

Human-centered NLP Fact-checking: Co-Designing with Fact-checkers using Matchmaking for AI

HOUJIANG LIU*, School of Information, University of Texas at Austin, USA

ANUBRATA DAS*, School of Information, University of Texas at Austin, USA

ALEXANDER BOLTZ*, Human Centered Design & Engineering, University of Washington, USA

DIDI ZHOU, School of Information, University of Texas at Austin, USA

DAISY PINAROC, School of Information, University of Texas at Austin, USA

MATTHEW LEASE, School of Information, University of Texas at Austin, USA

MIN KYUNG LEE, School of Information, University of Texas at Austin, USA

While many Natural Language Processing (NLP) techniques have been proposed for fact-checking, both academic research and fact-checking organizations report limited adoption of such NLP work due to poor alignment with fact-checker practices, values, and needs. To address this, we investigate a co-design method, *Matchmaking for AI*, to enable fact-checkers, designers, and NLP researchers to collaboratively identify what fact-checker needs should be addressed by technology, and to brainstorm ideas for potential solutions. Co-design sessions we conducted with 22 professional fact-checkers yielded a set of 11 design ideas that offer a “north star”, integrating fact-checker criteria into novel NLP design concepts. These concepts range from pre-bunking misinformation, efficient and personalized monitoring misinformation, proactively reducing fact-checker potential biases, and collaborative writing fact-check reports. Our work provides new insights into both human-centered fact-checking research and practice and AI co-design research.

CCS Concepts: • **Human-centered computing** → **HCI theory, concepts and models**; **Empirical studies in HCI**; *Computer supported cooperative work*.

Additional Key Words and Phrases: AI co-design, Natural language processing, Fact-checking, Misinformation, Disinformation

ACM Reference Format:

Houjiang Liu, Anubrata Das, Alexander Boltz, Didi Zhou, Daisy Pinaroc, Matthew Lease, and Min Kyung Lee. 2024. Human-centered NLP Fact-checking: Co-Designing with Fact-checkers using Matchmaking for AI. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2, Article 423 (November 2024), 44 pages. <https://doi.org/10.1145/3686962>

*Three authors contributed equally to this research.

Authors' Contact Information: Houjiang Liu, liu.ho@utexas.edu, School of Information, University of Texas at Austin, Austin, Texas, USA; Anubrata Das, anubrata.das@utexas.edu, School of Information, University of Texas at Austin, Austin, Texas, USA; Alexander Boltz, aboltz@uw.edu, Human Centered Design & Engineering, University of Washington, Seattle, Washington, USA; Didi Zhou, didizhou@utexas.edu, School of Information, University of Texas at Austin, Austin, Texas, USA; Daisy Pinaroc, drp2767@utexas.edu, School of Information, University of Texas at Austin, Austin, Texas, USA; Matthew Lease, ml@utexas.edu, School of Information, University of Texas at Austin, Austin, Texas, USA; Min Kyung Lee, minkyung.lee@austin.utexas.edu, School of Information, University of Texas at Austin, Austin, Texas, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2573-0142/2024/11-ART423

<https://doi.org/10.1145/3686962>

1 Introduction

The massive scale of intentional *disinformation* and inadvertent *misinformation* threaten modern society in many areas, such as law and order [142], politics [143], and healthcare [129]. To combat this, journalists and fact-checkers perform a key societal function by debunking fake news, online rumors, and conspiracy [36, 58]. Today there are over 400 teams of journalists and researchers in over 105 countries with 378 known, active fact-checking projects [127]. While such human fact-checking has proven to be effective in terms of various measures [101], fact-checking remains a largely manual affair today, limiting the scale of its effective reach and practical impact [89, 102].

To help address this scalability challenge, many Natural Language Processing (NLP) techniques and tools have been proposed for fact-checking, across tasks such as claim detection, evidence retrieval, and claim verification [1, 42, 134]. Beyond NLP, additional AI techniques have been proposed for fact-checking multi-modal content (images, videos, and audio) [2, 21, 88]. However, despite the clear need for greater scalability and this ever-growing body of AI work on fact-checking, to date there has been only limited adoption of AI-based methods and tools for fact-checking. As highlighted by both academic research [44, 62, 89] and fact-checking organization reports [8], this stems from insufficient alignment with fact-checker needs, practices, and values. In part, simplified problem formulations that permit full automation neglect important aspects of real-world fact-checking, which remains a highly complex process requiring subjective judgment and human knowledge to corroborate claims based on local contexts [44, 89]. Relatedly, AI-based fact-checking driven by technological challenges of cutting-edge AI neglects the need for design studies to better surface real fact-checker needs and translate these into human-centered AI tools responsive to these needs [32, 92].

In this work, we use co-design to engage fact-checkers directly in designing human-centered NLP tools to meet their needs. To limit the scope, we restrict our study to NLP technologies related to language-based fact-checking, omitting broader AI technologies related to multi-modal fact-checking. Fact-checkers, designers, and NLP researchers translate their needs through co-design into design ideas informed by state-of-the-art NLP techniques for fact-checking.

In particular, we propose *Matchmaking for AI* to conduct this co-design, extending Bly and Churchill [18]’s earlier *Matchmaking* concept to an AI co-design process. *Matchmaking for AI* distills AI techniques into AI probes—abstracted visual concepts and forms—that are easy for non-AI experts to understand and provide a step-by-step ideation process to map state-of-the-art AI techniques into user activities wherein those techniques may have the most impact. Prior work in human-centered AI has mainly incorporated stakeholder feedback in two ways: 1) enhancing existing design with human values [82, 131, 156, 157], and 2) brainstorming with stakeholders to translate their values into high-level design guidelines [50, 74, 123, 152]. However, there appears to be a potential gap – while numerous frameworks have been developed to integrate human values into AI development, brainstorming the “right” design concepts with stakeholders that are novel and feasible remains a challenge. Additionally, existing AI co-design practices for brainstorming ideas tend to be speculative and may not consistently focus on the practical feasibility of these ideas. Thus, *Matchmaking for AI* empowers stakeholders to co-design a suite of workable ideas that support their workflow.

Using *Matchmaking for AI*, we conducted co-design sessions with 22 professional fact-checkers, yielding 11 novel design ideas (to the best of our knowledge). Compared to existing NLP fact-checking that focused on claim selection and verification, our ideas point to a broader set of challenges across the fact-checking workflow and fact-checker goals. The ideas assist in information searching, processing, and writing tasks for efficient and personalized fact-checking; help fact-checkers proactively prepare for future misinformation; monitor their potential biases; and

support internal organization collaboration. In addition, by using AI probes to teach non-technical stakeholders “just enough” AI to collaborate effectively with specialists, *Matchmaking for AI* helped foster effective communications with AI researchers, to provide concrete design suggestions and relevant technical details, so that participants could more readily brainstorm feasible AI design ideas.

Our work contributes to the literature on human-centered fact-checking in Human-computer Interaction (HCI) and NLP by offering a suite of fact-checker-designed tool ideas that can guide future NLP-based fact-checking research. The ideas could lead to technology development with greater potential for adoption by professional fact-checkers in practice. Our work also contributes to the literature on AI co-design by offering the use case of *Matchmaking for AI* with fact-checkers, which facilitates co-design by fact-checkers, designers, and NLP researchers.

The remainder of this article is organized as follows. We first review key related work in Section 2. Then in Section 3, we describe Matchmaking for AI in greater detail, study protocol, recruitment of fact-checkers, and methods used to analyze co-design outcomes. In Section 4, we present 11 ideas co-designed by fact-checkers. Next, in Section 5, we reflect on its benefits, efficacy, limitations, and future improvements. Finally, we conclude with Section 6.

2 Related Work

2.1 Professional Fact-checking Today and Opportunities for Mixed-initiative Workflows

Fact-checking practices exhibit some common patterns and workflows across different organizations. For example, Graves [45] synthesizes a standard fact-checking process into five stages: 1) choosing claims to check, 2) contacting the speaker, 3) tracing false claims, 4) working with experts, and 5) writing the fact-check. Building on Graves [45]’s work, Micallef et al. [89] grounds a more detailed workflow by interviewing fact-checkers worldwide. Juneja and Mitra [62] explore both the human and technical infrastructure of fact-checking among different stakeholder groups (e.g., editors, fact-checkers, investigators and researchers, social media managers, and advocates). As reported by these studies, a key challenge is the difficulty for human fact-checkers in combating the massive scale of online misinformation.

In other work, a growing body of NLP research has sought to automate different aspects of fact-checking [47]. To date, most attention has been directed toward: 1) detecting claims and prioritizing claim checkworthiness [116, 117], 2) detecting previously fact-checked claims [66], and 3) assessing claims given textual evidence [113, 118]. However, as reported by recent academic studies [89, 102] and organization reports [8], most existing computational tools are still “fragmented” [89] and have “limited scope and use of custom solutions” [15, 89].

Two popular claim detection tools include *ClaimBuster*[51], an open-sourced API to extract claims from paragraphs, and *Full Fact Alpha*¹, which enables fact-checkers to search and filter checkable claims from news articles. Claim detection is also incorporated into speech recognition tools [1], social media monitoring tools², and tip-line services³. However, these tools struggle to correctly prioritize claims in alignment with human fact-checker’s definitions of “check-worthiness” [8, 114]. To investigate claims, fact-checkers use off-the-shelf search engines, such as Google and Bing, to surface related articles, but most topically relevant pages retrieved from these search engines usually do not contain useful content for fact-checking [8, 15]. For claim verification, existing tools include (but are not limited to): 1) verifying COVID-19 related claims based on official

¹<https://fullfact.org/about/ai/>

²<https://www.facebook.com/formedia/mjp/programs/third-party-fact-checking>

³<https://meedan.com/post/covid-19-whatsapp-bot-for-fact-checkers>

datasets [37]; 2) checking statistical macroeconomic claims by referring to government reports and national statistics [134]; and 3) checking scientific claims based on academic paper abstracts [139]. However, due to their limited scope, fact-checkers have not widely adopted these auto-verification tools.

As reported by empirical studies [44, 62, 89, 102], fact-checking remains a complex practice that inherently requires subjective judgment and human knowledge [8, 45, 89], encompassing a variety of tasks that remain beyond the capabilities of even state-of-the-art AI. For this reason, fact-checking tools appear most likely to impact practice if researchers and designers seek to develop assistive technologies that augment, enhance, and accelerate the work of professional fact-checkers, rather than seeking to completely automate fact-checking. We thus envision a *mixed-initiative process* [55]: fact-checking tasks are strategically divided into more fine-grained sub-tasks operated by human fact-checkers and AI [78, 92, 93]. Human-centered design and co-design are considered vital methodologies for combining the strengths of both human effort and AI to create complementary hybrid solutions [32, 33].

