecture should be a parallel *redundant* one, rather than iterative or *sequential*. Spending effort *selecting actors*, designing *context*, and aligning *subtasks* reduces latency, expense, and user effort with less impact on quality. Although the *adaptable architectures* and *quality thresholds* strategies can reduce latency, expense, and effort to design, they can induce quality issues best mitigated through *user guidance*. As reducing user effort is prioritized, this strategy should be deprioritized.

*Chain Design.* An objectives-driven chain for this customer service task pieces together the various suggested tactics into a cohesive chain such as the following. After indexing and retrieval, the selected chunks are fed into the generation portion of the chain. This portion has multiple *redundant generate* subtasks to answer the query based on the retrieved information. They are then passed to *redundant evaluate* subtasks, in which the highest scoring generation provides the output. The number and subject of these evaluations should be determined through testing. Finally, if all outputs do not meet an agreement threshold, alternative outcomes are provided such as requesting a rewording of the query or connecting to a representative. See Figure 14 for a visual representation.

## 6.3 Vignette 2: Fitness Corpus

For the second vignette, we consider a RAG system informed by a corpus of scientific papers to answer a wide variety of questions around fitness topics, such as exercise, nutrition, and sleep.

*Prioritized Objective: Subjectivity.* *Objectivity* in providing correct answers continues to be critical, but this application must incorporate uncertainty and open-endedness more than the first vignette, as the corpus of literature can support differing conclusions. The strategies of *diverse responses*, *validation*, and *user guidance* are especially helpful in addressing this uncertainty. To implement these strategies, the *user* can provide input at runtime to *validate* what output is desirable when there is uncertainty in the answer. Measuring uncertainty requires sourcing *diverse responses*, either

through measuring variation within a *redundant* architecture or measuring ease of persuasion in the debate *communicative* architecture.

*Prioritized Objective: Task General.* The RAG system must generalize to a wider variety of queries and topics than the first vignette. *Adaptable architectures* provide the flexibility to support such variety. As the variety of queries is not pre-defined, *branching* architectures are not a useful tactical instantiation of this strategy. Instead, *dynamic* and *communicative* architectures adapt at runtime to an unknown set of queries and topics. Finally, incorporating *user guidance* supports flexibility in result interpretation and reduces the need to specify a pre-defined set of topics or queries.

*Prioritized Objective: Quality.* Although quality is always desirable, in high-stakes domains like the medical field prioritizing this objective is critical even at higher expense, latency, or effort. For the purposes of this vignette, we assume few resource constraints. Of all the objectives, prioritizing an increase in quality has the widest range of options, as nearly all strategies and their tactical implementations improve quality. The exception is that *adaptable architectures* can induce errors and reduce transparency and consistency. In this health querying system, the *user* needs appropriate *context* to ensure the qualities of transparency and faithfulness. The chain's tactics can attribute outputs to verbatim sections of the corpus to improve the user's context.

*Chain Design.* By combining these prioritized tactics, an objectives-driven chain for this case could look like the following. The initial retrieval step allows for multiple chunks of the corpus to pass through. To accommodate for the potential disagreement from a variety of sources, the next step iteratively merges the corpus chunks into several distinct lines of argument, connected with references to verbatim parts of the corpus. At each step of this iterative *redundant* architecture, a pairwise comparison *evaluate* subtask determines if two chunks support similar takeaways. If they do, a subsequent subtask *merges* the two into a new chunk, collecting the indexes referencing specific parts of the corpus. This iterative loop runs for a *dynamically* determined amount of time until the arguments have converged. These arguments and their references then pass into the generation portion of the pipeline. Each argument informs a *generation* subtask that answers the query given that set of retrieved information. Then, a *communicative* debate commences in which each debater begins with this generated answer and its justification. Using the ease with which each debater is convinced it is incorrect, the chain outputs the different arguments, their uncertainty, the debate history, and references to the initial corpus for the *user* to *evaluate*. See Figure 14 for a visual representation of this chain.

## 6.4 Vignette Discussion

Through these two vignettes, we outline how the design space specializes chains to the task context by explicitly defining objectives and surfacing the cascading implications for strategies that then recommend beneficial tactic choices. In this section, we discuss how the process described in these vignettes fits into the broader design process.

*Tactic Suggestions Respond to Task Context.* Standard practice applies the same steps for any task context: retrieval and generation. With our design space, a designer can flexibly accommodate unique combinations of objectives for varied applications, as seen in the differing vignette designs.

*Incorporating Non-LLM Tools Can Reduce Costs.* The costs of using LLMs for all steps can be prohibitive and unnecessary. A designer could replace the semantic similarity retrieval step with LLMs by chunking the corpus and using *focus* subtasks to determine relevant parts of each chunk. Although the potential for LLMs for processing complex unstructured data is an active area of exploration [140], using LLMs for this step would cost much more than standard retrieval with

a large corpus. Our design space focuses on the LLM chaining literature, but there is immense opportunity and future work in incorporating tools beyond LLMs into these chains.

*The Design Space Filters and Inspires, Rather Than Prescribes.* Using the design space does not prescribe a definitively optimal chain, but it does serve the design process through filtering and generating ideas. Through the vignettes, we see that the design space prioritizes subsets of options within a task context. Given the enormous space of possible designs, this filtering is valuable guidance. Furthermore, using the design space allowed us to create unique chains. Due to the comprehensiveness of the design space over two literatures, it serves as inspiration during the brainstorming process for a wide variety of potentially unintuitive options validated through prior work.

*The Design Space Provides an Informed Starting Point, Rather Than a Finished Design.* The design space serves as a launching pad for other design decisions. Further testing can compare performance within the subset of tactic choices recommended by the design space (e.g., which type of *verification* performs best). Many of these changes depend on model performance (e.g., are diverse generations necessary, or is fluency high enough to have confidence in a single generation). Additionally, the design process will need to test specific tactical design decisions (e.g., what model to choose, prompt engineering, how many iterations, and for which specific qualities to evaluate). Although these chain designs may update after testing on specific corpora and models, they provide a starting point motivated by findings from prior work. Future work is needed to improve the granularity of design choice recommendations and develop evaluation mechanisms for comparing among design choices (Section 9.4).

## 7 Case Studies

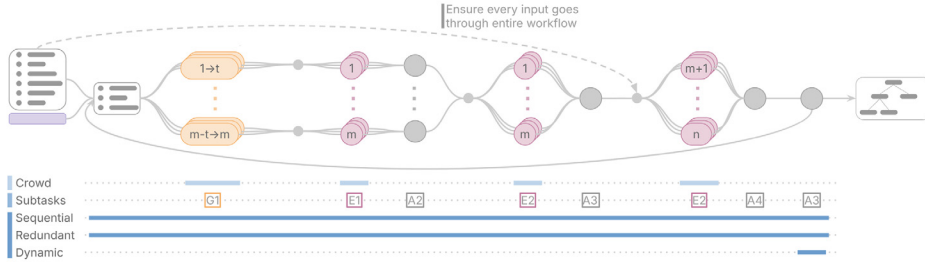
In order to better assess how crowdsourcing methods transfer to LLM chain development, we adapt crowdsourcing workflows to LLM chains in three case studies. This process of adapting workflows supports two investigations. First, it allows us to compare the performance of chaining to both zero-shot prompting and crowdsourcing. Second, through the process of adapting these workflows from crowdsourcing to chaining, we can see how the workflow design changes to accommodate for LLMs. These changes, found through practice, directly compare what *tactics* support the same *strategies* and *objectives* when using *crowdworkers* or *LLMs*.

To ensure we are suitably focused, we explore one strategy in more depth: supporting quality control through *user guidance* in which the user can directly manipulate workflow outputs. We focus on this strategy as it is implemented in crowdsourcing work, but it is under-explored in the LLM chaining space. Prior non-chaining LLM work explores controllability as a helpful attribute [79, 112], so here we seek to incorporate it within chains. With directly manipulable outputs, users can make “rapid, incremental, reversible operations whose impact on the object of interest is immediately visible” [143], increasing control over output elements and helping users explore alternative results [71]. By adapting this subset of the user guidance strategy from crowdsourcing, we explore how and if such a concept can transfer to chains, what modifications are necessary, and its relative advantages compared to controlling the outputs through zero-shot prompting.

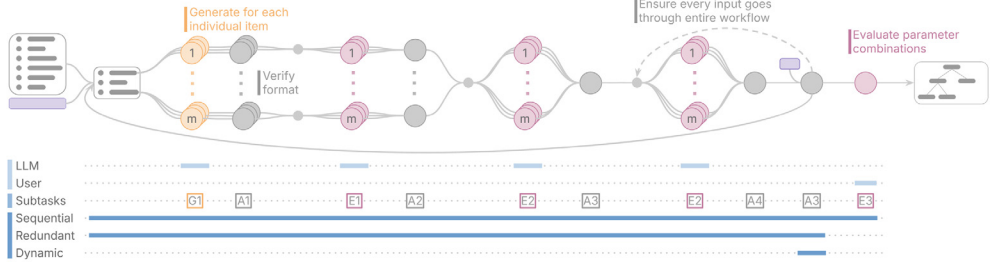
We choose three popular crowdsourcing tasks that incorporate different degrees of the *objectivity* and *subjectivity* design space objectives. We implement Cascade [32] to build taxonomies with manipulable structures, Soylent [19] to shorten text to a variety of target lengths, and Mechanical Novel [77] to control the elements of a story. Diagrams of the workflows, as well as our adaptations to use LLMs, can be found in Figures 15, 18, and 20.

For each case study, we first document how the workflow was changed to incorporate LLMs and map those changes to the design space. Common limitations we faced include LLM context

## Crowdsourced Cascade



## Cascade-LLM



## Legend

Parameters	Automated	Generate	Evaluate
Minimum category size	A1 Format verification	G1 Generate	E1 SelectBest
Merge threshold	A2 Vote threshold		E2 Categorize
	A3 Taxonomize		E3 Choose option
	A4 Remove underspecified		

Fig. 15. Case study 1: Cascade [32] for taxonomy creation. Crowdsourced Cascade uses various actors, subtasks, and architectures from the design space, as defined under the workflow diagram. Although automated events are not defined as subtasks in our design space, we indicate them here for clarity. Crowdsourcing workflows can inspire LLM chains, but the differing capabilities of crowdworkers and LLMs require changes to the design. We adapt the crowdsourcing workflow Cascade with elements from our design space to update the use of tactics to better work for LLMs. For example, we found that generating responses for individual items, rather than sets of items, is a more appropriate subtask because it creates better *response diversity*. Through these adaptations, we can control the output taxonomy shape with parameters.

lengths, inconsistent output formatting, lack of diversity in model responses, and models struggling to perform the workflow-defined subtasks. We adapted to these limitations with more granular subtasks, more diverse prompts, comparative tasks, and hard-coded validation steps. Across all three case studies, we vary prompts<sup>1</sup> at each parallel generation step to increase *response diversity*, update worker instructions and *context* through trial and error to get the right output formatting, and structure voting *evaluation* subtasks to choose among numbered options to minimize hallucinations. We share the prompts in [Supplementary materials](#).

We compare our LLM workflows to zero-shot baselines in terms of output control precision, salient qualities for each task, and the latency and number of API calls. To prioritize human-preferred LLM outputs, we use an instruction-tuned model: OpenAI’s API gpt-3.5-turbo with default parameters [128]. Compared to zero-shot prompting, chained workflows use more resources but can improve output quality. These workflows also provide a greater precision of output control than parameterized zero-shot baselines.

<sup>1</sup><https://github.com/madeleinegrunde/designingchains>

## 7.1 Case Study 1: Taxonomy Creation

For our first case study, we explore taxonomy creation, the task of organizing a set of items into a category tree to aid sensemaking. This task's outcome qualities are primarily *objective*, with some degree of *subjectivity*: categories need to be correct, but most items can fit into multiple categorizations of different levels of precision. We focus on the ability to control the granularity of these categorizations by controlling the taxonomy shape (Figure 17).

**7.1.1 Workflow Design.** We adapt the Cascade [32] crowdsourcing workflow (Figure 15), which takes as input a list of items and outputs a tree mapping items to categories. The workflow begins by selecting a subset of 32 items, *generating* category suggestions for each item, *evaluating* by selecting the best categories among the generated options, and *evaluating* by confirming which among all categories works for each item. The next step turns the item-category assignments into a taxonomy, deleting small categories and merging similar categories. The workflow then updates the categorizations for the entire list of items with this new assignment. The workflow *redundantly* iterates with new subsets until all items are categorized, building a growing set of categorizations.

To build Cascade-LLM, our version of Cascade powered by LLMs, we made multiple changes to the original workflow. We highlight some of the most salient modifications. To achieve diverse categorizations, we improved *response diversity* by giving different prompts to parallel generations and ensuring that every item is input into a category *generation* step. At first, the categories were too broad. Specifying category qualities via targeted prompts proved unsuccessful, so instead we added a programmatic *validation* step to delete categories that included “and,” “or,” commas, or slashes. As this hard-coded *validation* step worked, it was not necessary to include another LLM call. The original workflow gives actors sets of eight items for which to generate eight categories. However, the LLM would try categorizing multiple items at once, consistent with prior findings that LLMs struggle to disambiguate multiple subtasks in a prompt [114, 161]. To increase diversity, we *aligned subtask* design to fit model capabilities with one item categorization. Finally, the original workflow also requires parameters for the minimum category size and a merge threshold. To support *user guidance* through direct manipulation, we dynamically apply these parameters at the last taxonomization stage. To investigate different taxonomies with varied levels of precision, a user can manipulate the parameters of (1) the minimum category size and (2) the similarity threshold determining if two categories will merge (Figure 17).

**7.1.2 Evaluation Setup.** We compare Cascade-LLM against zero-shot baselines on three lists of nouns.

**Datasets.** Our experiments use the labels from the scenes in MIT Indoor Scenes [134], from object categories in MSCOCO [103], and from CIFAR100 [85]. These label sets are of size 67, 80, and 100, respectively. The original Cascade paper explores three datasets sourced from Quora.com question responses. Our LLM workflow is capable of processing these larger datasets, as only individual items need to be within the model's context size; however, the Quora.com data exceed the context limit for the zero-shot baselines.