2.2 Human-Centered Design and Co-Design for AI

Engaging stakeholders during AI development and design has increasingly piqued scholars' attention among HCI and AI communities [33, 159]. Empirical studies from different AI application domains, such as AI-assisted public and clinical decision-making [125, 132, 150], online advertising [5], as well as automated fact-checking [62, 89, 102], reported that the lack of stakeholder engagement in AI development might produce fairness-related harm for stakeholders at risk of poor AI performance. Additionally, the implementation and design of AI tools might fail to meet real stakeholder needs because they do not adapt to the socio-technical context and norms within specialized domains [74].

To address this, various human-centered approaches — including frameworks, methods, and design materials — have been proposed to broaden stakeholder participation in AI design and development. For example, Zhu et al. [157] introduced Value-Sensitive Algorithm Design to create “intelligent socialization algorithms for WikiProjects in Wikipedia.” Lee et al. [82] developed the WeBuildAI framework to design algorithms for “on-demand food donation transportation services” that balance stakeholder values. Also, HCI scholars propose using *AI as a design material* [40], in order to integrate AI elements into a traditional UX design process. A variety of design materials have been created to pair participatory AI with non-expert AI stakeholders, including designers [131, 156] and end-users [76]. For example, Zhou et al. [156] designed a user journey canvas that was incorporated with an AI development process (data annotation, model training, and inference). Designers can use this canvas to specify different elements for an AI design idea. Like Zhou et al. [156], Subramonyam et al. [131] created data-persona cards, model API cards, and explainability design guidelines, to engage AI practitioners (designers and engineers) to co-create AI experiences. Kuo et al. [76] also mapped an AI development lifecycle into comic boards, which help front-line workers understand how AI is developed and share detailed suggestions for designing AI systems.

Note that the aforementioned human-centered design methods are commonly employed to enhance design details by incorporating human values, particularly when a design concept is already established. In contrast, co-design is used to assist stakeholders in brainstorming ideas in situations where the AI design landscape is under-explored, translating social and ethical values as design guidance or distilling high-level user needs. For example, it has been used to co-imagine “tech futures” for a Utopian city with Black youth [50]; to understand how children imagine a fair AI librarian would behave [123]; to uncover values to inform AI design for journalism [74]; to imagine new rideshare platforms that consider driver well-being [152]; and to raise awareness of algorithmic controls on social media feeds [128].

Sanders and Stappers [111] define co-design as a co-creation activity, involving expert users who actively contribute to brainstorming new design ideas and informing future product development. This suggests that these users possess some level of technical expertise. As they begin to use existing tools with an innovative and expert mindset, co-design can effectively integrate their insights into the creation of more useful tools [111, 137]. However, by reviewing recent AI co-design practices, we learn that as most stakeholders are not familiar with AI, AI co-design often becomes speculative, primarily serving as a method for human value solicitation, rather than yielding specific design decisions with concrete AI technology being used [33]. Thus, there appears a potential gap in the existing AI design paradigm – while numerous human-centered frameworks have been created for incorporating human values into AI development, identifying the “right” design concepts upfront, particularly those that are innovative yet feasible for stakeholder-centered AI design remains a challenge. Additionally, existing AI co-design practices for brainstorming ideas tend to be speculative and may not consistently focus on the practical feasibility of these ideas.

2.3 Design through Matchmaking

In this work, we seek to bridge the above gap by empowering stakeholders to brainstorm AI ideas that are both novel and feasible – fact-checkers and AI researchers have deeper-level conversations to explore appropriate design choices with a better adoption of existing AI techniques, so that these design ideas are practically useful for fact-checkers and more likely to be developed in the real world. In particular, we introduce *Matchmaking for AI*, extending Bly and Churchill [18]’s Matchmaking concept as an AI co-design process, which maps AI techniques to user activities wherein those techniques may have the most impact. Kensing and Greenbaum [68] point out that to foster effective co-design, mutual learning is a necessary step for different stakeholders by sharing and synthesizing their domain knowledge to brainstorm novel design solutions. Different from studies that use co-design in a speculative approach without considering technological capabilities upfront, Bly and Churchill implement matchmaking with a series of mutual learning steps (described in Section 3.1), which help non-technical experts identify “technology affordance” and map this affordance to their work activities that can be easily supported, thus producing feasible design choices.

Inspired by this original process and an adapted version developed by van Dijk and Zimmerman [135], our extended *Matchmaking for AI* aims to foster such learning experiences between stakeholders and AI researchers. For example, by better understanding AI techniques, stakeholders could be more knowledgeable to inform what and how AI could augment domain-specific tasks instead of emulating human work [121]. Meanwhile, AI researchers could become more creative to innovate domain-specific AI techniques that situate stakeholder needs [32, 92].

3 Study Design

As prior human-centered design helps produce sufficient design details and traditional co-design facilitates idea brainstorming, *Matchmaking for AI* leverages synergistic interactions between stakeholder expertise and their understanding of AI capabilities to brainstorm feasible AI ideas. Our process actively familiarizes fact-checkers with fine-grained NLP techniques (e.g., classification, clustering, and summarization). With this learning experience, fact-checkers can envision how NLP might be usefully and realistically applied in their workflow and brainstorm how they can collaboratively work with AI automation across different fact-checking tasks.

In Section 3.1, we first describe the idea of Matchmaking for AI, including its original concept and our adaptation. Next, Section 3.2 presents the study protocol and materials used in our workshops with fact-checkers. We then describe the recruitment in Section 3.3, followed by the qualitative coding method used to analyze the workshops in Section 3.4.

3.1 Matchmaking for AI

Bly and Churchill [18] propose “Matchmaking” as a co-design concept to “incorporate user domain knowledge into early design when a technology prototype already exists”. The original concept includes four steps: 1) describing technology capabilities; 2) mapping those capabilities to associated work activities; 3) identifying work domains and specific tasks; and 4) verifying whether these tasks match technology capabilities. For the 2nd step, van Dijk and Zimmerman [135] argue that the mapping results of technology capabilities vs. work activities are too vast to be usefully navigated. To address this, they suggest that beyond merely mapping capabilities to activities, designers should further assess the rate at which mappings would be most useful. While prior matchmaking activities helped to identify stakeholders who might benefit from the new technology, designers, and developers still took an active role. We argue that this does not exploit the advantage of co-design, which democratizes the existing power structures of product development by empowering end-users with more autonomy to express their design choices. Thus, our *Matchmaking for AI* inherits the original mapping activity for discovering design spaces where existing techniques can have the most impact but empowers stakeholders to do so by eliciting their expertise, passion, and creativity.

Matchmaking for AI seeks to create “workable ideas” that are both grounded in participant needs and values, as well as informed by their understanding of AI capabilities. This is expected to produce results that translate high-level themes (e.g., user needs or design guidelines) into clear design concepts that are more feasible for implementation. However, this does not mean that one could simply use off-the-shelf AI models for practical implementation. Instead, the goal is to identify “right” design concepts that would serve as a “north star” for designers and NLP researchers. To generate such outcomes, we tailor a three-step matchmaking process that includes the following five aspects (a-e): a) need assessment, b) specific gaps identified from the existing tools and related prior work in the literature, c) tool features desired by our participants, d) mapping to existing AI techniques, and e) potential technical challenges for implementation. We discuss each of these three steps below.

3.1.1 Mapping Stakeholder Domain Expertise and AI Needs (Step 1). A growing body of empirical work in human-centered AI has called for more extensive research on aligning AI capabilities with domain expertise [14, 20, 73, 150]. Such alignments help designers and AI practitioners explore appropriate design spaces before developing AI applications. Our *Matchmaking for AI* starts with capturing details of stakeholder domain expertise and their AI needs. Inspired by Hoffman’s work [54] on “eliciting knowledge from experts”, we ask stakeholders to reflect on two aspects of domain expertise: 1) at a high level, they formalize a domain-specific workflow with its sub-tasks; and 2) at a low level, they externalize desired outcomes, values, and criteria of their decision-making process in each sub-task. The high level view enables participants to find potential design spaces where AI automation will be most useful, while the low level informs participants about their requirements, needs, and concrete parameters for developing practical AI solutions (e.g., input-outputs, training objectives, evaluation criteria, etc.).

Prior empirical studies in the domain provide a valuable starting point to learn about a stakeholder group and its activities. However, in a co-design environment, greater contextual information (e.g., use cases and user behavior) enables self-reflection for participants to better envision future new design opportunities [126]. Such “tacit knowledge or latent needs” are often difficult to express in words but can be documented via design activities, such as sketching, mapping, and modeling [137]. To support this, we create the *Domain Specialty Canvas* (Figure 1) based on the aforementioned low- and high-level aspects of domain expertise. This canvas helps fact-checkers to map their values and problems based on local work contexts.

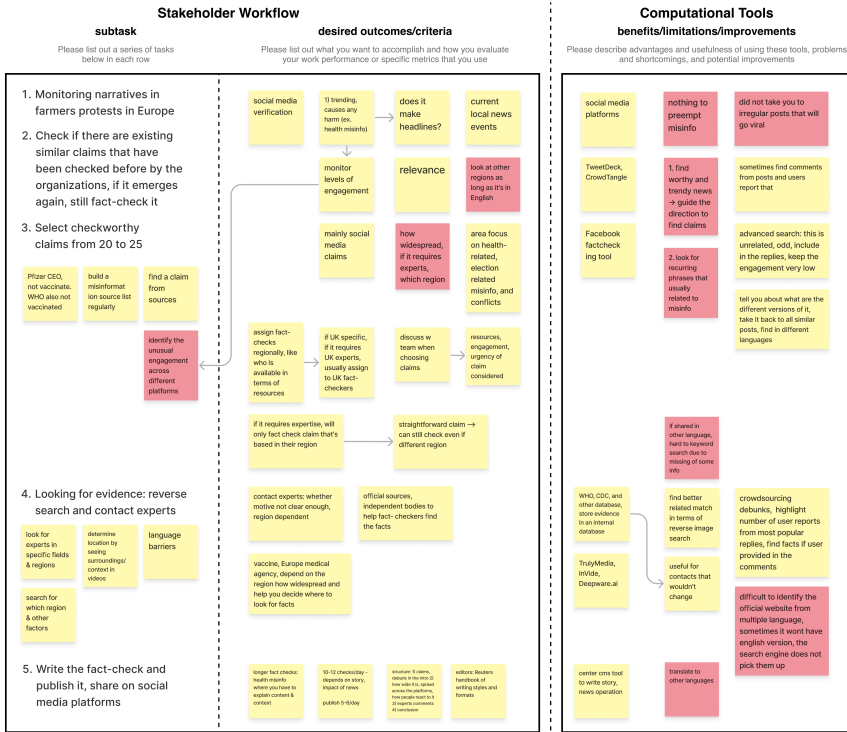


Fig. 1. The *Domain Specialty Canvas*: fact-checkers and design facilitators map out fact-checking workflow and computational tools. While fact-checkers were thinking aloud, designers facilitated distilling important concepts and recorded them (i.e., yellow sticky notes) in the canvas, followed by asking fact-checkers to affirm, revise, or add new content (i.e., red sticky notes).

We also use think-aloud protocols to drive the mapping process. During the think-aloud protocol, we first let stakeholders describe different work cases. This helps inductively abstract a typical workflow. We then complement the protocol with semi-constructed interviews that ask stakeholders how they make decisions and solve problems in each task within the workflow. Finally, we ask stakeholders to reflect on the limitations of existing AI-assisted tools they have used in their work.

3.1.2 Playing with AI Probes (Step 2). In previous human-centered design (see Section 2.2), AI abstractions are often created to help non-AI experts to familiarize themselves with AI capabilities and limitations [60, 131, 149]. Such abstractions include “visualizations, taxonomic vocabulary, or sensitizing concepts” that reveal technological capabilities as simplified concepts [149, 151]. In a similar vein, our second step leverages *AI Probes* as AI abstractions (Figure 2) to create a playground for stakeholders to have active conversations with AI researchers to become more familiar with AI concepts, behaviors, and possible errors.

As van Dijk and Zimmerman [135] discuss, it is difficult to map work activities to all possible techniques. In our matchmaking, *AI Probes* help facilitate this process. Specifically, each probe consists of two parts: 1) a visual design heuristic; and 2) an interactive WoZ (Wiz-of-Oz) demonstration. The visual design heuristic describes the basic elements of a half-baked AI idea as a pictorial diagram. It shows a standard NLP pipeline with inputs, outputs, and how inputs are processed to produce outputs. This diagram helps stakeholders understand AI from an end-user perspective

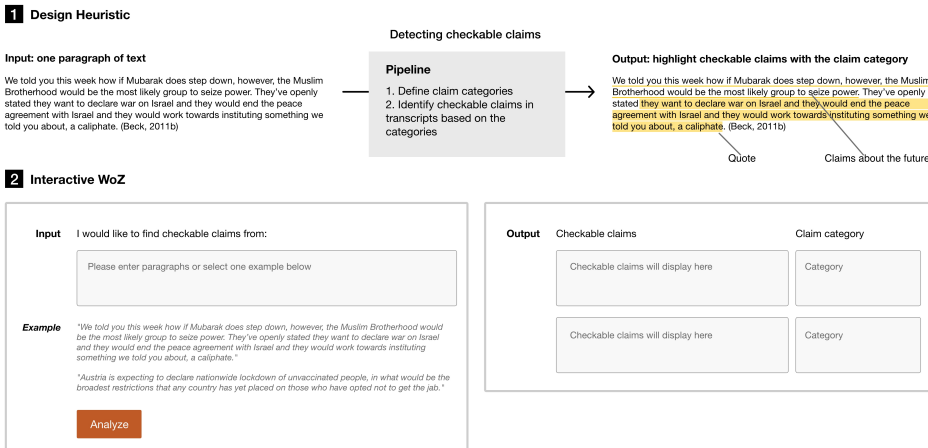


Fig. 2. An *AI Probe*: a half-baked AI idea consists of 1) design heuristic in the form of a textual diagram; and 2) abstracted NLP techniques in the form of an interactive Wizard-of-Oz simulation. This figure illustrates the first *AI Probe* (Detect checkable claims) presented to fact-checkers to help them learn the text classification technique.

without learning tricky technical details. The interactive WoZ demonstration is a design heuristic counterpart, which simulates an AI idea in a real-world context with simple functions and forms.

3.1.3 Brainstorming AI Design Elements (Step 3). AI probes help facilitate participant capacity for brainstorming through three steps. First, the visual heuristics of these probes showcase different NLP techniques (e.g., text classification, clustering, and generations), so that participants become familiar with nuanced capabilities of NLP. Second, by playing with the interactive WoZ, participants gain insights into how different NLP techniques can be applied to assist their work. This interaction creates a link between the possibilities offered by AI and their specific needs, helping to develop a mental model for aligning “technology affordance with work activity,” as outlined by Bly and Churchill [18]. Finally, in the later brainstorming sessions, participants can reflectively assess the existing relationship between NLP capabilities and user activities originated from the probes, as well as suggest novel relationships where NLP techniques could be repurposed for different applications than originally intended.

To outline a feasible AI idea, previous human-centered design studies use *Canvas mapping* [9, 131, 156] to carve out an AI lifecycle design. This includes: 1) creating a design brief with user personas; 2) specifying datasets and algorithms; and 3) designing prototypes with explainable user interfaces. Similar to these studies, we also create a *Co-Design AI Canvas* to assist fact-checkers. Our canvas is more succinct, including only AI elements (AI needs, techniques, input-outputs, and datasets) and design elements (functions, forms, and evaluative techniques). Fact-checkers have been familiarized with these elements from the previous two activities, so they use this canvas to brainstorm ideas by mapping AI techniques they learned from the AI Probes to the domain expertise, values, and problems specified from the Domain Specialty Canvas. Fact-checkers fill in these essential elements with the assistance of facilitators (see more details in Figure 3). In this step, fact-checkers can either extend previous AI Probes into a comprehensive tool by combining different AI techniques or imagine completely new tools. AI researchers help introduce new AI

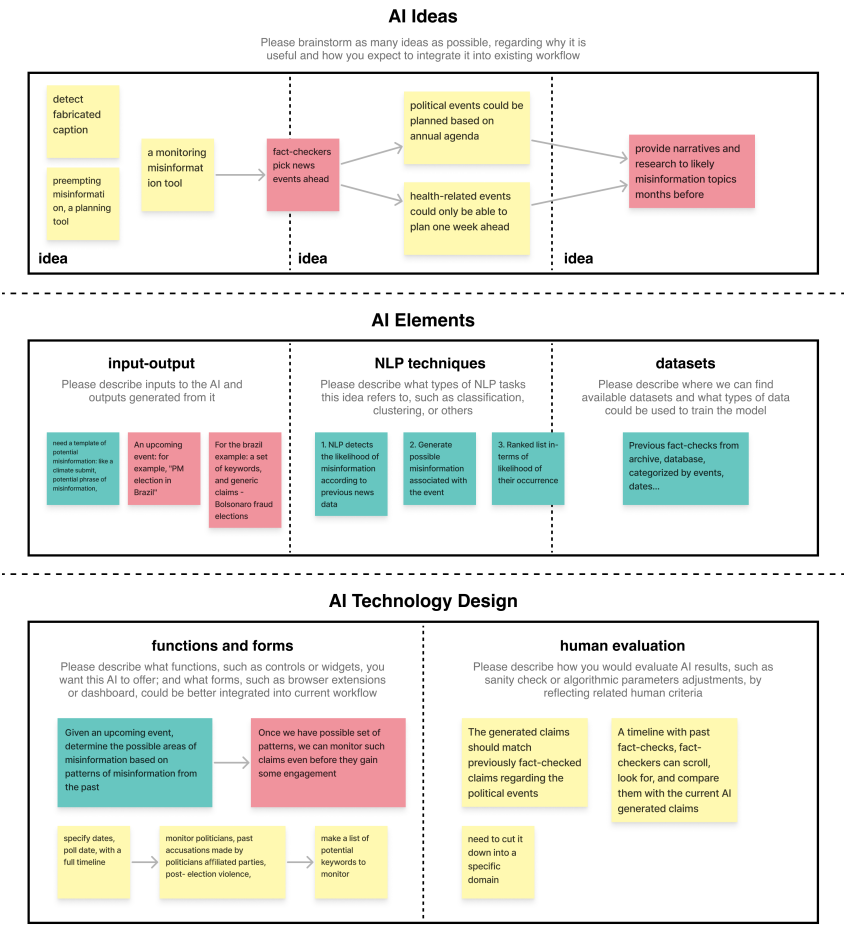


Fig. 3. The *Co-Design AI Canvas*: fact-checkers, designers, and AI researchers collaboratively brainstorm ideas by specifying different AI elements. This figure illustrates how Idea 1 (Section 4.1) is formulated. First, while fact-checkers were thinking aloud, designers helped them map out different seed ideas (i.e., yellow sticky notes). Continually, fact-checkers distilled the final one (i.e., red sticky notes and diagrams in the “AI ideas”). Later, AI researchers helped depict AI information (i.e., blue sticky notes in the “AI elements”). Meanwhile, fact-checkers provided concrete examples for these AI elements. Finally, all participants brainstormed “functions and form” and “human evaluation”.

techniques if stakeholders find techniques presented in the previous AI Probes cannot fully address their needs.

3.2 Study Protocol

In this section, we describe our workshop procedure and supplemental materials. They are developed to facilitate a two-phased workshop involving a recruited fact-checker, a designer, and an NLP researcher. The first phase was 1.5 hours long and focused on mapping fact-checker domain expertise (Step 1). The second phase was 2 hours long and included playing with AI Probes (Step 2)




Fig. 4. A fact-checking workflow synthesized from Graves [45], Micallef et al. [89] to guide participants to think-aloud their experiences.

and brainstorming NLP-based AI ideas (Step 3). We used Miro⁴ to present co-design materials and document participant feedback.

3.2.1 Workshop Procedure. For Step 1, we asked fact-checkers to bring specific claims they checked in the past, including: 1) a claim they found particularly difficult to check; 2) one whose veracity surprised them after checking; 3) their most recent check; and 4) a claim they found particularly interesting. After reflecting on these four claims, participants walked us through their typical workflow, selecting one of these claims as a working example. To facilitate this process, we prepared a fact-checking diagram (Figure 4) as a reference. This served as a useful memory aid to remind fact-checkers of the various sub-tasks they may neglect to mention.

For Step 2, we created eight specific *AI Probes*, which map existing state-of-the-art NLP techniques to different fact-checking tasks (Table 1). As discussed in Section 3.1.2, past NLP research proposed for fact-checking considered some potential ways NLP could support professional fact-checkers. However, it is important that fact-checkers verify these scenarios, as well as brainstorm themselves new scenarios or how NLP techniques might support other work activities. We organized visual design heuristics for eight probes onto a one-page “fact sheet” and encouraged fact-checkers to read it before the second workshop. We also asked them which NLP techniques they were most interested in, then had them experiment with the eight interactive WoZ demos, one-by-one.