**Baselines.** We compare our workflow to two baselines. The *zero-shot* baseline directly prompts the model to create a taxonomy. This baseline establishes how the model completes the task given no restrictions. The *zero-shot-target* baseline further specifies the target number of categories in the prompt. With this baseline, we compare the ability to control workflow outputs in a single prompt.

**7.1.3 Workflows Enable Direct Manipulation.** To measure the controllability of the taxonomy's shape, we measure how well methods can change the taxonomy size (the number of categories).

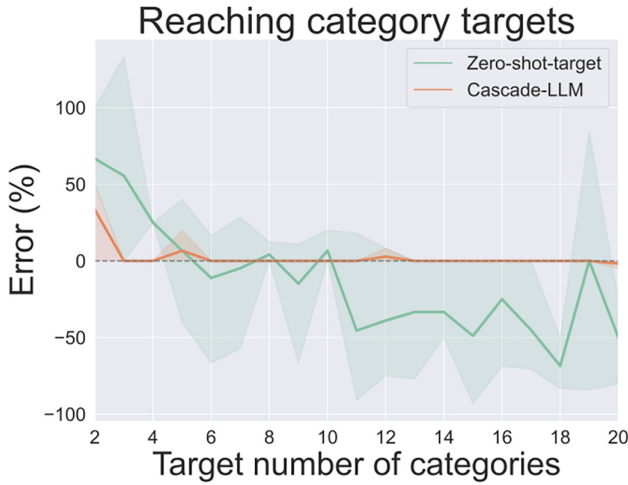


Fig. 16. Controllability of Cascade-LLM vs. zero-shot taxonomies. Manipulating Cascade-LLM outputs through given parameters can more precisely achieve a wide range of category sizes than indicating a target value within a zero-shot prompt. Error (%) refers to the percent error between the method's closest option and the target value. Across three datasets, Cascade-LLM exhibits low deviation from the target number of categories, while prompting with zero-shot gives noisy responses, often with too few categories. The shaded region indicates the 95% CI.

Given a target number of categories (range: 2–20), we measure the percent error between the method's closest option and the target value.

As seen in Figure 16, one run of Cascade-LLM almost always achieves a precise match with the desired number of categories. In contrast, zero-shot-target is close but usually unable to produce a taxonomy with the target number of categories. The zero-shot condition produces small taxonomies the size of which is not controllable (MIT Indoor Scenes: 2, MSCOCO: 8, CIFAR100: 4).

**7.1.4 Workflows Reduce Hallucinated and Forgotten Items.** We also evaluate the quality of the outputs through qualitative evaluation. Similar to the original paper, we hand-code taxonomies without knowledge of which method created it. First, we look for errors in items across all baselines. Specifically, we see if the output hallucinates or forgets any items. We then code the zero-shot and default-parameter Cascade-LLM baselines for the same mistakes considered in the original paper: duplicate categories, missing parent-child relationships, and incorrect parent-child relationships.

Cascade-LLM does not forget any items, but 83.72% of zero-shot taxonomies do not include all items. Cascade-LLM cannot hallucinate new items, while zero-shot baselines hallucinate 51 items across the 129 taxonomies. Cascade-LLM taxonomies have similar error rates to zero-shot baselines' taxonomies on duplicate categories and missing/incorrect category-item and parent-child pairs, yet are richer taxonomies with more categories and depth (Table 3).

**7.1.5 Workflows Have Higher Costs.** We measure generation time and the number of API calls to an LLM. Cascade-LLM takes more time (mean: 2,642.58 seconds, SD: 655.81) than the zero-shot (mean: 24.35 seconds, SD: 5.04) and zero-shot-target (mean: 21.69 seconds, SD: 4.49) baselines. Crowdsourced Cascade took 43 hours and 3 minutes to taxonomize three item sets [32]. Cascade-LLM makes more model calls (mean: 18,898.33, SD: 4,472.61) than the one call from the zero-shot and zero-shot-target baselines. Most of these calls are short in terms of token length, as they compare one item to one category. Crowdsourced Cascade reports 1,760 calls to crowdworkers per

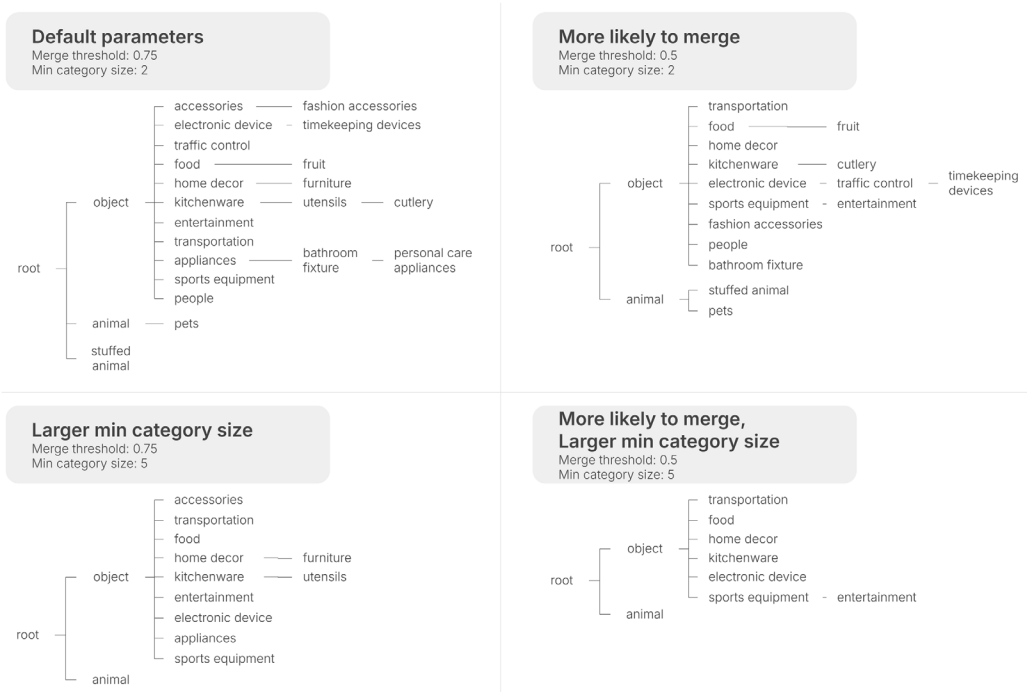


Fig. 17. Taxonomy examples using Cascade-LLM on MSCOCO labels [103]. Categorization into a taxonomy is a challenging task because the same item may be categorized with different levels of precision. Our implementation of the Cascade crowdsourcing workflow [32] enables manipulation of the similarity threshold needed to merge categories, along with the minimum category size. By changing these values, we can sort data into taxonomies of different levels of precision and compare options as the final stage of the workflow. This direct manipulation allows the user not only to validate any inaccuracies but also to make opinionated decisions, such as if “stuffed animal” should or should not be a sub-category of “animal.”

Table 3. Cascade-LLM vs. Zero-Shot Performance across Three Datasets

	MIT Indoor Scenes [134]		MSCOCO [103]		CIFAR100 [85]	
	Zero-Shot	Cascade-LLM	Zero-Shot	Cascade-LLM	Zero-Shot	Cascade-LLM
# of categories	2	21	8	24	4	19
Tree depth	1	3	1	4	1	3
Duplicate categories	0	0	0	0	0	0
Missing category-item	0	2	1	3	7	2
Incorrect category-item	0	0	1	2	2	2
Missing parent-child	N/A	0	N/A	1	N/A	0
Incorrect parent-child	N/A	0	N/A	1	N/A	1

Though the quality of categorizations is similar between methods, the Cascade-LLM generated taxonomies include more categories and have greater taxonomy depth.

iteration. For our goals, the increased cost in time and computation enables direct manipulation of more detailed categorizations without item forgetting or hallucination.

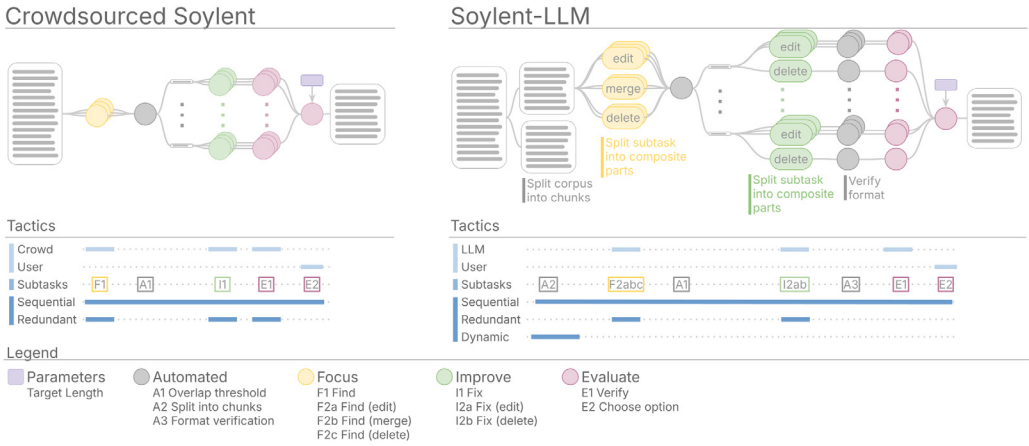


Fig. 18. Case study 2: Soylent [19] for text shortening. When using LLMs, the crowdsourcing workflow Soylent initially outputs low-quality responses and runs over context limits. In response, we build Soylent-LLM with updated tactics to have appropriately specified subtasks, to work within context limits, and to verify there were no hallucinated quotes.

## 7.2 Case Study 2: Text Shortening

The next case study is the task of shortening text. Unlike summarization, shortening the text involves making edits to create a condensed version of the text that attempts to preserve the full meaning. Shortening is a task primarily on the *subjectivity* side, but requiring elements of *objectivity*. Shortened versions must stay faithful to the original text, and there are many options for what phrases to shorten and how to shorten them. By implementing a workflow to generate many shortening options, an end user can control the final length of a paragraph dynamically.

**7.2.1 Workflow Design.** To shorten text we adapt the Find-Fix-Verify workflow (Figure 18) introduced in Soylent [19]. In the “Find” step, crowdworkers select areas that might be shortened. If at least two actors agree on a text span, that span progresses to the “Fix” step. In the “Fix” step, crowdworkers make edits to shorten the text. In the final “Verify” step, more crowdworkers verify the possible changes by rejecting by majority vote any edits deemed grammatically incorrect or that change the meaning of the sentence. These steps respectively map to the *focus*, *improve*, and *evaluate* subtask types in our design space. The workflow outputs edit options for multiple spans in the paragraph, which when combined create multiple possible outcome texts. Soylent defaults to showing the shortest possible option, but it also provides a length slider for the user to adjust the final output as desired. The slider enables the user to test different shortening phrasing options and lengths through a simple interaction.

We made multiple adjustments to produce our Soylent-LLM workflow. First, we ran into context limits and parallelization constraints with longer texts. In response, we apply a *dynamic* approach that splits the workflow into chunks of 10 sentences and applies the workflow to each. Second, Soylent allows actors to fix inputs by editing, merging, and deleting phrases. LLMs performed better when we split these tasks into separate prompts, so we further decompose Find and Fix into finer grained editing, merging, and deleting steps to get an *aligned* subtask size to the model’s capabilities. Finally, in the Find step, the LLM would sometimes “hallucinate” phrases to shorten that do not exist in the text or would change quotes. To mitigate these errors, we add several hard-coded *validation* steps: we ensure that the phrase to replace exists in the input text, that the replacement is shorter, and that no quotes were altered or added.

Recent work also implemented three unique variations of a Find-Fix-Verify workflow using LLM chains [161]. Their evaluation, using student preferences, found mixed results on effectiveness for creating better shortenings than a zero-shot baseline. We evaluate our workflow along different dimensions: controllability of output length, fluency, grammaticality, what type of content is deleted, and workflow resource usage. Unlike this prior work, we also retain the parallel generation and validation steps from the original paper.

**7.2.2 Evaluation Setup.** For evaluating SoyLent-LLM, we use texts from newspapers and from the original SoyLent paper [19]. We compare SoyLent-LLM with three zero-shot baselines.

**Dataset.** We use the XSUM summarization dataset [125]. We only use the input texts because our task is shortening, not matching gold-standard summarizations. We use 100 randomly sampled texts from XSUM, with the requirement that examples are more than 500 words long. For qualitative analysis, we use the texts from the SoyLent evaluation [19]. Through this dataset, we can compare the SoyLent-LLM shortenings to Crowdsourced SoyLent shortenings.

**Baselines.** The *zero-shot* baseline prompts the model to shorten the paragraph directly. For the *zero-shot-target* baseline, we specify the target output length in a zero-shot prompt. In the *zero-shot-ffv* baseline, we ask the LLM to explicitly perform Find-Fix-Verify steps in one prompt.

**7.2.3 Workflows Enable Length Manipulation.** We first compare how well the various LLM approaches reach a desired output length. We represent target lengths through word counts of 100–500 words. In initial experiments, we found that zero-shot baselines perform more competitively when asked to shorten to a given word count rather than to a length percentage and hence use word counts in targeted prompts. Nevertheless, we also measure accuracy in terms of a target percentage of length. We calculate the target percent length by dividing the target number of words by the input text’s number of words, and we calculate the actual percent length as the output text’s number of words divided by the input text’s number of words. We compare the percent difference (i.e., error) between the actual and target percent lengths.

We find that as text length increases, SoyLent-LLM can more precisely meet the target length (Figure 19). When shortening text, the zero-shot prompt condition consistently outputs approximately 100 word long shortenings (mean: 126.14, SD: 52.63). Thus, for longer target lengths, this method achieves lower performance. Zero-shot-target is similarly best at lengths close to 100 words, but it still shows a lack of precision in reaching longer target lengths. SoyLent-LLM and zero-shot-ffv both struggle to hit shorter targets. This result matches the SoyLent paper, in which output texts were between 78–90% of the original text length [19]. SoyLent-LLM can more precisely hit the larger target lengths than any of the baselines.

We find similar trends measuring percent error with the percent length reduction required to meet the target length. SoyLent-LLM and zero-shot-ffv can perform poorly with large reductions, in which the target % is a small fraction of the input text, while zero-shot and zero-shot-target perform worse with smaller reductions. In summary, SoyLent-LLM does a better job of tightening long texts but is less well-suited to extreme shortenings that cross over into summarization. In addition, SoyLent-LLM is unique in providing options among all targets without having to requery. Moreover, at each target length, there are often multiple possible output texts.

**7.2.4 Workflow Outputs Have Abrupt Transitions But Lower Perplexity.** We measure the average perplexity of output texts as a proxy for fluency. We use a different language model (Facebook’s OPT-1.3b) than the actor model (OpenAI’s gpt-3.5-turbo) for an external evaluation. We also perform a qualitative analysis of the shortening procedure. We use the five texts used for evaluation in the original SoyLent paper [19]. We compare the edits made with the shortest output from SoyLent-LLM,

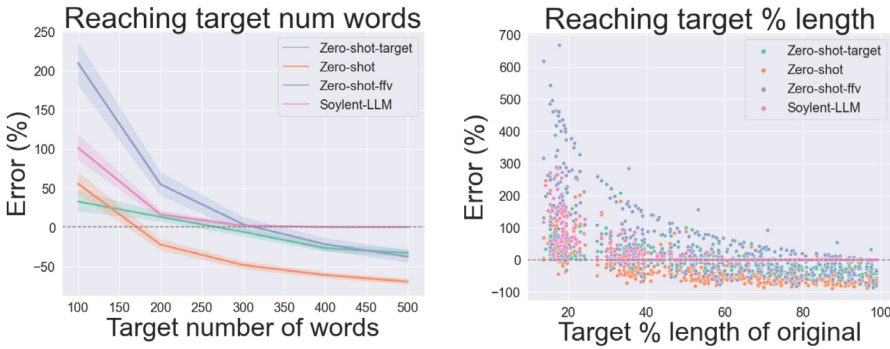


Fig. 19. Text shortening results for Soylent-LLM and zero-shot baselines. A 0% error indicates outputs that exactly match a target length, and the shaded region (left) indicates the 95% CI. Naïve zero-shot prompting without length guidance averages between 100–200 words, regardless of the desired target length. With length guidance added to the prompt, zero-shot-target does substantially better at reaching all targets, especially when the target length is short. Soylent-LLM struggles to hit short targets but is quite precise at longer lengths (left) and when reducing only a short amount relative to the original length (right). The zero-shot-ffv baseline, which instructs the model to perform Find-Fix-Verify steps within one prompt, behaves more similarly to the Soylent-LLM workflow but is consistently less accurate.

the baseline conditions, and the crowdsourced condition’s output provided by the authors of the original paper. Without knowing the condition, we coded the outputs for fluency and grammatical errors, following the qualitative evaluation procedure of the original paper. We additionally coded what content changed during shortening.

Soylent-LLM has the better result of a lower perplexity (mean: 11.69, SD: 3.77) than zero-shot (mean: 15.37, SD: 5.73), zero-shot-target (mean: 12.31, SD: 3.91), and zero-shot-ffv (mean: 16.19, SD: 5.68). This result surprised us, as zero-shot is fully generated by an LLM fine-tuned to provide fluent responses.

As with the original crowdsourcing workflow, some shortenings diminished fluency or introduced grammatical errors. Zero-shot, zero-shot-ffv, and Soylent-LLM shortenings all included passive voice but were otherwise grammatical. However, Soylent-LLM and zero-shot-ffv both could contain abrupt endings or transitions. This abruptness is due to separately editing parts of the input text. Crowdsourced Soylent induced grammatical errors as well as abrupt transitions. We found that zero-shot and Soylent-LLM both shorten by removing unnecessary phrasing and deleting examples and justifications. Zero-shot more successfully merges sentences than Soylent-LLM, while Soylent-LLM changes less text and more successfully retains the user’s voice, rather than formalizing it. Soylent-LLM’s changes affect a larger portion of the text than Crowdsourced Soylent. In sum, Crowdsourced Soylent made 42 edits to phrases and 12 to entire sentences, while Soylent-LLM made 12 edits to phrases and 36 to entire sentences.

**7.2.5 Workflows Have Higher Costs Than Zero-Shot Prompting.** We measure the time it takes to generate the outputs, as well as the number of API calls to the model. The number of API calls varies depending on the length of the input and the number of items that are selected at each step.

Soylent-LLM takes much longer (mean: 105.44 seconds, SD: 60.16) than the zero-shot (mean: 9.09 seconds, SD: 2.92), zero-shot-target (mean: 13.65 seconds, SD: 5.47), and zero-shot-ffv (mean: 43.55 seconds, SD: 23.16) baselines. Crowdsourced Soylent takes the longest with a reported range of 2,640–29,340 seconds (44–489 minutes) [19]. Compared to the one call to the model for the zero-shot

baselines, one Soylent-LLM run averages 127.11 calls to the model (SD 63.56). Crowdsourced Soylent reports incorporating 158–362 crowdworkers.

**7.2.6 Chaining Demonstrates Tradeoffs with CoT.** *Zero-shot-ffv* is a CoT [155] baseline, in which the LLM is prompted to use the steps of the Find-Fix-Verify chain all in one prompt-response set. As previously discussed, *zero-shot-ffv* follows the quality trends of Soylent-LLM more closely than those of the *zero-shot* baseline. Compared to Soylent-LLM, the *zero-shot-ffv* baseline’s quality and cost are lower. Therefore, the relative utility of doing a chain or a CoT depends on the restrictions on cost and the importance of high quality. We did not do a CoT baseline for the other two case studies, as context lengths did not enable the chain to be represented in one prompt and response. Therefore, another benefit of chains over CoT is the capability to do tasks and follow logical flows outside the context limit of an LLM. These restrictions will be subject to change as the context limits of LLMs grow.

### 7.3 Case Study 3: Short Story Generation

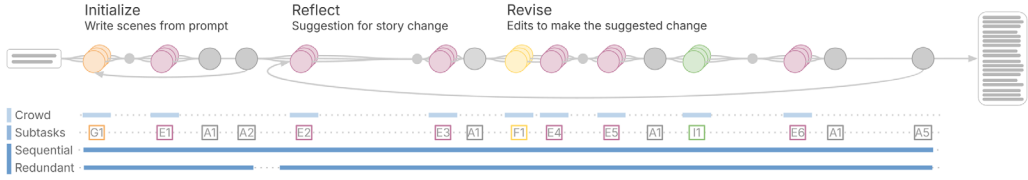
The final case study concerns story generation given a short text prompt. This task lies on the open-ended *subjectivity* side of the spectrum, as short stories have a wide design space of possible plots, characters, messages, imagery, and lessons. In this task, *user guidance* of the story involves higher-level control of different plot elements, with the ability to see counterfactuals (e.g., if a story element is included or not). We implement a workflow that, compared to zero-shot approaches, enables this exploration more precisely, while also creating stories that are preferred by crowdsourced raters for imagery, originality, and style.

**7.3.1 Workflow Design.** We adapt the Mechanical Novel (MN) [77] workflow for this task. This workflow involves three main steps (Figure 20). The first *generate* subtask “initializes” a story with multiple scenes based on a given story prompt (e.g., “Malcolm finds himself alone in a runaway hot air balloon and accidentally travels to a city in the sky.”). Then, the “reflect” and “revise” steps are iterated five times in a loop. In the “reflect” step, actors consider what changes might helpfully adjust the story. In the “revise” step, actors implement those changes. Each step has subtasks for parallel *generation* of suggestions and edits, as well as *evaluate* subtasks voting for the best generation. Crowdworkers additionally *focus* editing attention on the most important scenes.

We adapted the Crowdsourced MN to MN-LLM by incorporating fine-grained *user guidance* with manipulatable outputs. At the conclusion of the workflow, instead of receiving just one or more separate short stories, the user can mark check-boxes on an interface selecting different combinations of edits to view stories with different qualities and compare counterfactuals on editing decisions. To enable these outputs, we run the “revise” stage on texts of all combinations of suggestions from the “reflect” stage.

Beyond introducing control of story elements, the largest challenge was managing token context limitations. *Getting response* diversity with iterative improvement on texts often increases the text length [104]. To account for context length, we generate stories with four scenes rather than six and *iterate* through scenes rather than informing actors to select scenes. In our design space this change maps to using a *redundant* iterative architecture rather than a *sequential* architecture with a *focus* subtask. At the beginning of the “revise” step, we add a *generation* subtask that summarizes the current story to provide abbreviated context in the *worker instructions*. We included a hard-coded *validation* step with an *evaluate* subtask to disqualify edits that were more than 1.5 times the input text length. Unfortunately, this disqualified many highly descriptive texts and sometimes meant that no edits were made because all suggestions were too long.

## Crowdsourced Mechanical Novel



## Mechanical Novel-LLM

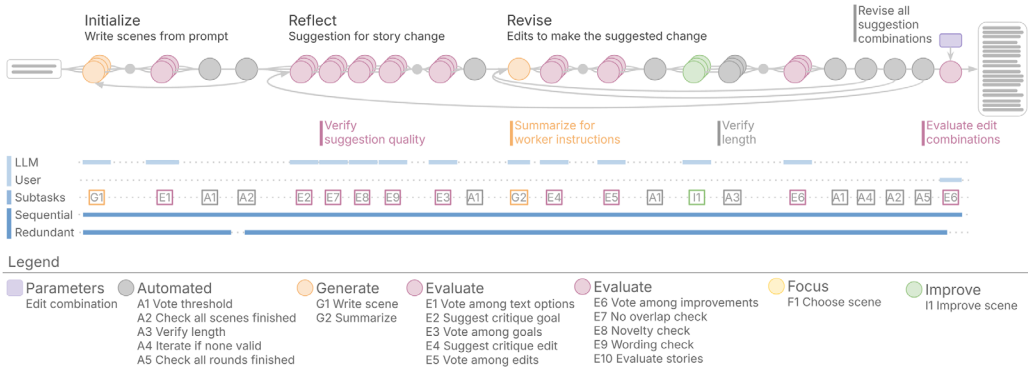


Fig. 20. Case study 3: Mechanical Novel [77] for short story generation. We alter the tactics used in the crowdsourcing workflow Mechanical Novel for LLMs. These adaptations include a new summarization step, more validation steps, and iterations over different input subsets. These changes fit the task within context limits, encourage diverse outputs, and enable control of plot elements.

Table 4. Crowdworke Preferences between MN-LLM and Zero-Shot Baseline Generated Stories

	Imagery	Coherency	Plot	Originality	Style	Technical	Overall
MN-LLM	124	97	93	137	118	97	112
Zero-shot	76	103	107	63	82	103	88

More crowdworkers voted for MN-LLM stories in terms of imagery, originality, style, and overall preference. However, crowdworkers judged the zero-shot stories to have a more coherent and complete plot, reflecting challenges with outputting coherent answers from workflows.

Our initial pilots generated editing suggestions that were very generic (e.g., “Add more detail to the story”) and were repetitive across iteration rounds. To improve the *diversity* of suggestion outputs, we tried changing the *worker instructions* to be contrastive by providing previously used “banned” suggestions. However, the outputs anecdotally seemed more likely to follow those of prior rounds after this change. We hypothesize that the length of the context and request to do two tasks in one (generate a suggestion and avoid the prior suggestions) led to a lack of diversity. Instead, we add several *validation steps* to promote suggestion diversity and quality. This *subtask* granularity worked better. In these *evaluate* subtasks, we prompted the LLM to ensure that (1) the suggestion did not overlap with prior suggestions or two seeded generic suggestions (“*Generic recommendation to introduce conflict.*” and “*Generic recommendation to add detail.*”), (2) the suggestion has not already

Table 5. Controllability of Story Plot Elements with MN-LLM and Zero-Shot Baselines

	Zero-Shot	MN-LLM
Avg. # combinations	3.80	20.40
Avg. length in words	276.38	642.93
Precision	77.05	97.74
Recall	88.68	90.00

We find that the MN-LLM workflow creates more combinations of long story options. Both approaches incorporate intended story elements at similar rates (recall), but the zero-shot baseline more often includes unintended story elements (precision).

been implemented in the story, and (3) that the suggestion is worded as a suggestion, not new text for the story.

**7.3.2 Evaluation Setup.** To evaluate MN-LLM, we use the story prompts and evaluation procedure from the original paper [77] and compare MN-LLM’s results to those of two baselines.

**Dataset.** Our dataset consists of the five short story prompts in the MN paper [77].

**Baselines.** The *zero-shot* baseline method prompts the model to generate five suggestions, then write one story that incorporates these suggestions. The *zero-shot-combo* baseline method prompts the model to generate five suggestions, then write different versions of a story based on combinations of those suggestions.

**7.3.3 Workflows Enable Control of Plot Elements.** We assess each output to evaluate the number and quality of suggestion combinations that emerge from MN-LLM and zero-shot-combo. To measure quality, we calculate the length of texts, as well as the precision and recall of the included suggestions. To measure precision, we calculate what percentage of included suggestions were supposed to be included. To measure recall, we calculate what percentage of suggestions that were supposed to be included were included. As MN-LLM is iterative, some later suggestions reference elements initialized from earlier suggestions, so when coding, we consider the parts of each suggestion that differ from one another.

As shown in Table 5, MN-LLM generates more, longer, and higher quality suggestion combinations than the zero-shot baseline. MN-LLM does not always generate the 33 story options afforded by the initial story and combinations of five suggestions, as sometimes no generated suggestions passed our validation steps, or all generated fixes were too long. However, it still generates more than 20 options on average, in comparison to zero-shot’s average of 3.8.

**7.3.4 Workflows Generate a Wider Variety of Suggestions.** Without knowing the condition, we coded the story suggestions in the MN-LLM and both zero-shot conditions by theme. MN-LLM suggestions were more varied in type than the zero-shot suggestions (Table 6). The MN-LLM validation steps removed nearly half of the generated suggestions and most often voted for adding plot, characters, or emotional details as the final suggestions.