For Step 3, we created a prompt protocol for each participant. The protocol included questions framed around a typical fact-checker workflow (documented in the *Domain Specialty Canvas*). For example: “Based on challenges identified in your workflow, do you see any opportunities to apply the AI techniques we have discussed?” Such questions oriented fact-checkers toward bridging AI techniques with their respective needs and to spur ideation. Fact-checkers then worked alongside facilitators to brainstorm ideas by filling in a set of AI design elements on the *Co-Design AI Canvas*.

3.2.2 AI Probe Design and Selection. We began by reviewing survey papers on automated fact-checking. This involved analyzing relevant NLP techniques adopted for automated fact-checking and existing computational fact-checking tools highlighted by nonprofit research organizations (e.g., the RAND Corporation⁵ and Credibility Coalition⁶) that contain resources for fact-checking. For the workshop, we choose specific probes (1, 2, 4, and 7) based on fact-checker unmet needs identified in recent empirical work (e.g., monitoring and prioritizing claims [89, 102] and augmenting ambiguous claims [122]). Research has reported that these needs are recognized but not adequately addressed by current tools [8]. We also designed probes (3, 5, 6, and 8) based on other use cases proposed by

⁴<https://miro.com/>

⁵<https://www.rand.org/research/projects/truth-decay/fighting-disinformation/search.html>

⁶<https://credibilitycoalition.org/credcatalog/>

No.	AI Probes	AI behaviors	Details	Related NLP research
<i>NLP helps to identify fact-checking related qualities</i>				
1	Detect checkable claims	Text Classification	Classify sentences into claims categories, including personal experiences, public opinions, quantitative/numeric claims, causation and correlation	claim detection[75]
2	Detect social media checkworthy claims	Text Classification	Identify checkworthy claims based on posts virality and highly-emotional content	virality and susceptibility[53], ideology detection[148]
3	Review the argumentation quality of a fact-check*	Text Classification	Evaluate the readability, coherence, and persuasiveness of writing	detect attackable sentence[61]
<i>NLP helps to prepare claims in a format that is convenient for fact-checking</i>				
4	Prepare stand-alone claims for fact-checking*	Text Classification + Text Generation	Add contextual information to a claim from previous paragraphs	decontextualization[27]
5	Identify ambiguous terms in a claim*	Text Classification + Question Generation	Generate sub-questions from ambiguous terms of a claim for later investigation	generate fact-check brief[25, 41], political ambiguity[23]
<i>NLP helps to extract additional data in an accessible, effective form to help fact-checking</i>				
6	Summarize claim-related articles*	Text Summarization	Summarize a claim's related articles into bullet points	extractive and abstractive fact-check explanations[67]
7	Retrieve quantitative related information	Text Retrieval + Text Summarization	Retrieve and format quantitative information consistently into table from the search results	unstructured data for analytical queries[85]
<i>NLP helps to determine veracity of claims for just in time claim-checking</i>				
8	Robocheck by evidence retrieval	Text Retrieval + Text Inference	Auto-verify the true and false of a claim by retrieving claim related information	detect fake news from crowd[93]

Table 1. Eight *AI Probes* were prepared to help stakeholders understand what existing NLP could offer to assist their work; ideas marked with * incorporate relatively new NLP techniques with potential value for fact-checking.

NLP researchers. While these methods have yet to be integrated into practical tools, we included them as AI probes to gauge fact-checker reactions and assessment of potential utility.

To develop probes to realistically simulate NLP capabilities for participant use, we employed a variety of methods. We utilized readily available APIs to summarize claim-related articles for Probe 6, as well as automatically verifying new claims for Probe 8. We also constructed basic Naive Bayes models for sorting claims into categories of *checkability* (Probe 1) and *check-worthiness* (Probe 2).

In cases where we could not easily generate outputs using advanced models (i.e., Probes 3, 4, 5, 7), a mix of accurate and inaccurate results was produced based on the test examples presented in NLP research (listed in Table 1).

3.3 Recruitment

No.	1st	2nd	Age Group	Gender	Ethnicity	Role	Organizational Context	Languages Fact-Checked	Years of Fact-Checking Experience	Country	Region
1	✓	✓	26-35	Female	White	Fact-checker	Ambiguous	English	1	USA	North America
2	✓	✓	36-45	Male	Black	Fact-checker	Independent	English	14	USA	North America
3	✓		46-55	Female	White	Editor	Independent	English	15	USA	North America
4	✓	✓	26-35	Female	White	Fact-checker	Independent	English & Afrikaans	0.5	South Africa	Africa
5	✓	✓	46-55	Male	Indian	Editor	Independent	English & Nepali	2.5	Nepal	Asia
6	✓	✓	26-35	Female	Indian	Fact-checker	Independent	English	3	South Africa	Africa
7	✓		26-35	Female	No response	Editor	Independent	Turkish	5	Turkey	Europe
8	✓	✓	18-25	Male	White	Fact-checker	Independent	English & Afrikaans	2	South Africa	Africa
9	✓	✓	46-55	Male	Pākehā/European-Australian/White	Fact-checker	Media	English	2	New Zealand	Oceania
10	✓		26-35	Female	Aryan	Fact-checker	Independent	Nepali	5	Nepal	Asia
11	✓	✓	18-25	Male	Hindu, Bhramin	Fact-checker	Independent	Nepali & English	0.3	Nepal	Asia
12	✓	✓	18-25	Male	Caucasian	Fact-checker	Media	English	2	USA	North America
13	✓	✓	26-35	Female	Black African	Fact-checker	Independent	English	2	South Africa	Africa
14	✓	✓	18-25	Female	White	Fact-checker	Media	English	1.5	USA	North America
15	✓	✓	26-35	Female	Pākehā/European-Australian/White	Fact-checker	Media	English	0.5	New Zealand	Oceania
16	✓	✓	18-35	Male	Prefer not to Answer	Fact-checker	Independent	English & Filipino	3	Philippines	Asia
17	✓	✓	26-35	Male	Hispanic	Fact-checker	Media	English & Spanish	2	Mexico	Latin America
18	✓	✓	26-35	Male	White	Fact-checker	Independent	English	7	USA	North America
19	✓	✓	56-65	Male	White	Editor	Independent	English	7	USA	North America
20	✓	✓	26-35	Male	White British	Fact-checker	Media	English	3.5	England	Europe
21	✓	✓	26-35	Male	Asian	Fact-checker	Media	English	1	India	Asia
22	✓	✓	56-65	Female	Prefer not to answer	Fact-checker	Independent	English	1	USA	North America

Table 2. Participant Demographics and Work History

We recruited ($n = 22$) professional fact-checkers, including 5 present or former editors and 17 full-time fact-checkers (Table 2). All editors had previous fact-checking experience and often guide their organizational fact-checkers. All participants participated in the first workshop session, and 20 participated in the second session. Our first workshop session lasted 1.5 hours, and the second lasted 2 hours. Each participant was compensated with Amazon Gift Cards or other redeemable options available in their areas.

Participants (11 women and 11 men) were drawn from 9 countries, including the USA, South Africa, Nepal, Turkey, New Zealand, Philippines, Mexico, England, and India. We intentionally sampled participants from diverse countries, fact-checking topics, organizations, ages, genders, etc. Among the participants, 11 (52.4%) identified as White, 5 (23.8%) as Asian, 2 (9.5%) as Black, 1 (4.8%) as Hispanic, and 2 (9.5%) did not self-identify. 5 (23.8%) identified as being in the 18-25 years old age range, 11 (52.4%) in the 26-35 range, 1 (4.8%) in the 36-45 range, 3 (14.3%) being between 46-55, and 1 (4.8%) in the 56-65 age range. Their fact-checking experience ranged from 6 months to 15 years, with an average of 3.8 years. 18 were from independent fact-checking organizations, while the remaining 4 were from news media outlets. Upon asking fact-checkers about their familiarity with AI, most of our participants had little knowledge of computational algorithms and programming, yet most had used different computational tools to assist their work. These include:

1) Open-source Intelligence (OSINT) tools, such as the InVid; 2) specific claim monitoring tools, such as CrowdTangle, TweetDeck, and Full Fact Alpha; 3) off-the-shelf search engines, including Google and Bing; and 4) databases, such as the Internet Archive and LexisNexis.

3.4 Analysis

We conducted all co-design sessions with recruited participants on Miro and Zoom and then transcribed the recordings for analysis. Building upon previous empirical studies [45, 89], we generated a set of codes for existing fact-checking practices to analyze transcripts from Step (1). Other codes were generated via a bottom-up thematic analysis approach based on how fact-checkers played with AI Probes in Step (2) and what ideas they brainstormed in Step (3) [98]. The qualitative coding was initially conducted by a single primary researcher and then reviewed by a second researcher. During the open coding, the primary coder examined transcripts on a sentence-by-sentence basis, examining whether any findings could be classified by existing literature or whether a new code needed to be created. All of these codes are divided into three parts: fact-checking criteria, stakeholder design experience, and co-designed ideas. We aggregate the fact-checking criteria into tables in Appendix A and B. Our Results (Section 4) primarily focus on AI ideas co-designed with fact-checkers. Our Discussion (Section 5) then summarizes fact-checker design experience as evidence to evaluate the utility of our co-design method. As our participants represent organizations in multiple countries, we also analyzed the potential influence of regions and organizational interests on participant needs and co-created ideas. Given that we have a few participants representing these different groups, we report them in Appendix D as exploratory findings.

3.5 Positionality Statement

Our research team has members with diverse backgrounds, including HCI, design, NLP, information retrieval, crowdsourcing, and AI ethics. The main workshop facilitators include one UX professional with years of industry experience and another UX researcher who has received academic training in HCI. The two NLP researchers who participated in the second workshop activity have expertise in explainable AI and NLP applications for misinformation. The interdisciplinarity of our team allowed us to take a human-centered approach to develop *Matchmaking for AI* and our co-design study. We take a constructive stance, believing that NLP has the potential to assist in fact-checking.

As both designers and NLP researchers are also co-authors of this paper, we recognize that our perspectives and roles naturally influenced the research process. Bearing this reflexivity in mind, we sought to minimize researcher-induced biases. First, we invited an expert with over 15 years of experience in fact-checking research and misinformation reporting to refine our workshop scripts and study procedures. Second, during the matchmaking, we encouraged participants to brainstorm ideas by pinpointing problems reported in the *Domain Specialty Canvas*, ensuring that ideas originated directly from participants. We further discuss the potential limitations of our study in Section 5.3.