**7.3.5 The MN-LLM Workflow Improves Imagery, Originality, and Style.** Following the evaluation methods of the MN paper, for each story prompt we gave the output from MN-LLM and the zero-shot baseline, in randomized order, to 200 crowdworkers on Amazon Mechanical Turk using

Table 6. MN-LLM and Zero-Shot Generate Suggestion Types for Story Elements

Suggestion Type	Zero-Shot	MN-LLM	Passed	Voted
Add plot events	25	30	13	6
Introduce character	14	18	13	3
Introduce significant object	6	1	1	0
Add plot twist	4	11	10	3
Add detail to interactions	1	6	3	1
Add emotional detail	0	17	8	3
Add background details	0	8	5	0
Add a different perspective	0	5	2	2
Add a moral lesson	0	2	1	1
Add story detail	0	2	1	0

The MN-LLM workflow first generates a suggestion, which must pass a series of verification steps. The workflow then votes on valid suggestions to decide which to incorporate. At all stages, MN-LLM provides a wider variety of story suggestions than zero-shot.

EasyTurk [84]. We asked crowdworkers to vote between stories from the SoyLent-LLM and zero-shot baselines along seven story writing dimensions [77]. We paid crowdworkers \$1.50 USD, the equivalent of \$18 per hour.

The MN workflow improved the quality of the output narratives, although sometimes at the expense of a concise plot. We find that MN-LLM stories were more highly rated along the imagery, originality, and style, and overall dimensions. Zero-shot-generated stories gained more votes in the coherency, plot, and technical dimensions (Table 4).

**7.3.6 Workflows Can Struggle with Coherence.** Without knowing the condition, we also reviewed the MN-LLM and zero-shot conditions for plot irregularities. Long-form plot coherence of creative LLM outputs is difficult to achieve [121, 166], and we find some plot irregularities in the MN-LLM outputs. Three stories revealed the same information multiple times. Two plots had disjointed elements, one in which the same pronoun “she” referred to two characters in confusing ways and one with a time and location skip due to a scene that was never selected for editing. Such plot irregularities occur in the MN-LLM outputs because of the iteration through scenes and the use of a story summary as context. Both of these design choices were made to avoid exceeding LLM context limits.

**7.3.7 Workflows Have Higher Costs.** We measure the time and number of API calls required to run the workflow compared to running the baselines. Compared to zero-shot (mean: 21.88 seconds, SD: 5.89), MN-LLM takes much longer in seconds to complete (mean: 1,452.97 seconds, SD: 930.26). Compared to zero-shot’s one call, MN-LLM makes 2,398 LLM calls on average. The crowdsourced MN paper does not report latency or the number of crowdworkers [77].

## 8 Discussion

The approach of decomposing a task into a series of subtasks is conceptually the same for both crowdsourcing workflows and LLM chains. Through overlap in the literature and our case study observations, we find that the same *objectives* and *tactics* apply to both, with the exception of the choice of actors. We also find that the *strategies* found in crowdsourcing apply to chaining, as do their impact on *objectives*. Despite these similarities, LLMs and people have different abilities,

strengths, and weaknesses, so we sometimes need different approaches for employing *tactics* to effectively implement the same *strategies*.

Thus, *strategies and their effect on objectives transfer directly from crowdsourcing to LLMs. However, the tactic designs to support strategies can vary significantly and could benefit from improved development and testing processes.* Understanding how to effectively use *tactics* to implement *strategies* is not only a more focused goal than the connection between *tactics* and *objectives*, but it is also more generalizable. For example, learning how to effectively support LLMs in validating their own outputs would be applicable to many different tasks and objectives.

In this section, we discuss further the similarities and differences between LLM chains and crowdsourcing workflows, lessons that exploring this intersection taught us about supporting *strategies* with *tactics* in LLM chain design, and unknowns that mark promising avenues for future discovery.

### 8.1 LLM Stochasticity Is Different Than Human Diversity.

One of the key benefits of crowdsourcing workflows is that they incorporate multiple people, such that crowdworkers can fill in each other's deficits [69, 75, 126, 171]. Therefore, requesting multiple crowdworkers for the same or sequential subtasks inherently leads to differing outputs and considerations. The diversity of perspectives happens inherently by querying different actors but is limited to available crowdworkers and thus affected by demographic and geographic disparities on crowdsourcing platforms [40].

LLMs also have the capability to output different responses, as their training data include information produced by many people on the Internet [17]. To induce diversity, prior work has tried methods like adjusting prompts, temperature, and model types [15, 43, 129, 130]. In some cases, one can automatically calculate the optimal mix of LLMs for a task [113]. Our case studies suggest that LLMs need more variation in prompting and explicit intervention to get *diverse responses*, although we did not test this systematically. Intuitively, this makes sense because there is great diversity in the human population, but models are known to mimic the majority voice to minimize their objectives, ignoring marginalized opinions that make up a smaller percentage of the training data [52].

Using a high temperature and prompting the LLM with a variety of different instructions can improve diversity somewhat but is unlikely to match that of crowds. Although this difference between actors may at some level seem obvious, it impacts the tactics needed to effectively use diverse responses. Diverse responses are used both to source variety and to check agreement. In crowdsourcing, the underlying tactic behind both of these goals is identical: run the same subtask with multiple actors. However, different tactics better explicitly address these applications for LLMs.

*Variety.* Diverse responses are often used to support tasks with a variety of subjective outputs. Consider one step in MN that generates suggestions for updating the story. As this is an iterative workflow, we want a variety of story update suggestions both at each iteration and across iterations. Unfortunately, even with variety in generation prompt wording, we found a low diversity of suggestion content (Table 6 “Zero-Shot” column). After repeated challenges in prompt engineering to increase the variety of suggestions, we eventually intervened with an added *validation* step of multiple *evaluate* subtasks to ensure diversity of story updates. This approach was effective, but inefficient, as it filters out many *generated* responses. Role-based *communicative* architectures are also promising in giving prompts extra power to increase variety. Although giving identity based roles should be treated with caution [152], giving other roles to LLMs could be a more powerful way of inducing a variety of responses [130]. Many roles in

crowdsourcing workflows were based on organizational positions for relevant domains such as journalism [3] or project development [150].

To source variety, the best choice of *redundant* architecture can differ when using crowdworkers or LLMs. Prior crowdsourcing work has found that parallel and iterative architectures source different degrees of diversity and quality of responses in crowdsourcing depending on the task [10, 53, 82, 104]. For example, iterative architectures have less variation and a higher average quality of outputs but have a lower maximum quality output [104]. Some of these tradeoffs will hold with LLMs, but differences may emerge as well. For instance, some work suggests that, relative to crowdworkers who experience anchoring bias on iterative architectures, LLMs can perform divergent thinking to generate something explicitly different than provided examples [158, 161]. Similar studies on LLM chains would better inform the effect of different redundant architectures on variety and quality.

*Agreement.* Crowdsourcing work has used agreement among crowdworkers as a proxy for confidence vs. uncertainty [27, 101]. This agreement can be used to *validate* the quality of outputs, so understanding how agreement relates to quality is of critical importance. Here, we focus on agreement among outputs of identical tasks, but the difference between a different person validating the outputs of another and a model sequentially validating its own outputs should also be considered.

There are many potential methods for measuring confidence for LLMs, including a confidence score, checking consistency across small prompt variations, and querying multiple times for taking the majority answer [43, 100, 129, 164]. The effect of each of these approaches on the quality of answers is not well understood [43, 100, 129, 164]. It is unclear the degree to which just the consistency in one LLM’s outputs is a proxy for confidence and accuracy, vs. consistency across outputs from varied prompts or models.

This more tenuous connection between agreement and correctness may have some impact on chosen tactics. For example, *communicative* architectures in crowdsourcing were often based on the premise that agreement among people mattered, but that may not transfer to LLMs. Consider debate-based *communicative* architectures. Some crowdsourcing studies only bring up debate if there’s disagreement [27, 28, 42], but LLMs may benefit more from inducing debate for all inputs, rather than if two outputs disagree.

Therefore, the differences between human diversity and LLM stochasticity impact the design of *redundant* and *communicative* architectures to support divergent thinking or agreement.

## 8.2 LLMs Validate Response Quality Less Reliably Than Crowdworkers

*Validation* in crowdsourcing workflows often involves asking a different actor to verify if errors or desirable qualities exist in outputs. However, both the literature and our case studies show mixed results on the effectiveness of this strategy for chaining.

While some studies show that prompting a model to reflect can help achieve higher accuracy outputs for *objective* tasks [51, 114, 164, 168], others caution against the same model validating its own outputs, as LLMs exhibit high confidence for incorrect prior outputs [43, 100]. Some approaches for more robust validation include: validating with a different LLM [100], debate-based *communicative* structures [43, 100], and structural validation such as similarity across parallel *redundant* streams [101]. One paper found that the “ease of persuasion,” or how quickly an LLM relented in a debate, was more indicative of factual accuracy than other measures of confidence [43].

Through our case studies, we found several helpful practices for adapting validation steps to work with LLMs.

First, validating “quality” is open-ended, so it is helpful to specify aspects of quality to check instead of using one general “quality” validation step. Crowdsourced Soylent already split the validation into separately checking for “grammar” and for “meaning” changes [19]. However, we needed to make these checks even more targeted to adapt to LLM-specific flaws. Crowdsourced Soylent did not need to check that the rewritten text was shorter, but we added a length-checking step to Soylent-LLM after struggling to ensure a shorter length through prompting alone. We also explicitly validated that the model did not change direct quotes, a nuance that is understood in crowdworkers’ common sense. However, the subphrases from Soylent-LLM were for the most part fluent, deprioritizing the importance of the grammaticality check. In another example, crowdsourced MN involves one *evaluate* step in which crowdworkers choose the “best” suggestion. We added three validation checks to MN-LLM before this step to account for common errors in chosen suggestions. We explicitly check that LLM responses (1) do not overlap with prior suggestions, (2) are not already implemented in the story, and (3) are worded as a suggestion rather than as new text. Although we did not specify any in our case studies, LLMs may require additional validation of output structures, as text-based responses cannot be restricted by interface design the way crowdworker interfaces can.

Validation should address specific elements of quality that are aligned with actor capabilities. In our case studies, we determined what elements of quality to validate by evaluating intermediate outputs in a trial and error process. Determining these qualities can be tricky, especially as the capabilities of LLMs rapidly change. To determine what elements of quality require validation, intuition needs to be tempered by empirical testing and subtask evaluation.

Even after the set of qualities under consideration has been determined, we found LLMs could still struggle with successfully validating these qualities. To adapt to these issues, we used several approaches. First, when possible we made *evaluate* steps programmatic rather than LLM-based. This approach is most reliable, but there are a limited number of domains where programmatic checks are feasible. Second, we continued to break qualities down into even simpler parts when needed to *align subtask* difficulty with actor capabilities. Third, we *simplify context* to the minimum needed. Considering the story suggestion quality checks above, we found that including multiple prior suggestions or the entire story when not necessary reduced the quality of validation checks.

Some types of validation remain difficult for LLMs. Anecdotally, we found that some more abstract validation tasks, such as identifying changes in meaning, were challenging to enact with LLMs. In Crowdsourced Soylent, the “Verify” step evaluates if the fix changes the meaning of the statement. This requires nuance for what elements of the meaning are necessary. Overall, supporting specificity and simplicity in validation steps were key for adapting these workflows to use LLMs.

### 8.3 Chains for Subjective Tasks Can Use Crowdsourcing Approaches to Enable Users to Navigate the Space of Outputs

Some tasks have a large *subjective* space of possible correct outputs, due to uncertainty or creativity, and could benefit from the user giving direct input. Under these conditions, it is helpful to draw on longstanding work on mixed-initiative user interfaces to incorporate *users* into the workflow [62]. Such an approach allows for parts of the workflow to be automated while others involve *user* input.

By enabling user initiative on workflow direction or outputs, users can navigate the space of outputs—either to explicitly take control or to learn through the process of exploration. For example, this could be helpful for determining classification label sets. Instead of pre-defining all possible rules in instructions, workflows can surface uncertain areas for users to make final decisions [27]. In more creative ventures like prototyping, incorporating user feedback to align the output to their vision can avoid technically correct but irrelevant outputs [87].

User control through prompting is brittle [15], so workflows can enable more fine grained control. LLM chains can draw from user interaction methods from crowdsourcing workflows. Several large types of user interaction emerge from the literature.

First, the user may act as a step in the workflow for validation or for clarifying context when triggered by ties [133] or crowdworkers [86, 137]. The main design decisions for this approach are deciding what and when the user will validate. The structure and timing of user interaction can matter, for example when supplying actors with needed context [137].

Second, the user and crowdworkers communicate through a shared interface in a relatively unstructured cycle of interactions while creating a shared creative output. Hurdles here include making user review of outputs *easier* and reducing *latency* [92, 126]. Metaphorian is a chaining paper that follows a structured version of this in which users work on an interface to develop metaphors, but are limited to a few pre-defined actions [78].

Third, a user may review options between iterations of the workflow [29, 121]. This approach can help with exploring design spaces [87, 108, 130] and determining when the final quality threshold is met [157].

Fourth, a user may directly manipulate the output to see the breadth of user responses and make final decisions about the output with less effort than doing the task directly [19, 27, 29, 122]. This approach enables users to surface uncertainty in the data and express control over the final output.

Many chaining papers discuss a generic possibility of a user replacing steps or intervening if needed [139, 153, 163, 166], but few actually implement this with people [78, 121, 130]. We can see from crowdsourcing that there are many ways of incorporating the user that bring myriad benefits. In our case studies, we take direct manipulation from crowdsourcing and apply it to chaining, increasing the controllability of LLM outputs.

#### 8.4 On-Demand Availability of LLMs Reduces Barriers to Real-Time Crowdsourcing Architectures

Unlike crowdworkers, LLMs are available to be queried at all times with low *latency*. While, in one sense, this observation is obvious, it has deep implications on tactics.

The unpredictable availability of crowdworkers presented challenges for workflows in which crowdworkers must stay available for extended periods of time or must be called on-demand [18]. While incentive schemes could provide some real-time response, incentive tuning and implementation added significant complexity [93, 94, 156]. As a result, several crowdsourced workflows—while achieving impressive results on paper—were too complex for widespread deployment [28, 150]. The availability of LLMs reduces the barriers to these workflow designs and their benefits.

Methods like retainer models that incentivize keeping crowds were important for crowdsourcing workflows that involve real-time interaction for users [18] and for *communicative* workflows like debate in which actors must remain online as other actors respond to them [28, 42]. These workflows improve *latency* for users and improve quality and *accuracy* for objective tasks, respectively, but the challenges of retaining crowdworkers increased barriers in crowdsourcing. Similarly, for mixed-initiative interfaces, delayed crowdworker responses increased the *latency* of interactions [92]. Reduced latency of worker responses can make user interactions more real-time.

In other workflows, such as Flash Organizations [150], crowdworkers recruit other crowdworkers to complete subtasks that are dynamically created. Finding the right crowdworkers for the job took a median of 14 minutes with a light screening process and 15 hours with a more intense screening process [150]. In contrast, the Orchestrator LLM of the AutoGen [158] chain “recruits” with other calls to LLMs, which can happen flexibly and quickly [158].

Therefore, using LLMs as actors reduces barriers to using powerful *communicative* workflow regimes.

### 8.5 LLMs Need Different, Often Simpler, Subtasks

Throughout our case studies, we often needed to break crowdsourcing *subtasks* into multiple, more specified *subtasks*. Multiple factors necessitated these changes: natural language is difficult to structure (e.g., thus, we query for one item category at a time rather than eight in Cascade), LLMs have difficulty balancing multiple requirements in a subtask [63, 114, 161] (e.g., in response, we split Soylent’s “Fix” step into separate edit, merge, and delete queries), and LLMs do not understand length-related constraints (e.g., to address this, we add a hard-coded *validation* step after a *generate* subtask in MN to check for length).

In our case studies, LLMs performed better on highly specified subtasks for chains rather than on fewer, more complicated subtasks. However, this may change as newer, more powerful LLMs are introduced.

### 8.6 Instructions Should Provide the Minimum Necessary Context

Providing global *context* to subtasks is widely agreed to be a critical element of subtask design and output coherence, but it can increase instruction complexity [9, 111, 137, 159]. Crowdsourcing research more broadly has studied the process of improving worker instruction, including by ensuring comprehension through tests [106] and communicating misconceptions to requesters [21]. However, findings from the literature survey and case studies suggest that LLMs can respond poorly to long, complicated instructions [43, 114, 129, 161]. This context limitation should be re-evaluated over time as LLM context windows grow.

The nuances of context inclusion appeared in our case studies. For example, we wanted to provide MN-LLM the entire short story thus far as context, as occurs in crowdsourced MN. To generate *diverse* suggestions for improvement we gave the LLM previous suggestions as context for what it should avoid. However, adding this requirement to the overall prompt anecdotally increased the output’s chance of anchoring to prior suggestions. Therefore, we made a separate evaluate *subtask* that determined if the new suggestion matched previous suggestions out of context from the rest of the story.

We recommend designers take into consideration that LLMs may not adequately process all the nuances of complicated instructions. Simplifying the subtask size in turn simplifies instructions and, often, the amount of necessary context.

### 8.7 Chains Can Be Useful Despite Increased Cost

Although chains bring many improvements over zero-shot performance in our case studies, they also incur sizeable increases in expense and latency. Without thoughtful design, AI support tools can make easy tasks easier and hard tasks harder [144]. When should one choose between chaining or a zero-shot approach for potential applications? We consider the effects of cognitive load, the degree of user agency, the interruptibility of the overall workflow, and limitations in model capabilities.

The increased effort and latency of a chain’s runtime can be offset by quicker and less effortful interpretation of the outputs. For instance, in writing editing, it is easier to review highlighted minor edits rather than entirely new large blocks of text.

Chains may be helpful, despite higher costs, when the goal is spurring users’ innovation, learning, and design, rather than giving responses outright. For example, Metaphorian allows users to create scientific metaphors in an interactive interface [78]. The goal of such a system is not only to output the metaphors but also to be exploratory and spur thinking. Compared to zero-shot creation of metaphors, the paper finds users retain a greater sense of agency, metaphors increase in originality and coherence, and user satisfaction is higher when using an interactive chain.

The impact of increased latency on the utility of chains lessens when the broader workflow can be interrupted without major disruption. Chains that do not require user guidance, or only require it in spaced-out intervals, can run in the background while the user performs other tasks. Some tasks, such as in-line coding assistance, are not conducive to this type of interruption. However, others, such as short-story generation, can incorporate a long runtime, thus reducing the impact of longer latency on the user experience.

Finally, another reason to use chains despite the cost is if the task cannot be completed without a chain. At the moment, long-form writing and generating large code repositories are outside the capabilities of models and require a chain. It is worth noting that these capabilities will change over time. Thus, thoughtfulness about objectives and the broader task context can motivate the use of chains when a higher utility offsets increased costs.

## 9 Future Opportunities for Chaining

In this section, we highlight practices in crowdsourcing that we consider likely to be worthwhile to transfer to chaining. We encourage future work on these topics.

### 9.1 Subjective Tasks with Dynamic and Communicative Architectures

Crowdsourcing studies currently support a wider range of tasks for *subjective* work with a wider range of approaches than chaining studies. We expect these approaches to transfer to chaining and bring myriad benefits.

Future work could attempt *subjective* tasks that are under-explored in chaining, such as: prototyping [87, 95, 150], categorizing [22, 27, 70], topic concept maps [60, 107, 111], project development [115], or more specific tasks like itinerary creation [171], researching a purchase [82], and conceptual blends [33].

Crowdsourcing work also uses a broader range of *communicative* and *dynamic* architectures to support *subjective* work. Giving actors roles during *communication* can assist both crowdsourcing workflows [3, 136, 150] and LLM chains [98, 158]. In crowdsourcing work, some studies go beyond giving a role in instructions and instead create free-form role-based [126, 150, 172], hierarchical [3, 136], and state-sharing [92, 93, 95, 126] *communicative* architectures.

*Dynamic* architectures provide more flexibility when contingencies arise [75, 82, 86, 150] and adapt to edge cases [28, 86, 124]. LLM chains rarely implement *dynamic* architectures for *subjective* creative work [78, 161, 168, 173]. Crowdsourcing work applies *dynamic* architectures to a wider variety of *subjective* creative tasks [33, 82, 92, 136] and includes more complex dynamic workflows such as multi-level map-reduce [82] and dynamically choosing between workflows using their maximum expected value of utility [53]. Decision theoretic techniques, such as Partially Observable Markov Decision Processes, yield great efficiencies in crowdsourcing [38, 39], but have thus far barely been considered for LLM chaining [113].

We recommend pursuing a broader range of *subjective* creative tasks with more *communicative* (e.g., establishing a hierarchy of roles) and *dynamic* (e.g., multi-level map-reduce) architectures.

### 9.2 Integrating the User

Although many chaining papers discuss involving humans into the chain hypothetically [139, 153, 158, 163, 166], few implement this user oversight [145, 149]. Our case studies showcase one example of interactions through directly manipulable outputs, but the varied ways of incorporating the user in crowdsourcing workflows could continue to inspire similar approaches in chaining.

Explorable interfaces surface disagreements to the user rather than resolving them in the workflow itself [27, 29, 122, 157]. Actor-initiated guidance allows actors to define when to alert the user, focusing user attention only on what matters to the actors [86, 137]. In *subjective* tasks, quick

iteration and choosing among uncertain options reduces the effort it takes for the user to prototype creative outputs and allows the user to guide ideation [92, 126]. For example, unstructured text can be interpreted in many ways depending on an analyst's research questions, so recent work such as Concept Induction allows analysts to steer the process of LLM-aided text analysis [90].

We can also consider how our case study examples could be expanded to incorporate the user differently. With inspiration from the crowdsourcing workflow Revolt [27], a user of Cascade-LLM could review edge cases that only match one category label of two that were merged. Soylent-LLM could replace the verify step with a user, as the user may recognize subtle undesirable changes in fixed phrases that are difficult to pre-define through prompts. MN-LLM could implement the *quality thresholds* strategy by having the user review after each iteration if the generation process should continue. The user could also give more oversight on the edit suggestions chosen. Incorporating the user in more ways reduces *expense* by calling LLMs fewer times and by directing the calls of LLMs that do occur toward more important needs. However, it does require more *effort* on the user's part and could increase *latency* if the user is slower than the equivalent call to an LLM.

We recommend research into the opportunities and limitations of user guidance, especially for subjective tasks, as initial findings conflict on if using generative AI tools leads to divergent thinking [2] or design fixation [151].

### 9.3 Studying Chains More Systematically

More systematic studies on optimal chain design would be valuable. Several studies already approach questions around model choice [43, 113], incorporating roles into prompts [130, 158], and using self-reflection or debate to validate outputs [43, 100]. Although the outcomes may be different, systematic crowdsourcing studies could provide inspiration for some of these explorations. Prior crowdsourcing studies investigate the effect of context timing, architecture, and content on task output [111, 137]. Others measure tradeoffs of iterative and parallel workflows (e.g., effect on latency [111], quality averages and variance [53, 82, 104], risk of cascading errors [53]) in various task contexts, as well as good practices for each (e.g., appropriate number of branches or iterations [10, 70]).

By specifying a strategy of interest, the design space can orient studies for understanding optimal tactical design decisions. For example, studying debate designs could look at debate to establish *adaptability* for tasks with many edge cases, debate to *validate* outputs, or debate to encourage a *diversity* of creative ideas for a *subjective* task. Evaluating the use of tactic designs within the context of specific strategies will lead to more generalizable findings than evaluating solely on task performance.

We recommend systematic tests for tactical design decisions informed by prior crowdsourcing work, objectives, and strategies.

### 9.4 LLM Workflow Design Processes Are Unstructured and Design Support Is Needed

Since LLMs are cheaper and faster to query [148], there are more opportunities to find optimal designs through systematic testing. While the lower barrier to entry may allow LLM chain designers to iterate more quickly, it doesn't guarantee that they focus their attention on the right things while iterating. Since LLMs display unpredictable outputs in response to low-level changes, designers may end up fixating on minor tweaks like prompt rewordings rather than exploring the design space of strategies and tactics. To avoid potential design fixation that tunnels into technical detail rather than exploring higher-level design directions, LLM chain designers would benefit from tools that scaffold their design and iteration process. Although running an LLM chain is cheaper and faster than running a crowdsourcing workflow, the design process is unstructured. There is a need for design support.

We see many promising directions to help build chains efficiently and effectively, given that iterations are often based on extensive qualitative analysis and enlightened trial-and-error.

As there are many overlapping tactics among workflows, the flexibility and reusability of modules can support the development of future workflows [73]. A repository of “battle-tested” workflow operators could allow for sharing of quality parts for reuse in new workflows. Such a platform could include operator-level test suites to assist with validation that would monitor operator output quality in new content areas or with newly released LLMs. The platform design could draw on crowdsourcing to create such test suites. Corresponding workflow design tools could suggest different options to achieve a goal by collecting multiple approaches for sub-parts of the workflow.

Other work engineers more optimal LLM prompts by iteratively generating and scoring test prompt candidates using mathematical or natural language optimization functions [165, 177]. Future work could adapt these approaches within the context of chaining, for example to optimize for diverse responses.

Comparing many workflow designs can be difficult because of the requirement that the output of a previous step match the input of the next step. Therefore, removing or adding steps can be a non-trivial process. More formalized input and output indicators, potentially with matching prompts, could facilitate easier chain development. This approximates the approach of Flash Teams which documents the inputs and outputs of steps [136].

Sketching aids practitioners in rapidly instantiating ideas to explore the possible solution space rather than committing too early to one design direction [26, 41, 167]. Incorporating sketching could aid LLM chain designers in arriving at better solutions, rather than getting stuck in local optima. Sketches are quick, but they also should be low-fidelity and minimal in detail. While LLMs afford quick instantiation, they produce high-fidelity output that can distract designers from their primary prototyping goals [72, 89]. Opportunities to support this designing perspective include providing interfaces amenable to rapid chain construction and reconfiguration [14, 159] and suggesting strategies and tactics when a user may be tunneling on technical details.

We recommend continued work on chain design support tools such as repositories of workflow operators, prompting optimization methods for chaining, formalized input-output indicators, and sketching-inspired workflow design approaches.

## 10 Limitations

Our work addresses crowdsourcing workflows and LLM chains, but it is limited in scope beyond these fields and in discussion of broader impacts. Additionally, our recommendations are contingent upon an imperfect understanding of LLM capabilities, and we support the utility of our case studies through theory instead of a user study.

*Scope.* Our work targets LLM chain design by adapting crowdsourcing workflow techniques. However, there are opportunities to look beyond LLMs to multi-modal models and other tools as the workers [59, 109, 146] and beyond crowdsourcing to other task decomposition methods as a source of inspiration [12, 13, 116].

Additionally, there may be limitations to the scope of work covered by our review of the LLM chaining literature. Since LLM chaining is a relatively new area, researchers have not converged upon established terminology as in the crowdsourcing domain. Our keyword search may not have captured some work that did not use one of our keyword search terms.

*Broader Impacts.* We focus on the ability of workflows to serve a beneficial role in supporting various tasks. However, there are potential risks in using LLMs and crowdworkers as support tools that we do not explicitly consider. Several papers in the design space cite concerns around a loss of jobs and unintended consequences of using LLMs, especially for creative tasks [121, 130,

158]. Crowdsourcing work similarly raised concerns around replacing expert workers and around labor rights for crowdworkers [81]. Broader discussions about privacy [132, 158], ownership [78, 121], and perceived sense of agency we also leave to future work. Finally, although workflows can improve output quality, results will still be imperfect, so errors from implemented chains could still have harmful impacts [46, 55].

*Rapid Changes of LLM Capabilities.* The capabilities of LLMs are rapidly changing, so researchers' understanding of and subsequent recommendations to work with these capabilities may fall out of sync with new model developments. We recommend future work to monitor the efficacy of chain designs as new models are released, with potentially differing capabilities and failure modes.