4 Results

In this section, we describe the ideas generated from our co-design workshops that brought together AI experts, designers, and professional fact-checkers. Given our definition of “workable ideas” (Section 3.1), for each idea, we provide details on stakeholder needs, gaps in existing tools and related work in AI research, desired tool features, and potential technical and research challenges. These details are synthesized from participant notes on the *Domain Specialty Canvas* and the *Co-design AI Canvas*, along with their reflections on the *AI Probes*. While there are parallels to the workable ideas in related AI research (described in Table 3), most of these ideas require building

datasets specific to the tasks and additional domain-specific modeling work for implementation. Moreover, we identify if there is a need for more research in AI capabilities to implement them, in turn providing a road map for AI researchers interested in fact-checking. We expect further technical and design work will be needed to develop, implement, and deploy these ideas in practice.

Table 3 summarizes the 11 co-designed ideas, which span the fact-checking lifecycle from claim monitoring, selection, and investigation to the final fact-check writing stage. Compared with the dominant cases of using NLP techniques for fact-checking — detecting and ranking claim check-worthiness, retrieving previously fact-checked claims and evidence, and auto-verifying claims [32, 92] — we find other interesting use cases, such as forecasting disinformation (Idea 1, Section 4.1), contextualizing ambiguous claims (Idea 4, Section 4.4), human-AI teaming for generating fact-check briefs (Idea 7, Section 4.7) and co-writing fact-check reports (Idea 10, Section 4.10). We categorize these use cases as design ideas and highlight their novelty based on: 1) whether this design idea is completely new to fact-checking; or 2) whether our participants uncovered new needs not identified in prior work. In the following sections, we present ideas that either reflect new design or address unmet needs. We also synthesize ideas from participants that appear significant even if they are not particularly novel. We include these in Appendix C. In Appendix D, we present an exploratory analysis on how different regions and organizational interests influence participant needs and co-created ideas.

4.1 Forecasting Disinformation

Disinformation and misinformation frequently reoccur. A simple example is when old fake news is “recycled” (often verbatim or with only minor tweaks) and re-distributed anew. Slightly more sophisticated is when certain types of news events (e.g., an election or a pandemic) are routinely followed by regular patterns of disinformation (e.g., “*the presidential candidate is associated with <insert any crime>*”) [133]. Participants noted that before major events like presidential elections, they often release “fact sheets” or “misinformation fact-check guides” to the public to help “inoculate” them against potential fake news exposure. However, this process typically requires extensive analysis of social media engagement, such as identifying what false news the public is likely to believe as true (P4, 19) or relies on fact-checker intuition to predict potential issues (P4). Thus, participants wanted to understand if and how AI could systematically assist in pre-bunking. Specifically, they wanted AI to help identify which facts are more crucial to report on and prevent the possible recurrence of similar false claims: “If three months in advance I’m already aware of misinformation topics that would spread, I will have my coverage ready to counter those narratives... [so that] these misinformation claims don’t actually gain traction” (P21).

Generating misinformation has mainly been seen by AI researchers as a method to mimic possible misuse of LLMs [97] or to evaluate the effectiveness of new misinformation detection tools [57]. In contrast, our study participants have redirected the use of a generative model toward a more beneficial use scenario, i.e., forecasting potential fake news for recurring news events to be better prepared for fact-checking potential misinformation.

During the matchmaking Step (3), participants (P4, 21) and NLP researchers brainstormed the use of generative models to create such potential misinformation ahead of an upcoming event, e.g., given context of the event type. Models might generate a set of possible fake claims or a ranking ordered by their predicted probability of occurrence based on historical data. For example, to forecast possible fake news claimed by candidates for an election, P21 said the AI tool should be able to “identify possible set of [discourse] patterns or narratives on [historical] fake results, rumors, and previously fact-checked claims.” Additionally, it would be beneficial to provide data insights, such as summarizing how voters from various geographical regions have significantly engaged with these claims on social media (connected to Idea 9, Section 4.9). Such historical evidence and

Fact-checking stages	No.	Design ideas	New idea	New need	Related AI research
Forecasting claims	1	Forecasting Disinformation	✓	✓	[57, 97]
Monitoring claims	2	Identifying Broader Disinformation Narratives		✓	[6, 49, 59]
	3	Dynamic Credibility Monitoring of Social Media Users	✓	✓	[43, 93, 103, 154, 155]
Selecting claims	4	Finding and Providing Context for Ambiguous Claims	✓		[25, 27]
	5	Personalized Claim Filtering and Selection		✓	[3, 4, 6, 48, 49, 65, 91, 117, 136]
	6	Personal Bias Warning System for Claim Selection	✓	✓	[6, 49]
Investigating facts	7	Human-AI Teaming for Fast Generation of Fact-check Briefs	✓	✓	[24, 25, 41, 94]
	8	Identifying Official Databases and Formulating Queries for Verifying Quantitative Claims		✓	[10, 64, 110]
Writing fact-checks	9	Understanding Reader Engagement for Fact-checking Reports	✓		[153]
	10	AI Assistance in Writing Fact-checking Reports	✓	✓	[24]
	11	AI Assistance in Critiquing and Editing Fact-checking Reports		✓	[138]

Table 3. Matchmaking-for-AI outcomes across design ideas, AI research, tools, and resources. This table describes the outcome of our co-design process. Our process led to identification of 9 novel needs for fact-checkers and 7 new ideas that can inform future tools for fact-checking. Additionally, we provide pointers to the literature related to the design ideas that can be used by developers as a starting point.

insights could bolster their trust in use of AI: “I can trace it back to understand: Are there any facts on these things [AI predictions]?” Regarding specific tool features, participants imagined that a timeline would be particularly useful, allowing them to scroll, look for, and compare previous claims, fact-checks, and historical data with the AI-generated claims.

While language models have been used to generate variants of existing claims (e.g., counterfactual claims [109]), predicting future fake news would be significantly different for model training and evaluation. Researchers could curate new datasets that pair examples of past news events with associated fake news related to those events. Participants also suggested a way to evaluate

the prebunking models before trusting model outcomes in practical scenarios. For example, P21 suggested creating a test set by putting aside claims from a known source, training a model on older claims, and testing if the model could produce unseen claims in the test set. Even when the model is not 100% accurate, the participants deemed the prebunking model helpful to guide their intuition in predicting what claims might occur in the future. In the context of misinformation in political speech, P21 said, “it still helps you understand the candidate’s background, chase information on this topic ...”. Participants thought it could be helpful in trying to anticipate all possible misinformation that might occur.

4.2 Identifying Broader Disinformation Narratives

Individual claims often reflect larger narrative patterns in thematically connected claims. For example, a theme of *distrust in technology creators and providers* might connect claims such as: a) Bill Gates is using the COVID-19 vaccine to implant microchips, b) self-driving cars self-destructing, and c) Twitter uses improper means to censor content (P7). Several participants (P1, 7, 19, 21, 22) refer these narratives as “conspiracy theories” or “hoax” that can easily provoke social media users for sharing. Recognizing such narratives are important, assisting fact-checkers to: 1) monitor disinformation by understanding more on what to search for (see Idea 5, Section 4.5), and 2) develop counter-strategies against these narratives when writing fact-checks.

Communication studies [38, 39] suggest that people often believe and share disinformation narratives for psychological or social needs, such as the desire of curiosity, avoiding uncertainty, or group belonging. Our participants (P7, 21) note that by recognizing these narratives, they can “understand motivations [people] share and how disinformation affects different groups”, and informs them to write “about the truth to convince people more easily” (P7). Fact-checking reports informed by the broader narrative might help alleviate feelings of collective insecurity or discomfort caused by the disinformation. However, P7 points out that existing tools do not provide insights into broader narratives or connections between different ones; they have to manually piece together potential narratives, or perform a manual retrospective analysis from previously fact-checked claims.

Participants were interested in whether NLP methods could help identify such larger narratives. Specifically, two settings were envisioned: 1) given a set of claims, identify the broader narrative; and 2) given a narrative statement, generate claims to fit the narrative.

For the first, as with the previous forecasting idea (Section 4.1), generating such claims could help fact-checkers prepare for anticipatory pre-bunking. P7 gave an example of generation from a narrative, “vaccines are bad for people”, with possible false claims, such as “vaccines are a game of US and China” or “Bill Gates called for the withdrawal of vaccines.” NLP techniques such as topic modeling or unsupervised clustering could help group claims together to identify claims that have similar narratives [6, 59]. Furthermore, generative models built to provide natural language descriptions of such clusters could help describe the narratives.

For the second setting, a generative language model could be trained to produce example claims given a narrative. For example, a state-of-the-art generative model (such as GPT-4⁷) might be applied in a zero-shot or few-shot manner [144]. However, because generative AI providers are increasingly wrapping their service offerings with “guardrails” to prevent adversarial use (such as generating disinformation, as envisioned here), it may be necessary to use an open-source model or obtain privileged API access without the guardrails. In addition, evaluating the quality of generated claims would require human review and/or an annotated dataset.

⁷<https://openai.com/research/gpt-4>

4.3 Dynamic Credibility Monitoring of Social Media Users

As an example of the 80/20 rule, a disproportionately large amount of misinformation tends to be spread by a relatively small handful of online sources (e.g., social media users). Participants reported that they often manually curate and monitor content from known, repeat offenders. This led to the suggestion (P5, 19-21) of developing a “credibility checker” tool that could automatically estimate and continually update the predicted credibility of such users (e.g., based on the content they post, their social network of followers and who they follow, etc.). To assess the authenticity of social media profiles, related prior work has commonly employed tools that are used to identify social media bots, such as the Bot sentinel or Botometer [52, 104]. In our study, participants advocated for the value of long-term tracking of personal credibility of real individuals.

Note that the *Global Disinformation Index* (GDI)⁸ [124] already provides a “neutral, independent, transparent index of a website’s risk of disinforming readers.” Similarly, existing claim detection methods already pay attention to individual posts or known malicious sources at an organizational level [155]. As participants (P5, 14, 20) noted based on the 1st and 2nd AI probes, claim detection techniques having individual credibility ratings could be used to achieve this. The key technical challenge would be modeling and estimation of user credibility, and the frequency of ingesting new data and updating model estimates for each user.

Political and ethical considerations around such credibility monitoring were also discussed. For example, GDI has faced intense criticism and litigation in response to its domain-level source credibility ratings. There are also risks of harm from assigning low credibility scores due to AI detection errors, especially if those scores became public. Additionally, there are questions of how easily a reformed user could overcome a low credibility score given a past history of bad behavior. During the matchmaking Step (3), however, participants did not flag any strong ethical concern towards building such a tool, as many already monitor potentially problematic sources of misinformation internally. P14’s organization has a long history of monitoring and sharing politician records regarding sharing false claims. Similarly, P19 stated that they released information about social media users who are probable sources of misinformation. Additionally, fact-checking programs from Twitter (now X) and Facebook were frequently mentioned by participants as community efforts committed to overseeing misinformation spreaders.

4.4 Finding and Providing Context for Ambiguous Claims

Textual claims on social media often lack sufficient context to be clearly understood [122]. Claim checking such texts often requires additional disambiguation and world knowledge. Participants raised the idea of using AI to augment a claim with contextual information so that the claim’s terms, concepts, and/or arguments are less ambiguous. This stemmed in part from our AI probe using stand-alone claims, leading some participants to add information about what needed to be disambiguated. P5 explained that “what misinformation does is it removes the context”. For example, in the claim that “Midtown crime is up by 30% the last quarter” (taken from [122]), “midtown” could refer to Midtown Manhattan in New York City, or several cities in the US state of Florida. P1 voiced that the contextual information could “help us determine whether or not that there was a basis for the claim and if that’s something we needed to do a deeper dive into”. This contextual information could include background knowledge (e.g., time and location - P5) and cultural context (e.g., who is the targeted audience and who will be interested in the claim - P5, 11, 15).

Existing NLP research has explored methods to augment these ambiguous claims with more contextual information. These include rewriting sentences to maintain the same meaning by summarizing previous text [27] or retrieving knowledge from an external database [147]. However,

⁸<https://www.disinformationindex.org/>

these techniques have not yet been implemented in current fact-checking tools. In this study, we discuss design elements necessary for the practical implementation of claim disambiguation. Participants suggested NLP tooling here might: 1) identify the spans in a claim that need disambiguation, 2) formulate queries for additional information required, 3) retrieve the relevant information, and 4) rephrase the claim accordingly. Complex or ambiguous arguments could also be decomposed into simpler sub-claims using question generation models with human supervision [25]. Inference beyond the sentence boundary may require anaphora resolution⁹ [100] and entity-linking¹⁰ methods [99]. While the recent advances in NLP show promise for these tasks, their effectiveness in fact-checking requires further investigation. One approach is to formulate an effective workflow for Human-AI collaboration. For example, queries might be handled internally to the tool, externally via another tool (e.g., search engine, knowledge base), or provided and rephrased via a human-in-the-loop (HITL) approach [34]. A dataset consisting of original and rephrased claims would also be invaluable to train and evaluate tools for this task [27].

Disambiguation tools might falsely rephrase claims. To avoid this issue, P1 envisioned that a disambiguation tool must provide references used for generating the new claim, which makes it easier for fact-checkers to conduct sanity checks. Additionally, P5 said that even if the tool cannot successfully rephrase a claim, it would still be useful to at least highlight the parts of the claim that need further clarification.

4.5 Personalized Claim Filtering and Selection

Because the number of claims that could be checked far exceeds the capacity of human fact-checking, identifying the most checkworthy claims is crucial [114]. Moreover, the diverse criteria used to gauge checkworthiness, and their relative emphasis, varies across individual fact-checkers as well as their organizations (Details of such differences are described in Appendix D). Examples of such criteria include: severity, topic, location, political leaning, targeted audience, etc. [46]. While prior NLP work has formulated a checkworthiness prediction task [4, 12], the one-size-fits-all nature of this formulation neglects the variance described above.

Consequently, participants expressed a desire for greater control and personalization in claim filtering and selection tools (See Appendices A and B for a full list of personalization criteria). For example, some participants (P2-3, 8-9) believed that a claim should be more checkable if there is already publicly available evidence, and thus should be ranked higher by a claim monitoring tool. This would require information retrieval of related evidence in scoring checkworthiness. Similarly, P1 and P21 wished to compare the *spread velocity* across different claims. If a claim reaches a certain group of audience faster than other claims, P1 and P21 may be more interested in checking that claim. This would require creating and integrating a model of such claim spreads.

Participants also suggested labeling claims for checkworthiness based on individual or organizational interest. Labels would be collected through fact-checker daily work to train personalized prediction models. They also discussed providing manual feedback to prediction models in order to similarly personalize them. P12 said she wanted to be able to weight different criteria and let the model learn from her feedback. Similarly, P8 expressed a willingness to “remove or adjust the checkworthiness ranking if a claim is not worthy of a high score.” However, P8 also expressed concern about such a trained model suffering distribution shift over time: by labeling “claims... made immediately after a breaking news event, the model might weigh related claims as more checkworthy” long after that story was no longer newsworthy.

⁹Anaphora resolution requires resolving references to entities through different forms in a sentence (i.e., pronouns, abbreviations)

¹⁰Entity linking helps to identify the unique entities (i.e., people, places, government bodies) in text

Though expressed with regard to claim selection, a more general theme regarding human-AI interaction is personalized user training and calibration. For example, to make sense of the checkworthiness scores predicted by AI, P20 and P21 suggested an onboarding process. As P21 remarked, “the combined [checkworthiness] score will be helpful once we have used the tool enough to set our own threshold for those ratings.” Regarding human-AI calibration, P20 also mentioned that because “there’s some room for error that’s taking into account these [AI] scores”, he would like the AI to convey how reliable its predictions are, such as by sharing its error rate akin to model cards [90].

4.6 Personal Bias Warning System for Claim Selection

Participants were highly conscientious regarding the risk of personal biases influencing their work (see Appendix A). Repeatedly, there was mention of the importance of maintaining strict impartiality, despite acknowledging the loftiness of this ideal. For example, P1 remarked, “there were definitely examples where I think my own personal interests and backgrounds kind of led me to practice that over some more local stuff that would have been harder for people to individually fact check ... maybe that’s [an area] where my own personal interests kind of affected my agenda.”

Similar to the media bias chart, which has been commonly used by journalists [120], participants envisioned an AI tool that could be an active partner in recognizing, classifying, and quantifying personal biases in claim selection. Key features of this tool included: 1) making fact-checkers aware of topical distribution of the claims they have selected (P4, 5), 2) flagging any potential bias towards selecting claims to fact-check depending on their source (e.g., political, regional, or racial) (P1), and 3) identifying claim selection patterns of individual fact-checkers by clustering them across attributes of interest (P1).

P5 was positive toward such an AI tool, noting, “we talk about the algorithm bias as well, [as] the way they [were] built, but I think the machine will be better than human in detail, or in not being as bias[ed].” P4 also thought a tool that “not limiting us in terms of topic area would be super helpful.” Such a ‘bias detection’ AI could help participants reflect on their own personal biases in claim selection, backing up assumptions with tangible data, and enabling them to counteract such biases. However, some participants also worried about how such a tool would produce results that might misrepresent the real situation if it takes a simplistic assumption about biases. For example, P15 said “if all the false claims are coming from one party and not the other party, then we shouldn’t try to create a false equivalence.”

Currently, NLP researchers use media bias charts¹¹ and crowdsourced political ideology ratings to create datasets to determine the political ideology of articles and news outlets [26]. Models trained on these datasets could be used to categorize sources collected by a fact-checker and to examine whether these sources are neutrally selected. Topical analysis could also help find common themes from claims selected. Note that identifying and mitigating personal bias is challenging, such as oversimplifying the bias problem by forcefully trying to restore a balance of coverage, as P15 mentioned above. Further research is required at the intersection of bias and fact-checking more broadly. We summarize other aspects of potential biases mentioned by our participants in Appendix A.

4.7 Human-AI Teaming for Fast Generation of Fact-check Briefs

Participants described fact-checking *briefs* as short, early drafts used internally to pitch and negotiate which claims would be fact-checked (P8, 9, 15). As P9 described it, a typical brief consists of “a rough indication [of] what the topic is, what we think the verdict is, and here’s what we want to

¹¹<https://www.allsides.com/media-bias>

do.” P14 begins fact-checking by first creating a fact-check brief, then pitching it to her editors in a team meeting before gaining approval to pursue a more detailed investigation. P15, a fact-checking intern, said that one of her main tasks is to prepare fact-checking briefs for senior fact-checkers. P9 and P15 both imagined that the AI can help improve efficiency by generating briefs in the same format as human-prepared briefs. While a prior NLP study has investigated automatic generation of fact-checking briefs [41], the notion of briefs in that work (designed for crowd workers) is disconnected – in content, format, and use – from the notion of briefs being discussed here as used by professional fact-checking organizations.

Beyond AI-generated briefs, participants also envisioned a hybrid, human-AI collaborative approach. In this vein, participants imagined that AI could generate questions to help fact-checkers research a topic from multiple angles (closer in spirit to [25] than [41]). For example, imagine one wished to fact-check the following claim: “People in the Muslim Brotherhood openly stated they want to declare war on Israel” (taken from [45]). P6 said that potential AI-generated questions that would “essentially give [her] the next steps” in her research process include: “Who in the Muslim Brotherhood claim this statement?” and “When and where is the claim made?”

Beyond aiding negotiation of which fact-checks to perform, fact-checkers also envisioned a larger role for such briefs in helping to guide the ultimate fact-check itself. P15 wanted the AI to retrieve evidence from different sources, so she could then decide which directions to pursue. Participants also noted that questions lacking sufficient evidence could be a signal indicating the claim to be false. Additionally, P4 thought this AI could help in “reducing the amount of time it takes sifting through information to find something that’s relevant,” as well as helping her “think about things that we might sort of take for granted as known already” to avoid self-presumptions (related in spirit to detecting and mitigating self-bias, as discussed in Section 4.6).

4.8 Identifying Official Databases and Formulating Queries for Verifying Quantitative Claims

Access to statistical information is essential for fact-checking quantitative claims. Checking complex claims that merge statistics from various sources requires manual effort to contextualize these statistics by examining data from different sources. Participants noted that identifying the appropriate database is often a time-consuming task (P2, P18). P2 said, “[it] slows us down ... it’s a long drawn out process.” To address this issue, participants imagined a tool that could help infer the official databases required to fact-check the claim and formulate queries for them (e.g., government databases and research statistics published by reputable agencies). For example, when a news article cited data from the Center for Disease Control and Prevention (CDC), P18 asked “what [specific] database do I need to go to [from the CDC]? And what parameters do I need to input in order to get the data that was just cited by the New York Times?”

Participants imagined an NLP model to first detect the specific parts of a text referring to quantitative data, and then generate queries to verify that data (including any additional metadata, such as location and date). For example, for the claim, “the Supreme Court allowed warrant-less home searches within 100 miles of the US border”, outputs of interest would include: 1) labeled and extracted parts of the claim, 2) what original source the article refers to, 3) possible source of the statistics (i.e., public warrants issued by Supreme Court), and 4) other metadata (i.e., warrants effective date).