*Lack of User Evaluation.* We build on longstanding HCI theory and practice that demonstrates the benefit of direct manipulation interfaces [71, 143]. However, our work does not conduct a user evaluation to explicitly test the benefits and drawbacks of our direct manipulation case study implementations. Future work may more deeply explore how users interact with LLM chains via direct manipulation and if these interactions positively impact user control and task outcomes.

## 11 Conclusion

By revisiting the established field of crowdsourcing workflows, we find inspiration and direction for the nascent area of LLM chaining. To facilitate adaptation, we contribute a design space for LLM chain designers that incorporates findings from the crowdsourcing workflow and LLM chaining literatures. This design space covers the objectives supported by and the tactical building blocks of prior workflows and chains. We also surface strategies employed by designers to use tactics to support objectives. Through three case studies, we build an understanding of how crowdsourcing workflows can be adapted for LLM chains.

Finally, we synthesize our findings to discuss tactical design decisions for LLM chains through the context of their comparison with crowdsourcing workflows. We find LLM chains need additional support in subtask design and validation due in part to differences between LLM stochasticity and human diversity. On the other hand, the comparatively low cost and latency of LLM calls allow for faster iteration and for workflow techniques that were studied but never practical with human crowds. We hope these findings guide future work toward more approaches to user guidance and subjective tasks, the systematic discovery of good practices for tactic design, and artifacts to support the chain design process.

LLM chains present remarkable flexibility, which comes along with both promising capabilities but also substantial complexity. Our work helps LLM chain designers navigate this vast design space by building on the established field of crowdsourcing.

## Acknowledgments

We thank Michael Bernstein and Joy Kim for providing outputs from the original Soylent and MN papers. Thank you to Clement Zheng for your input on creating corpus tables similar to prior work [16]. We also thank the Allen Institute for Artificial Intelligence for access to Amazon Mechanical Turk resources. We thank OpenAI for offering research credits to the GPT-3.5 API.

## References

- [1] LangChain. 2024. Build a retrieval augmented generation (RAG) app. Retrieved from <https://python.langchain.com/v0.2/docs/tutorials/rag/>
- [2] Chiou Li-Yuan, Hung Peng-Kai, Liang Rung-Huei, and Wang Chun-Teng. 2023. Designing with AI: An exploration of co-ideation with image generators. In *Proceedings of the ACM Designing Interactive Systems Conference (DIS '23)*. ACM, New York, NY, USA, 1941–1954.

- [3] Elena Agapie, Jaime Teevan, and Andrés Monroy-Hernández. 2015. Crowdsourcing in the field: A case study using local crowds for event reporting. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 3, 2–11.
- [4] Salman Ahmad, Alexis Battle, Zahan Malkani, and Sepander Kamvar. 2011. The Jabberwocky programming environment for structured social computing. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 53–64.
- [5] Christopher Alexander. 1977. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
- [6] Garrett Allen, Gaole He, and Ujwal Gadiraju. 2023. Power-up! What Can Generative Models Do for Human Computation Workflows?
- [7] Abdullah Alshaibani, Sylvia Carrell, Li-Hsin Tseng, Jungmin Shin, and Alexander Quinn. 2020. Privacy-preserving face redaction using crowdsourcing. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 8, 13–22.
- [8] Abdullah Alshaibani and Alexander J. Quinn. 2021. Pterodactyl: Two-step redaction of images for robust face deidentification. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 9, 27–34.
- [9] Vamshi Ambati, Stephan Vogel, and Jaime Carbonell. 2012. Collaborative workflow for crowdsourcing translation. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 1191–1194.
- [10] Paul André, Robert E. Kraut, and Aniket Kittur. 2014. Effects of simultaneous and sequential work structures on distributed collaborative interdependent tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 139–148.
- [11] Paul André, Haoqi Zhang, Juho Kim, Lydia Chilton, Steven Dow, and Robert Miller. 2013. Community clustering: Leveraging an academic crowd to form coherent conference sessions. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 1, 9–16.
- [12] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 39–48.
- [13] Obinna Anya. 2015. Bridge the gap! What can work design in crowdwork learn from work design theories?. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 612–627.
- [14] Ian Arawjo, Priyan Vaithilingam, Martin Wattenberg, and Elena Glassman. 2023. ChainForge: An open-source visual programming environment for prompt engineering. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 1–3.
- [15] Simran Arora, Avaniika Narayan, Mayee F. Chen, Laurel Orr, Neel Guha, Kush Bhatia, Ines Chami, and Christopher Re. 2023. Ask me anything: A simple strategy for prompting language models. In *Proceedings of the 11th International Conference on Learning Representations*. Retrieved from <https://openreview.net/forum?id=bhUPJnS2g0X>
- [16] S. Sandra Bae, Clement Zheng, Mary Etta West, Ellen Yi-Luen Do, Samuel Huron, and Danielle Albers Szaifir. 2022. Making data tangible: A cross-disciplinary design space for data physicalization. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, New York, NY, USA, Article 81, 18 pages. DOI: <https://doi.org/10.1145/3491102.3501939>
- [17] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.
- [18] Michael S. Bernstein, Joel Brandt, Robert C. Miller, and David R. Karger. 2011. Crowds in two seconds: Enabling realtime crowd-powered interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 33–42.
- [19] Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A word processor with a crowd inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, 313–322.
- [20] Danielle Bragg, Naomi Caselli, John W. Gallagher, Miriam Goldberg, Courtney J. Oka, and William Thies. 2021. ASL sea battle: Gamifying sign language data collection. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–13.
- [21] Jonathan Bragg, Mausam, and Daniel S. Weld. 2018. Sprout: Crowd-powered task design for crowdsourcing. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. Retrieved from <https://api.semanticscholar.org/corpusid:51948377>
- [22] Jonathan Bragg and Daniel Weld. 2013. Crowdsourcing multi-label classification for taxonomy creation. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 1, 25–33.
- [23] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101.

- [24] Victor S. Burszty, David Demeter, Doug Downey, and Larry Birnbaum. 2022. Learning to perform complex tasks through compositional fine-tuning of language models. arXiv:2210.12607. Retrieved from <https://doi.org/10.48550/arXiv.2210.12607>
- [25] Crystal Butler, Stephanie Michalowicz, Lakshmi Subramanian, and Winslow Burleson. 2017. More than a feeling: The MiFace framework for defining facial communication mappings. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, 773–786.
- [26] Bill Buxton. 2010. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann.
- [27] Joseph Chee Chang, Saleema Amershi, and Ece Kamar. 2017. Revolt: Collaborative crowdsourcing for labeling machine learning datasets. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2334–2346.
- [28] Quanze Chen, Jonathan Bragg, Lydia B. Chilton, and Dan S. Weld. 2019. Cicero: Multi-turn, contextual argumentation for accurate crowdsourcing. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–14.
- [29] Justin Cheng and Michael S. Bernstein. 2015. Flock: Hybrid crowd-machine learning classifiers. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 600–611.
- [30] Justin Cheng, Jaime Teevan, Shamsi T. Iqbal, and Michael S. Bernstein. 2015. Break it down: A comparison of macro-and microtasks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 4061–4064.
- [31] Lydia B. Chilton, Juho Kim, Paul André, Felicia Cordeiro, James A. Landay, Daniel S. Weld, Steven P. Dow, Robert C. Miller, and Haoqi Zhang. 2014. Frenzy: Collaborative data organization for creating conference sessions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1255–1264.
- [32] Lydia B. Chilton, Greg Little, Darren Edge, Daniel S. Weld, and James A. Landay. 2013. Cascade: Crowdsourcing taxonomy creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1999–2008.
- [33] Lydia B. Chilton, Savvas Petridis, and Maneesh Agrawala. 2019. VisiBlends: A flexible workflow for visual blends. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–14.
- [34] Maria Christoforaki and Panagiotis Ipeirotis. 2014. Step: A scalable testing and evaluation platform. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 2, 41–49.
- [35] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. arXiv:2110.14168. Retrieved from <https://doi.org/10.48550/arXiv.2110.14168>
- [36] Antonia Creswell and Murray Shanahan. 2022. Faithful reasoning using large language models. arXiv:2208.14271. Retrieved from <https://doi.org/10.48550/arXiv.2208.14271>
- [37] Antonia Creswell, Murray Shanahan, and Irina Higgins. 2022. Selection-inference: Exploiting large language models for interpretable logical reasoning. arXiv:2205.09712. Retrieved from <https://doi.org/10.48550/arXiv.2205.09712>
- [38] Peng Dai, Christopher H. Lin, and Daniel S. Weld. 2013. POMDP-based control of workflows for crowdsourcing. *Artificial Intelligence* 202 (2013), 52–85.
- [39] Peng Dai and Daniel Weld. 2010. Decision-theoretic control of crowd-sourced workflows. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 24, 1168–1174.
- [40] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and dynamics of mechanical Turk workers. In *Proceedings of the 11th ACM International Conference on Web Search and Data Mining*, 135–143.
- [41] Steven P. Dow, Alana Glassco, Jonathan Kass, Melissa Schwarz, Daniel L. Schwartz, and Scott R. Klemmer. 2011. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Transactions on Computer-Human Interaction* 17, 4, Article 18 (Dec. 2011), 24 pages. DOI : <https://doi.org/10.1145/1879831.1879836>
- [42] Ryan Drapeau, Lydia Chilton, Jonathan Bragg, and Daniel Weld. 2016. Microtalk: Using argumentation to improve crowdsourcing accuracy. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 4, 32–41.
- [43] Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. Improving factuality and reasoning in language models through multiagent debate. In *ICML 2024 Conference*. OpenReview. net. Retrieved from <https://openreview.net/forum?id=zj7YuTE4t8>
- [44] Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jian, Bill Yuchen Lin, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena D. Hwang, et al. 2023. Faith and fate: Limits of transformers on compositionality. In *37th Conference on Neural Information Processing Systems (NeurIPS 2023)*. Retrieved from [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/deb3c28192f979302c157cb653c15e90-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/deb3c28192f979302c157cb653c15e90-Paper-Conference.pdf)
- [45] Nouha Dziri, Sivan Milton, Mo Yu, Osmar R. Zaiane, and Siva Reddy. 2022. On the origin of hallucinations in conversational models: Is it the datasets or the models? arXiv:2204.07931. Retrieved from <https://doi.org/10.48550/arXiv.2204.07931>

- [46] Florin Eggmann, Roland Weiger, NicolaU Zitzmann, and MarkusB Blatz. 2023. Implications of large language models such as chatgpt for dental medicine. *Journal of Esthetic and Restorative Dentistry : official Publication of the American Academy of Esthetic Dentistry .. [Et Al.]* 35, 7 (2023), 1098–1102. DOI : <https://doi.org/10.1111/jerd.13046> <https://doi.org/37017291>
- [47] Jacob Eisenstein, Daniel Andor, Bernd Bohnet, Michael Collins, and David Mimno. 2022. Honest students from untrusted teachers: Learning an interpretable question-answering pipeline from a pretrained language model. arXiv:2210.02498. Retrieved from <https://doi.org/10.48550/arXiv.2210.02498>
- [48] Suhana Bedi, Scott L. Fleming, Chia-ChunChiang, Keith Morse, Aswathi M. Kumar, Birju Patel, Jenelle A. Jindal, Conor Davenport, Craig Yamaguchi, and Nigam H. Shah. 2024. QUEST-AI: A system for question generation, verification, and refinement using AI for USMLE-Style Exams. *medRxiv* 2023–04. DOI : <https://doi.org/10.1101/2023.04.25.23288588>
- [49] Suzanne Fricke. 2018. Semantic scholar. *Journal of the Medical Library Association* 106, 1 (2018), 145.
- [50] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. 1994. *Design Patterns Elements of Reusable Object-Oriented Software*. Addison-Wesley.
- [51] Zelalem Gero, Chandan Singh, Hao Cheng, Tristan Naumann, Michel Galley, Jianfeng Gao, and Hoifung Poon. 2023. Self-verification improves few-shot clinical information extraction. arXiv:2306.00024. Retrieved from <https://doi.org/10.48550/arXiv.2306.00024>
- [52] Mitchell L. Gordon, Michelle S. Lam, Joon Sung Park, Kayur Patel, Jeff Hancock, Tatsunori Hashimoto, and Michael S. Bernstein. 2022. Jury learning: Integrating dissenting voices into machine learning models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, New York, NY, USA, Article 115, 19 pages. DOI : <https://doi.org/10.1145/3491102.3502004>
- [53] Shinsuke Goto, Toru Ishida, and Donghui Lin. 2016. Understanding crowdsourcing workflow: Modeling and optimizing iterative and parallel processes. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 4. 52–58.
- [54] Sai Gouravajhala, Jinyeong Yim, Karthik Desingh, Yanda Huang, Odest Jenkins, and Walter Lasecki. 2018. Eureka: Enhanced understanding of real environments via crowd assistance. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 6, 31–40.
- [55] Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. Not what you’ve signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. In *AISeC ’23: Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, 79–90. DOI : <https://doi.org/10.1145/3605764.3623985>
- [56] Anhong Guo, Xiang’Anthony’ Chen, Haoran Qi, Samuel White, Suman Ghosh, Chieko Asakawa, and Jeffrey P. Bigham. 2016. Vizlens: A robust and interactive screen reader for interfaces in the real world. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, 651–664.
- [57] Anhong Guo, Jeeun Kim, Xiang’Anthony’ Chen, Tom Yeh, Scott E. Hudson, Jennifer Mankoff, and Jeffrey P. Bigham. 2017. Facade: Auto-generating tactile interfaces to appliances. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 5826–5838.
- [58] Anhong Guo, Junhan Kong, Michael Rivera, Frank F. Xu, and Jeffrey P. Bigham. 2019. Statelens: A reverse engineering solution for making existing dynamic touchscreens accessible. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, 371–385.
- [59] Tanmay Gupta and Aniruddha Kembhavi. 2023. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14953–14962.
- [60] Nathan Hahn, Joseph Chang, Ji Eun Kim, and Aniket Kittur. 2016. The knowledge accelerator: Big picture thinking in small pieces. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2258–2270.
- [61] Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. Surface form competition: Why the highest probability answer isn’t always right. arXiv:2104.08315. Retrieved from <https://doi.org/10.48550/arXiv.2104.08315>
- [62] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 159–166.
- [63] Cheng-Yu Hsieh, Si-An Chen, Chun-Liang Li, Yasuhisa Fujii, Alexander Ratner, Chen-Yu Lee, Ranjay Krishna, and Tomas Pfister. 2023. Tool documentation enables zero-shot tool-usage with large language models. arXiv:2308.00675. Retrieved from <https://doi.org/10.48550/arXiv.2308.00675>
- [64] Qing Huang, Yishun Wu, Zhenchang Xing, He Jiang, Yu Cheng, and Huan Jin. 2023. Adaptive intellect unleashed: The feasibility of knowledge transfer in large language models. arXiv:2308.04788. Retrieved from <https://doi.org/10.48550/arXiv.2308.04788>
- [65] Qing Huang, Jiahui Zhu, Zhilong Li, Zhenchang Xing, Changjing Wang, and Xiwei Xu. 2023. PCR-chain: Partial code reuse assisted by hierarchical chaining of prompts on frozen copilot. In *Proceedings of the 2023 IEEE/ACM 45th International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*. IEEE, 1–5.