Building such an AI model would require integration across several standard NLP tasks. First, a sequence-tagging model [17] is required to detect quantitative parts of a claim and tag referred sources. Then, a generative model could be used to construct necessary queries for verifying the claim using that sequence tag from the tagged spans. To build and evaluate such a tool, a supporting

dataset would provide a structured output containing: a) the official source, b) the exact name of the database to be consulted, and c) related metadata.

Recognizing the complexity involved, participants also acknowledge that a solution need not be fully automated to be useful. Specifically, P18 pointed out that using the AI tool to track and extract the sources of information is intended to simplify the search process, not to have AI directly provide answers to fact-checkers. They stressed that fact-checkers are still required to confirm the origins of all information themselves. During the discussion, NLP researchers also suggested a HITL approach [34], such as engaging online crowdworkers to help find desired information across diverse web databases. This would reduce work for professional fact-checkers while using crowdworkers to scaffold desired functionality until an automation solution becomes feasible.

4.9 Understanding Reader Engagement for Fact-checking Reports

Fact-checkers dedicated to curbing the spread of dis/misinformation care deeply about measuring the reach and impact of their work. Like other journalists, they want to understand who their readership is (and is not) and how they engage with posted content. Given the abundance of empirical studies that analyze reader engagement of fact-checking in the context of academic journalism [69, 106], there is a corresponding need for advanced tools to assist professional fact-checkers in analyzing reader behavioral data. Most participants reported using web analytics tools (e.g., Google Analytics¹² and Trendolizer¹³) to measure the reach and spread of their work, but these tools are often limited to examining the readership of fact-checking reports.

Our participants (P4, 7-9, 13) wanted to analyze how readers of different demographics engage with their work, and to compare the impact of their work to other fact-checkers within or across organizations. Regarding the readership composition, P7 wanted to know, “what areas it’s spreading or the common characteristics of people who are engaging with... [Theoretical targeting] really helps us reach different groups of people while spreading our accurate information so we try to change our narratives to reach [more] people ... [or] you can just maybe convince people [more] easily.” P4 echoed that, “We have a claim, and we want to have some storytelling about the contextual information surrounded, and what kinds of potential audience, demographics, attraction, and attention it attained. On social media, this figure is very important.” Akin to search engine optimization (SEO), such understanding would enable fact-checkers to make specific changes to their fact-check reports to better reach and impact their intended audience(s).

As noted above, measuring impacts on live traffic is common by instrumenting webpages to collect analytics quantifying user engagement with content. However, this only works after content has already been posted, and conducting A/B experimentation of variant presentations with live traffic is always risky. Conducting surveys or focus groups with target audiences provides more focused and qualitative feedback, but is also post-hoc and involves greater latency.

Akin to forecasting disinformation (Section 4.1), forecasting spread and audience reactions (as a function of fact-check content or presentation) has potential to enable fact-checkers to predict such impacts before content is posted, and to tailor it to optimize impact. For example, network analysis can be performed to predict spread [153]. Other work has correlated content analysis of text with audience reactions to guide best practices in general [69]. Perhaps most intriguing from an NLP perspective is recent work seeking to simulate the reactions of different demographic groups via large language models [7, 70, 112]. Such work is quite nascent but early work is encouraging.

¹²<https://developers.google.com/analytics>

¹³<http://get.trendolizer.com/>

4.10 AI Assistance in Writing Fact-checking Reports

Participants imagined NLP tooling to aid them in writing fact-check reports, especially the repetitive, standardized portions. As P3 put it succinctly, “We spend too much time writing and copying.” With the rapid rise and spread of large language models (LLMs) such as GPT-4, it has become increasingly common to have humans revise AI-generated texts rather than writing texts from scratch. To do so effectively, researchers found that LLMs must be tailored to meet domain-specific needs (i.e., fact-checking) and a claim-driven extractive step helps improve abstractive summarization [107].

Because fact-check reports usually have a consistent writing structure within organizations (see Appendix B), fact-checkers first want to specify the structure, then use AI to generate content accordingly. For example, P2 wanted AI first to generate descriptions of who said the claim and the source, such as “say it in a tweet that ...the White House claimed that...,” then explain the evidence “there was no vaccine available at the time...” More generally, participants (P3, 9, 19, 20) stated they would like to provide the claim, its verdict, and the evidence they have collected to verify the claim. The generated report would contain a title, ClaimReview¹⁴ tags, an archived URL of the claim, blockquote fields linking the claim, and captioned images with the verdicts. Additional desired features included producing a description of that claim, possible suggestions for writing supports or counterarguments to the claim, and synthesis of any other raw information provided by the checker.

While today’s LLMs can take natural language prompts as input and provide suggestions for writing different parts of a report, open research challenges include performing complex and implicit multi-step reasoning (e.g., generating fact-checking reports [24], retrieving relevant and accurate external information to enhance generation [84], appropriately citing such evidence [63], and enforcing other specific formatting requirements).

4.11 AI Assistance in Critiquing and Editing Fact-checking Reports

Whereas the previous idea envisioned AI generating a rough copy of content for humans to refine, this idea reverses these roles, with humans writing polished content and the AI taking on the role of the human editor, critiquing and editing content. Inspired by our 3rd AI probe to review the argumentation quality of a fact-check, participants envisioned that AI models could help with judging their writing quality. They imagined that a model could reliably judge fact-checking reports across dimensions such as readability, coherency, clarity, and quality of the argument presented (more such criteria is listed in Appendix B). P4 said, “if we’re writing in English, everything has to be accessible to people who have a limited educational background and whose first language isn’t English.” P6 also echoed that “[you need to] ensure that concepts are thoroughly explained. Inflation, for example, not just talking about it [as a concept], but actually explaining what, when, and how.”

Moreover, the participants imagined AI tools will take domain-specific content and presentation requirements into consideration and allow customization to capture their organization’s writing constraints. For example, P9 highlighted that: “every organization seems to present their checks in different styles. If there was some flexibility in this model, where you could input your own organization’s style that would be cool.”

Participants were also interested in subjective aspects of the writing, such as, the strength of presented evidence or pointing out silos in their arguments. P9 said, “[the tool can check] if all aspects of the claim have been supported with clear evidence and they are clearly linked together.” Moreover, P8 said because journalists can sometimes be oblivious to their confirmation biases, if AI

¹⁴<https://schema.org/ClaimReview>

could provide adversarial challenges to the written arguments, it could potentially improve the reporting to maintain greater objectivity.

While the use of LLMs to generate draft content is increasingly popular, the idea of AI automatically critiquing and refining professionally produced, polished content appears to be more novel, challenging, and subjective. In contrast with familiar automated checking of spelling and grammar, and perhaps even automated student essay grading in standardized testing, sophisticated editing of professional content would seem to take such AI-assistance to an entirely new level. Further consideration of domain-specific content and presentation requirements would make this task even more challenging yet.

Current research in AI-assisted writing has looked into more subjective domains such as story writing [80, 146] by using LLMs. Prior research [80, 146] suggests the promise of using LLMs as effective collaborators for different writing tasks. However, more research is required for AI-based critiques for subjective aspects of a fact-check report. As a starting point for editing fact-checks, one could explore few-shot learning [19] with a relatively small expert annotated parallel data (a corpus that contains both the initial text and the edited version of the text) that helps edit the initial text to a more-polished version appropriate for fact-checks.

5 Discussion

In this section, we discuss the benefits, efficacy, limitations and improvements of *Matchmaking for AI*, including its implications for human-centered fact-checking and general AI co-design. First, we present a summary of co-designed ideas for fact-checking to illustrate their novelty, extending previous empirical studies. We also explain that *Matchmaking for AI* is an example of translational HCI practice as we produce design ideas useful for practitioners. Additionally, our findings suggest that matchmaking helps brainstorm ideas with concrete design suggestions and relevant technical details. Finally, we articulate its limitations and future improvements, helping researchers better adapt this co-design approach to other AI application domains.

5.1 Benefits of Co-designing AI via Matchmaking

5.1.1 Summary of Co-Designed Ideas for Fact-checking. In our study, fact-checkers generated new design ideas that address various needs that arise across the fact-checking process. These ideas either satisfy novel needs that were not mentioned in prior ethnographic studies or offer design solutions for existing needs documented in prior work, which existing tools or NLP research do not address.

The majority of co-designed ideas assisted in information searching, processing, and writing tasks for efficient and personalized fact-checking. Much of existing NLP research focuses on supporting claim selection and automatic claim verification. In our study, fact-checkers also generated an idea related to claim selection (Idea 5), highlighting important yet missing selection criteria that these tools must accommodate. Other co-designed ideas covered tasks in the fact-checking workflow beyond claim selection. Fact-checkers wanted to leverage NLP to identify broader disinformation narratives (Idea 2), to dynamically monitor the credibility of social media users (Idea 3), to find and provide context for ambiguous claims (Idea 4) so that fact-checkers can process information more efficiently and make a judgment on the credibility of the claim themselves. Writing the fact-check briefs and reports was another important area where fact-checkers thought NLP (and human-AI teaming) could make their work efficient (Ideas 8-11).

Our co-design session also revealed needs to support diverse goals of fact-checkers beyond conducting the fact-checking task. They wanted help in forecasting disinformation (Idea 1) so that they could proactively prepare resources to combat future misinformation. They also wanted assistance in monitoring their own potential biases in claim selection (Idea 6). Finally, because

much of fact-checking work is now done by a group of fact-checkers, they thought that NLP could help them quickly generate fact-check briefs for internal discussion on what misinformation to prioritize and focus on (Idea 7). Additionally, we learned that fact-checkers' needs and perceived importance of the co-design ideas might differ depending on their regional focus and organizational interests. We describe these exploratory findings in Appendix D.

5.1.2 Matchmaking Addressing the HCI Translational Gap. HCI scholars have discussed the importance of producing translational work to bridge the gap between research and practice [28–30, 96]. Our study is an example of translational HCI practice for fact-checking. Matchmaking for AI translates knowledge from previous empirical fact-checking research [45, 62, 89] to establish AI design requirements, and imagines useful fact-checking tools by adopting state-of-the-art AI techniques. By doing so, we address the misalignment between the practices of AI-based fact-checking driven by technological challenges and fact-checker needs, practices, and values. Our results encourage building, human-centered AI tools that are helpful to fact-checkers.

As noted by both Colusso et al. [30] and Churchill [28], ethnographic research in HCI produces theories and thematic understandings to explain human practices and problems of using technologies. This knowledge often fails to describe needs with sufficient detail required for design. Additionally, practitioners rarely have time to read ethnographic studies and are often unfamiliar with academic theories, concepts, and terminology. Norman [96] concludes that practitioners want to understand: 1) what stakeholders actually need; and 2) how the existing tools can be improved for their intended use, or adapted to serve different use cases which they were originally intended for.