- [66] Qing Huang, Zhou Zou, Zhenchang Xing, Zhenkang Zuo, Xiwei Xu, and Qinghua Lu. 2023. AI chain on large language model for unsupervised control flow graph generation for statically-typed partial code. arXiv:2306.00757. Retrieved from <https://doi.org/10.48550/arXiv.2306.00757>
- [67] Ting-Hao Huang, Joseph Chee Chang, and Jeffrey P. Bigham. 2018. Evorus: A crowd-powered conversational assistant built to automate itself over time. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–13.
- [68] Ting-Hao Huang, Walter Lasecki, and Jeffrey Bigham. 2015. Guardian: A crowd-powered spoken dialog system for web APIs. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 3, 62–71.
- [69] Yi-Ching Huang, Jiunn-Chia Huang, Hao-Chuan Wang, and Jane Hsu. 2017. Supporting ESL writing by prompting crowdsourced structural feedback. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 5, 71–78.
- [70] Jordan S. Huffaker, Jonathan K. Kummerfeld, Walter S. Lasecki, and Mark S. Ackerman. 2020. Crowdsourced detection of emotionally manipulative language. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14.
- [71] Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. 1985. Direct manipulation interfaces. *Human-Computer Interaction* 1, 4 (1985), 311–338.
- [72] Ellen Jiang, Kristen Olson, Edwin Toh, Alejandra Molina, Aaron Donsbach, Michael Terry, and Carrie J. Cai. 2022. PromptMaker: Prompt-based prototyping with large language models. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. ACM, New York, NY, USA, Article 35, 8 pages. DOI: <https://doi.org/10.1145/3491101.3503564>
- [73] Martin Josifoski, Lars Klein, Maxime Peyrard, Yifei Li, Saibo Geng, Julian Paul Schnitzler, Yuxing Yao, Jiheng Wei, Debjit Paul, and Robert West. 2023. Flows: Building blocks of reasoning and collaborating AI. arXiv:2308.01285. Retrieved from <https://doi.org/10.48550/arXiv.2308.01285>
- [74] Harmanpreet Kaur, Mitchell Gordon, Yiwei Yang, Jeffrey Bigham, Jaime Teevan, Ece Kamar, and Walter Lasecki. 2017. Crowdmask: Using crowds to preserve privacy in crowd-powered systems via progressive filtering. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 5, 89–98.
- [75] Joy Kim and Andres Monroy-Hernandez. 2016. Stora: Summarizing social media content based on narrative theory using crowdsourcing. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, 1018–1027.
- [76] Juho Kim, Phu Tran Nguyen, Sarah Weir, Philip J. Guo, Robert C. Miller, and Krzysztof Z. Gajos. 2014. Crowdsourcing step-by-step information extraction to enhance existing how-to videos. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 4017–4026.
- [77] Joy Kim, Sarah Sterman, Allegra Argent Beal Cohen, and Michael S Bernstein. 2017. Mechanical novel: Crowdsourcing complex work through reflection and revision. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 233–245.
- [78] Jeongyeon Kim, Sangho Suh, Lydia B. Chilton, and Haijun Xia. 2023. Metaphorian: Leveraging large language models to support extended metaphor creation for science writing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, 115–135.
- [79] Tae Soo Kim, Yoonjoo Lee, Minsuk Chang, and Juho Kim. 2023. Cells, generators, and lenses: Design framework for object-oriented interaction with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 1–18.
- [80] Aniket Kittur, Susheel Khamkar, Paul André, and Robert Kraut. 2012. CrowdWeaver: Visually managing complex crowd work. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 1033–1036.
- [81] Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, 1301–1318.
- [82] Aniket Kittur, Boris Smus, Susheel Khamkar, and Robert E. Kraut. 2011. Crowdforge: Crowdsourcing complex work. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 43–52.
- [83] Masaki Kobayashi, Hiromi Morita, Masaki Matsubara, Nobuyuki Shimizu, and Atsuyuki Morishima. 2018. An empirical study on short- and long-term effects of self-correction in crowdsourced microtasks. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 6, 79–87.
- [84] Ranjay Krishna. 2019. EasyTurk: A wrapper for custom AMT tasks. Retrieved from <https://github.com/ranjaykrishna/easyturk>
- [85] Alex Krizhevsky et al. 2009. Learning multiple layers of features from tiny images.
- [86] Anand Kulkarni, Matthew Can, and Björn Hartmann. 2012. Collaboratively crowdsourcing workflows with turko-matic. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 1003–1012.

- [87] Anand Kulkarni, Prayag Narula, David Rolnitzky, and Nathan Kontny. 2014. Wish: Amplifying creative ability with expert crowds. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 2, 112–120.
- [88] Vivian Lai, Chacha Chen, Alison Smith-Renner, Q. Vera Liao, and Chenhao Tan. 2023. Towards a science of human-AI decision making: An overview of design space in empirical human-subject studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1369–1385.
- [89] Michelle S. Lam, Zixian Ma, Anne Li, Izequiel Freitas, Dakuo Wang, James A. Landay, and Michael S. Bernstein. 2023. Model sketching: Centering concepts in early-stage machine learning model design. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. ACM, New York, NY, USA, Article 741, 24 pages. DOI : <https://doi.org/10.1145/3544548.3581290>
- [90] Michelle S. Lam, Janice Teoh, James A. Landay, Jeffrey Heer, and Michael S. Bernstein. 2024. Concept induction: Analyzing unstructured text with high-level concepts using LLoM. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–28.
- [91] LangChain. 2022. LangChain. Retrieved from <https://www.langchain.com/>
- [92] Walter S. Lasecki, Juho Kim, Nick Rafter, Onkur Sen, Jeffrey P. Bigham, and Michael S. Bernstein. 2015. Apparition: Crowdsourced user interfaces that come to life as you sketch them. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 1925–1934.
- [93] Walter S. Lasecki, Rachel Wesley, Jeffrey Nichols, Anand Kulkarni, James F. Allen, and Jeffrey P. Bigham. 2013. Chorus: A crowd-powered conversational assistant. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, 151–162.
- [94] Thomas D. LaToza, W. Ben Towne, Christian M. Adriano, and André Van Der Hoek. 2014. Microtask programming: Building software with a crowd. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, 43–54.
- [95] Sang Won Lee, Yujin Zhang, Isabelle Wong, Yiwei Yang, Stephanie D. O’Keefe, and Walter S. Lasecki. 2017. Sketch-express: Remixing animations for more effective crowd-powered prototyping of interactive interfaces. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, 817–828.
- [96] Yoav Levine, Itay Dalmedigos, Ori Ram, Yoel Zeldes, Daniel Jannai, Dor Muhlgay, Yoni Osin, Opher Lieber, Barak Lenz, Shai Shalev-Shwartz, et al. 2022. Standing on the shoulders of giant frozen language models. arXiv:2204.10019. Retrieved from <https://doi.org/10.48550/arXiv.2204.10019>
- [97] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. 2023. Teach LLMs to personalize—An approach inspired by writing education. arXiv:2308.07968. Retrieved from <https://doi.org/10.48550/arXiv.2308.07968>
- [98] Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2024. CAMEL: Communicative agents for “mind” exploration of large language model society. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 2264, 18 pages.
- [99] Jiyi Li. 2024. A comparative study on annotation quality of crowdsourcing and LLM via label aggregation. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. DOI : <https://doi.org/10.1109/ICASSP48485.2024.10447803>
- [100] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Zhaopeng Tu, and Shuming Shi. 2023. Encouraging divergent thinking in large language models through multi-agent debate. arXiv:2305.19118. Retrieved from <https://doi.org/10.48550/arXiv.2305.19118>
- [101] Beatrice Liem, Haoqi Zhang, and Yiling Chen. 2011. An iterative dual pathway structure for speech-to-text transcription. In *Proceedings of the 11th AAAI Conference on Human Computation*. AAAI Press, 37–42.
- [102] Chi-Chin Lin, Yi-Ching Huang, and Jane Yung-jen Hsu. 2014. Crowdsourced explanations for humorous internet memes based on linguistic theories. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 2, 143–150.
- [103] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Proceedings of the 13th European Conference on Computer Vision (ECCV '24), Part V 13*. Springer, 740–755.
- [104] Greg Little, Lydia B. Chilton, Max Goldman, and Robert C. Miller. 2010. Exploring iterative and parallel human computation processes. In *Proceedings of the ACM SIGKDD Workshop on Human Computation*, 68–76.
- [105] Greg Little, Lydia B. Chilton, Max Goldman, and Robert C. Miller. 2010. TurkIt: Human computation algorithms on mechanical Turk. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, 57–66.
- [106] Angli Liu, Stephen Soderland, Jonathan Bragg, C. H. Lin, Xiao Ling, and Daniel S. Weld. 2016. Effective crowd annotation for relation extraction. In *North American Chapter of the Association for Computational Linguistics*. Retrieved from <https://api.semanticscholar.org/CorpusID:10705630>

- [107] Ching Liu, Juho Kim, and Hao-Chuan Wang. 2018. ConceptScape: Collaborative concept mapping for video learning. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–12.
- [108] Ryan Liu, Howard Yen, Raja Marjeh, Thomas L. Griffiths, and Ranjay Krishna. 2023. Improving interpersonal communication by simulating audiences with language models. arXiv:2311.00687. Retrieved from <https://doi.org/10.48550/arXiv.2311.00687>
- [109] Vivian Liu, Han Qiao, and Lydia Chilton. 2022. Opal: Multimodal image generation for news illustration. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, 1–17.
- [110] Tao Long, Dorothy Zhang, Grace Li, Batool Taraif, Samia Menon, Kynneddy Simone Smith, Sitong Wang, Katy Ilonka Gero, and Lydia B Chilton. 2023. Tweetorial hooks: Generative AI tools to motivate science on social media. arXiv:2305.12265. Retrieved from <https://doi.org/10.48550/arXiv.2305.12265>
- [111] Kurt Luther, Nathan Hahn, Steven Dow, and Aniket Kittur. 2015. Crowdlines: Supporting synthesis of diverse information sources through crowdsourced outlines. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 3, 110–119.
- [112] Stephen MacNeil, Andrew Tran, Joanne Kim, Ziheng Huang, Seth Bernstein, and Dan Mogil. 2023. Prompt middleware: Mapping prompts for large language models to UI affordances. arXiv:2307.01142. Retrieved from <https://doi.org/10.48550/arXiv.2307.01142>
- [113] Pranjal Aggarwal, AmanMadaan, Ankit Anand, Srividya Pranavi Potharaju, Swaroop Mishra, Pei Zhou, Aditya Gupta, Dheeraj Rajagopal, Karthik Kappaganthu, Yiming Yang, Shyam Upadhyay, Manaal Faruqui and Mausam. 2024. AutoMix: Automatically Mixing Language Models. In *Advances in Neural Information Processing Systems*. A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.). Curran Associates, Inc, 131000–131034. Retrieved from [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/ecda225cb187b40ea8edc1f46b03ffda-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/ecda225cb187b40ea8edc1f46b03ffda-Paper-Conference.pdf)
- [114] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2024. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, Vol. 36, 46534–46594.
- [115] Narges Mahyar, Michael R. James, Michelle M. Ng, Reginald A. Wu, and Steven P. Dow. 2018. CommunityCrit: Inviting the public to improve and evaluate urban design ideas through micro-activities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–14.
- [116] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. 2019. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. arXiv:1904.12584. Retrieved from <https://doi.org/10.48550/arXiv.1904.12584>
- [117] Adam Marcus, Eugene Wu, David Karger, Samuel Madden, and Robert Miller. 2011. Human-powered sorts and joins. In *Proceedings of the 38th International Conference on Very Large Data Bases*.
- [118] Adam Marcus, Eugene Wu, David R. Karger, Samuel Madden, and Robert C. Miller. 2011. Demonstration of Qurk: A query processor for humanoperators. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, 1315–1318.
- [119] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. arXiv:2005.00661.
- [120] Tyler McDonnell, Matthew Lease, Mucahid Kutlu, and Tamer Elsayed. 2016. Why is that relevant? collecting annotator rationales for relevance judgments. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 4, 139–148. Retrieved from <https://doi.org/10.48550/arXiv.2005.00661>
- [121] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-writing screenplays and theatre scripts with language models: evaluation by industry professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–34.
- [122] Vikram Mohanty, Kareem Abdol-Hamid, Courtney Ebersohl, and Kurt Luther. 2019. Second opinion: Supporting last-mile person identification with crowdsourcing and face recognition. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7, 86–96.
- [123] Varun Nair, Elliot Schumacher, and Anitha Kannan. 2023. Generating medically-accurate summaries of patient-provider dialogue: A multi-stage approach using large language models. arXiv:2305.05982. Retrieved from <https://doi.org/10.48550/arXiv.2305.05982>
- [124] Varun Nair, Elliot Schumacher, Geoffrey Tso, and Anitha Kannan. 2023. DERA: Enhancing large language model completions with dialog-enabled resolving agents. arXiv:2303.17071. Retrieved from <https://doi.org/10.48550/arXiv.2303.17071>
- [125] Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. arXiv:1808.08745. Retrieved from <https://doi.org/10.48550/arXiv.1808.08745>

- [126] Michael Nebeling, Alexandra To, Anhong Guo, Adrian A. de Freitas, Jaime Teevan, Steven P. Dow, and Jeffrey P. Bigham. 2016. WearWrite: Crowd-assisted writing from smartwatches. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 3834–3846.
- [127] Feng Nie, Meixi Chen, Zhirui Zhang, and Xu Cheng. 2022. Improving few-shot performance of language models via nearest neighbor calibration. arXiv:2212.02216. Retrieved from <https://doi.org/10.48550/arXiv.2212.02216>
- [128] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, Vol. 35, 27730–27744.
- [129] Aditya G. Parameswaran, Shreya Shankar, Parth Asawa, Naman Jain, and Yujie Wang. 2023. Revisiting prompt engineering via declarative crowdsourcing. arXiv:2308.03854. Retrieved from <https://doi.org/10.48550/arXiv.2308.03854>
- [130] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, 1–18.
- [131] Amy Pavel, Colorado Reed, Björn Hartmann, and Maneesh Agrawala. 2014. Video digests: A browsable, skimmable format for informational lecture videos. In *Proceedings of the UIST*, Vol. 10. Citeseer, 2642918–2647400.
- [132] Charith Peris, Christophe Dupuy, Jimit Majmudar, Rahul Parikh, Sami Smaili, Richard Zemel, and Rahul Gupta. 2023. Privacy in the time of language models. In *Proceedings of the 16th ACM International Conference on Web Search and Data Mining*, 1291–1292.
- [133] Mary Pietrowicz, Danish Chopra, Amin Sadeghi, Puneet Chandra, Brian Bailey, and Karrie Karahalios. 2013. Crowd-Band: An Automated Crowdsourcing Sound Composition System. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 1, 121–129.
- [134] Ariadna Quattoni and Antonio Torralba. 2009. Recognizing indoor scenes. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 413–420.
- [135] Justin Reppert, Ben Rachbach, Charlie George, Luke Stebbing Jungwon Byun, Maggie Appleton, and Andreas Stuhlmüller. 2023. Iterated Decomposition: Improving Science Q&A by Supervising Reasoning Processes. arXiv:2301.01751. Retrieved from <https://doi.org/10.48550/arXiv.2301.01751>
- [136] Daniela Retelny, Sébastien Robaszkiewicz, Alexandra To, Walter S. Lasecki, Jay Patel, Negar Rahmati, Tulsee Doshi, Melissa Valentine, and Michael S. Bernstein. 2014. Expert crowdsourcing with flash teams. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, 75–85.
- [137] Niloufar Salehi, Jaime Teevan, Shamsi Iqbal, and Ece Kamar. 2017. Communicating context to the crowd for complex writing tasks. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 1890–1901.
- [138] Elliot Salisbury, Ece Kamar, and Meredith Morris. 2017. Toward scalable social alt text: Conversational crowdsourcing as a tool for refining vision-to-language technology for the blind. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 5, 147–156.
- [139] Timo Schick, Jane Dwivedi-Yu, Zhengbao Jiang, Fabio Petroni, Patrick Lewis, Gautier Izacard, Qingfei You, Christoforos Nalmpantis, Edouard Grave, and Sebastian Riedel. 2022. Peer: A collaborative language model. arXiv:2208.11663. Retrieved from <https://doi.org/10.48550/arXiv.2208.11663>
- [140] Shreya Shankar, Aditya G. Parameswaran, and Eugene Wu. 2024. DocETL: Agentic query rewriting and evaluation for complex document processing. arXiv:2410.12189. Retrieved from <https://doi.org/10.48550/arXiv.2410.12189>
- [141] Rulin Shao, Jacqueline He, Akari Asai, Weijia Shi, Tim Dettmers, Sewon Min, Luke Zettlemoyer, and Pang Wei Koh. 2024. Scaling retrieval-based language models with a trillion-token datastore. In *38th Conference on Neural Information Processing Systems (NeurIPS 2024)*. Retrieved from [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/a5d8aba27dfef4e849e8cb03fb87a954-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/a5d8aba27dfef4e849e8cb03fb87a954-Paper-Conference.pdf)
- [142] Yang Shi, Xingyu Lan, Jingwen Li, Zhaorui Li, and Nan Cao. 2021. Communicating with motion: A design space for animated visual narratives in data videos. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–13.
- [143] Ben Shneiderman. 1982. The future of interactive systems and the emergence of direct manipulation. *Behaviour & Information Technology* 1, 3 (1982), 237–256.
- [144] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2025. Ironies of generative AI: Understanding and mitigating productivity loss in human-AI interactions. *International Journal of Human-Computer Interaction* 41, 5 (2025), 2898–2919. DOI: <https://doi.org/10.1080/10447318.2024.2405782>
- [145] Otilia Stretcu, Edward Vendrow, Kenji Hata, Krishnamurthy Viswanathan, Vittorio Ferrari, Sasan Tavakkol, Wenlei Zhou, Aditya Avinash, Emming Luo, Neil Gordon Alldrin, et al. 2023. Agile modeling: From concept to classifier in minutes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22323–22334.

- [146] Didac Suris, Sachit Menon, and Carl Vondrick. 2023. Vipergpt: Visual inference via python execution for reasoning. Retrieved from [https://openaccess.thecvf.com/content/ICCV2023/papers/Suris\\_ViperGPT\\_Visual\\_Inference\\_via\\_Python\\_Execution\\_for\\_Reasoning\\_ICCV\\_2023\\_paper.pdf](https://openaccess.thecvf.com/content/ICCV2023/papers/Suris_ViperGPT_Visual_Inference_via_Python_Execution_for_Reasoning_ICCV_2023_paper.pdf)
- [147] Jaime Teevan, Shamsi T. Iqbal, and Curtis Von Vech. 2016. Supporting collaborative writing with microtasks. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2657–2668.
- [148] Christian Terwiesch and Karl Ulrich. 2023. M.B.A. Students vs. ChatGPT: Who comes up with more innovative ideas? *The Wall Street Journal*. Retrieved from <https://www.wsj.com/tech/ai/mba-students-vs-chatgpt-innovation-679edf3b>
- [149] Imad Eddine Toubal, Aditya Avinash, Neil Gordon Alldrin, Jan Dlabal, Wenlei Zhou, Enming Luo, Otilia Stretcu, Hao Xiong, Chun-Ta Lu, Howard Zhou, et al. 2024. Modeling collaborator: Enabling subjective vision classification with minimal human effort via LLM tool-use. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17553–17563.
- [150] Melissa A. Valentine, Daniela Retelny, Alexandra To, Negar Rahmati, Tulsee Doshi, and Michael S. Bernstein. 2017. Flash organizations: Crowdsourcing complex work by structuring crowds as organizations. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 3523–3537.
- [151] 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 *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, Article 380, Pages 1–18. DOI: <https://doi.org/10.1145/3613904.3642919>
- [152] Angelina Wang, Jamie Morgenstern, and John P. Dickerson. 2024. Large language models cannot replace human participants because they cannot portray identity groups. arXiv:2402.01908. Retrieved from <https://doi.org/10.48550/arXiv.2402.01908>
- [153] Boshi Wang, Xiang Deng, and Huan Sun. 2022. Iteratively prompt pre-trained language models for chain of thought. arXiv:2203.08383. Retrieved from <https://doi.org/10.48550/arXiv.2203.08383>
- [154] Sitong Wang, Savvas Petridis, Taeahn Kwon, Xiaojuan Ma, and Lydia B. Chilton. 2023. PopBlends: Strategies for conceptual blending with large language models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19.
- [155] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, Vol. 35, 24824–24837.
- [156] Sarah Weir, Juho Kim, Krzysztof Z. Gajos, and Robert C. Miller. 2015. Learnersourcing subgoal labels for how-to videos. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work Social Computing*. ACM, New York, NY, 405–416. DOI: <https://doi.org/10.1145/2675133.2675219>
- [157] Wesley Willett, Jeffrey Heer, and Maneesh Agrawala. 2012. Strategies for crowdsourcing social data analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 227–236.
- [158] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. AutoGen: Enabling next-gen LLM applications via multi-agent conversation framework. arXiv:2308.08155. Retrieved from <https://doi.org/10.48550/arXiv.2308.08155>
- [159] Tongshuang Wu, Ellen Jiang, Aaron Donsbach, Jeff Gray, Alejandra Molina, Michael Terry, and Carrie J. Cai. 2022. Prompt chainer: Chaining large language model prompts through visual programming. In *Proceedings of the CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–10.
- [160] Tongshuang Wu, Michael Terry, and Carrie Jun Cai. 2022. AI chains: Transparent and controllable human-AI interaction by chaining large language model prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–22.
- [161] Tongshuang Wu, Haiyi Zhu, Maya Albayrak, Alexis Axon, Amanda Bertsch, Wenxing Deng, Ziqi Ding, Bill Guo, Sireesh Gururaja, Tzu-Sheng Kuo, et al. 2023. LLMs as workers in human-computational algorithms? Replicating crowdsourcing pipelines with LLMs. arXiv:2307.10168.
- [162] Tongshuang Sherry Wu, Michael Terry, and Carrie J. Cai. 2021. AI Chains: Transparent and controllable human-AI interaction by chaining large language model prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. DOI: <https://doi.org/10.1145/3491102.3517582>
- [163] Yiran Wu, Feiran Jia, Shaokun Zhang, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2024. MathChat: Converse to tackle challenging math problems with LLM agents. In *Proceedings of the ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- [164] Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, and Michael Qizhe Xie. 2023. Self-evaluation guided beam search for reasoning. In *37th Conference on Neural Information Processing Systems (NeurIPS 2023)*. Retrieved from [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/81fde95c4dc79188a69ce5b24d63010b-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/81fde95c4dc79188a69ce5b24d63010b-Paper-Conference.pdf)
- [165] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. 2023. Large language models as optimizers. arXiv:2309.03409. Retrieved from <https://doi.org/10.48550/arXiv.2309.03409>

- [166] Kevin Yang, Nanyun Peng, Yuandong Tian, and Dan Klein. 2022. Re3: Generating longer stories with recursive reprompting and revision. arXiv:2210.06774. Retrieved from <https://doi.org/10.48550/arXiv.2210.06774>
- [167] Qian Yang, Justin Cranshaw, Saleema Amershi, Shamsi T. Iqbal, and Jaime Teevan. 2019. Sketching NLP: A case study of exploring the right things to design with language intelligence. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 1–12. DOI: <https://doi.org/10.1145/3290605.3300415>
- [168] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, Vol. 36, 11809–11822.
- [169] Omar Zaidan and Chris Callison-Burch. 2011. Crowdsourcing translation: Professional quality from non-professionals. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies*, 1220–1229.
- [170] J. D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: How non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–21.
- [171] Haoqi Zhang, Edith Law, Rob Miller, Krzysztof Gajos, David Parkes, and Eric Horvitz. 2012. Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 217–226.
- [172] Jieyu Zhang, Ranjay Krishna, Ahmed H. Awadallah, and Chi Wang. 2023. EcoAssistant: Using LLM assistant more affordably and accurately. arXiv:2310.03046. Retrieved from <https://doi.org/10.48550/arXiv.2310.03046>
- [173] Zheng Zhang, Jie Gao, Ranjodh Singh Dhaliwal, and Toby Jia-Jun Li. 2023. VISAR: A human-AI argumentative writing assistant with visual programming and rapid draft prototyping. In *UIST '23: Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, Article 5, 1–30. DOI: <https://doi.org/10.1145/3586183.3606800>
- [174] Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. Retrieved from <https://www.chatgptthero.io/wp-content/uploads/2023/12/2210.03493.pdf>
- [175] Diliara Zharikova, Daniel Kornev, Fedor Ignatov, Maxim Talimanchuk, Dmitry Evseev, Ksenya Petukhova, Veronika Smilga, Dmitry Karpov, Yana Shishkina, Dmitry Kosenko, et al. 2023. DeepPavlov dream: Platform for building generative AI assistants. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Vol. 3: System Demonstrations*, 599–607.
- [176] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. arXiv:2205.10625. Retrieved from <https://doi.org/10.48550/arXiv.2205.10625>
- [177] Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitit, Harris Chan, and Jimmy Ba. 2023. Large language models are human-level prompt engineers. In *ICLR. OpenReview. net*. Retrieved from <https://openreview.net/forum?id=92gvk82DE->

## A Paper

To show the design space's completeness in representing the existing literature, we revisited each paper and categorized it within the design space. We describe the nuances of that process here.

We record the year of publication, such that if we collected a paper on arXiv in 2023 and it was published in a conference in 2024, we record the year as 2024.

We show the quality–cost objective spectrum as four columns, in order to show the increasing and decreasing of quality and cost. A paper may both increase and decrease quality or cost at once (e.g., the output is higher quality, but there are reported quality issues with cascading errors). We count the ability to do a task at all as an increase in quality and an equal match with baseline capabilities as neither an increase nor a decrease. Another challenge for comparison is that the baseline differs across papers. For instance, some papers compare quality to expert outputs, others to novice crowdworkers, and others to zero-shot LLM outputs. We label changes in quality and cost as they are discussed in the paper, rather than use a consistent cross-paper standard. For instance, if the chain could logically be expected to be more expensive than a single call to an LLM, but this is not explicitly discussed in the paper, we do not mark it as increasing costs. We take this approach to avoid speculation and to show trends along with what effects are considered worth reporting.

We used our discretion in classifying papers as only *subjective*, if the main focus of the task is highly subjective, such as with creative writing. There may still be *objective* elements, such as

wanting a generated story to be grammatically correct. This subjectivity in the categorization itself is why we express it as a spectrum in the design space.

In the design space text, we mention some automated elements that perform similar functions (e.g., rule-based heuristics for *evaluate*, sourcing outside material for *generate*, and algorithmic *merges*). We include these in the text descriptions to reference how workflows may use *actors* or automated elements to complete tasks. However, we do not label a paper as including a subtask of these types if it is only deterministically automated.

Finally, if there are multiple workflows tested in the paper, we include takeaways and elements from all in the coding.

Received 3 May 2024; revised 30 October 2024; accepted 23 December 2024