Matchmaking for AI helps stakeholders formulate “need-based statements” (written at the beginning of each idea) to be useful for practitioners. In Step (1) (see Section 3.1.1), participants are empowered to articulate what needs to be done to meet their desired outcomes and criteria specialized in each separated fact-checking task. Thus, these “need-based statements” have clear goals, reducing potential design conflicts for practitioners. Additionally, because we conduct Step (1) as a mapping activity by guiding participants in small steps to “construct and express deeper levels of knowledge as their experiences” [137], participants help uncover new needs (summarized in Section 5.1.1) that are not directly identified from the prior work.

In Step (3) (see Section 3.1.3), participants prioritize the most important needs to be addressed and co-design ideas with NLP researchers to address these needs. First, they brainstorm add-on functionalities to improve existing tools for their intended use. For example, in Idea 5 (Section 4.5), participants brainstorm concrete functions that enable them to flexibly filter and select claims on existing claim selection tools. Also, they update general annotation schemes [75] of claim detection based on organizational interest. Additionally, they point out new use cases that both designers and NLP researchers have not thought of, inspiring new adoption of AI techniques, such as using fake news generators to forecast disinformation ahead of an event (Idea 1, Section 4.1) and using common narrative patterns as queries to retrieve previously checked or group unchecked claims (Idea 2, Section 4.2).

In summary, our work complements Colusso et al. [30]’s translational model by employing Matchmaking for AI as an example of translational practice in fact-checking. We produce design ideas with different levels of maturity, providing a “north star” to guide practitioners for practical adoption.

5.1.3 Implications of Co-designed Ideas on NLP Research. Interdisciplinary research across HCI and NLP has shown the importance of incorporating human feedback into AI development and evaluation [83, 141]. Previous studies [32, 77, 130] also describe that empirical evaluations of human-AI decision-making require more fine-grained metrics to measure AI performance (e.g.,

accuracy, efficiency, and calibrated trust). These metrics should be framed by stakeholder domain expertise so mixed-initiative AI systems can be more reliably measured [20]. In this study, we provide specific and nuanced definitions of these metrics, which could either improve model development or evaluation. For example, if AI researchers want to build models to forecast social media engagements for online misinformation, they need to carefully consider what it means for a social media post to be *viral*. How can we, as researchers and developers, quantify a *viral* post and allow customization for different contexts? In Appendix A, we describe four unique definitions of virality mentioned by different participants, including virality between social media vs. traditional media, attention raised by public figures, and sensationalized or hyperbolic language that grabs public attention or is easily misinterpreted. These details of often abstract and ill-defined values help developers refine the annotation scheme for labelling viral claims across different media environments and operationalize metrics to test AI performance between in-distribution and out-of-distribution datasets.

5.2 Efficacy of Co-designing AI via Matchmaking

5.2.1 Fostering Effective Communications with NLP Experts for Idea Brainstorming. While prior co-design produces high-level design guidelines or imaginary AI solutions, their outcomes often lack concrete design suggestions and relevant technical details. By bringing NLP researchers into stakeholder conversations, matchmaking constructs a progressive idea brainstorming flow for participants to: 1) explore the maximum capabilities of AI; 2) align AI capabilities and limitations with human criteria; and 3) propose applicable design recommendations.

First, participants brainstormed new AI possibilities with a better understanding of what existing AI offers. For example, some participants who had used off-the-shelf claim detection tools noticed high AI accuracy in determining whether a sentence constitutes a claim. NLP researchers explained that because a claim syntax can be well defined with objective annotations, a claim detection model gains a high accuracy by training on a gold dataset. Thus, participants (P3, 8, and 14) imagined NLP tools might further identify claims that are politically aggressive, provoking emotions like fear or calls to action. They presented exemplar claims containing these attributes and discussed with NLP researchers how to annotate these semantic accurately (results in Appendix A). Additionally, while experimenting with the AI Probe reviewing the argumentation quality of fact-checks, P9 compared it with the existing commercial writing assistants. He proposed designing a similar AI assistant to generate a meta-review of the fact-checking percentage that successfully addressed different ambiguous aspects of a claim and how persuasive the fact-check explains evidence (see Idea 11, Section 4.11).

Meanwhile, by knowing where and why AI might possibly fail, participants align AI capabilities and limitations with human criteria. For example, many participants (P8, 20, and 21) discussed with NLP researchers how to align AI decision thresholds and errors of claim check-worthiness with humans. As previously mentioned in Idea 5 (Section 4.5), identifying a single, most check-worthy claim from a large volume of claims requires detailed examinations on different selection criteria. However, participants reported that results from existing claim selection tools often fail to meet their standard. This dissatisfaction might be attributed to the issues of trust and the lack of AI explainability. Because selecting which claims to check first is a critical task for fact-checkers, AI decision-making process should be explained so that they can discern how and why certain claims were prioritized over others [32]. This process helps promote appropriate trust by aligning AI predictions with user expectations [79, 87].

During the workshop, participants suggested several design recommendations, including the use of AI explanations, to foster human-AI appropriate trust. For example, while iteratively testing different inputs of the AI Probe detecting social media check-worthy claims, P20 proposed to

set a human-AI decision threshold by aligning the model decision and its error rate with human judgement. Only when the AI predicted score is over a minimal human-AI threshold would he look at AI-suggested claims. Similar to P20, P21 said once fact-checkers have used the tool enough to set their own threshold, such as the highest threshold for worthiness and unworthiness, the AI advice would become more effective and useful for humans. Additionally, participants (P8, 14) brainstormed interactive functions to adopt feature-based AI explanations for Idea 5 (Section 4.5). After learning that certain parts of the claim exert a greater influence in the final AI prediction, P8 proposed to rank claims and highlight important parts of the claim to explain its ranking: "I think ranking them is even better, then you could highlight important texts explaining why something has been given a certain priority over other claims. Potential harm [is a priority] because it's [highlighted as] a health claim."

5.2.2 Implications of Stakeholder Algorithm Aversion on AI Co-design. During our matchmaking, many fact-checkers initially exhibited algorithm aversion [35, 81]. We found that having participants experiment with AI Probes and communicate with NLP researchers helped reduce their aversion and restore a constructive flow of idea brainstorming.

Observed levels of algorithm aversion varied significantly across participants. For example, many participants (P2, 4, 6, 11, 12, 15, 16) believed that algorithm accuracy is inferior to human accuracy, so human oversight would be imperative when using AI. While talking about the AI Probe that robo-checks quantitative claims, P12 wanted to further verify those statistics himself: "I just want to make sure it wasn't pulling numbers from inaccurate parts of whatever the document is". P13 also highlighted that AI checking the writing quality of fact-checks is not comparable to editors' eyes: "even if it looks promising, you [still] need a human eye in the form of a human sub-editor." Some participants (P5, 11, 16) were more hesitant to use AI because they were aware of certain AI biases. For example, P11 said "machine learning is the process of continuous learning by the computer itself, so what we put into the process [datasets] results in a certain bias." Additionally, some participants (P5, 8, 9) expressed stronger concerns about over-relying on AI. P5 clarified, "the machine can support you, but you should not completely rely on it."

Although participants exhibited some algorithm aversion, we observed a changing perspective from matchmaking Step (1) to Step (2). For example, P5 who was previously averse, expressed his willingness to use AI after experimenting with the AI Probes. In Step (1), he said, "Humans can be biased, and I'm aware [of that], but I think what's more [severe] is algorithm bias." After P5 saw example outputs from the AI Probes, he thought AI could perform certain tasks accurately, such as capturing claim syntax (e.g., classifying if the text is a claim, an opinion, or a prediction) and identifying quantitative information. Similar to P5, P16 was previously hesitant to have AI monitoring claims because he thought AI is "rigid" and only humans can flexibly determine what is check-worthiness based on the local news context. But he was interested in the AI Probe checking the argumentation quality of a fact-check: "I think this can be trustworthy. But when using this tool, we still need editors to touch on these copies."

As more HCI research explores AI co-design, we believe that participant aversion to AI could be either a barrier or an elixir, influencing the quality of co-design outcomes. Although this aversion might lower participant initial interest, framing it as a design problem can, in turn, elicit participation. As suggested by Hou and Jung [56], the transition from algorithm aversion to appreciation requires researchers to provide *Expert Power* for users (i.e., the result of being more knowledgeable, competent, and knowing what better action to take). As shown in our ideas and participant design experiences, Matchmaking for AI helps participants cultivate such Expert Power by engaging them in active conversations with AI experts. This reflects the traditional human-centered design principle held by the design community that user trust in technology

should be mutually communicated and agreed on between users (who are served by the technology) and developers (who create the technology), rather than singly produced or transferred from either side [96, 111, 126].

5.2.3 Implications of Stakeholder-centered AI Design in Fact-checking. Recently, there is an increasing trend in HCI and design research to use stakeholder-centered AI approaches. The goal is to more effectively tailor AI development to meet the specific needs of various domain-specific stakeholders. In our study, we focus on a specific group of stakeholders – fact-checkers. Due to their varied backgrounds, the ideas we co-designed with them display a range of AI literacy and acceptance levels. We acknowledge that this diversity influenced by their cultural perceptions and sensitivities towards AI might lead to uneven adoption rates for these ideas in the future. To enhance this, we suggest future stakeholder-centered AI design for fact-checking could be developed in two directions: 1) employing an open-sourced design strategy to build AI infrastructure and practice; and 2) using an iterative design process to customize AI solutions, taking into account their unique cultural contexts. We describe these two directions as follows.

Compared with other stakeholder-centered AI studies that focus on designing AI tools for gig-workers (e.g., Uber drivers [152], food courier [86]), public services providers (e.g., frontline workers [76]), or other domain-specific specialists (e.g., clinicians [71]), we point out that there is a unique opportunity for fact-checkers to develop common AI infrastructure around the globe based on their international network and collaborative practice. In our study, participants frequently mentioned that the existing International Fact-checking Network (IFCN) helps establish a code of conduct for a standard fact-checking workflow. Additionally, some of the existing computational infrastructure was also grounded by this collaborative work, such as the CoronaVirus check database¹⁵ and the ClaimReview tagging system¹⁶. Thus, co-design ideas generated from our study could be further developed as similar open-source AI infrastructure.

In our study, *Matchmaking for AI* is employed through a linear design process to gather fact-checker initial AI needs and ideas. However, design in practice is often non-linear, involving iterative reframing of both the design goals and the solutions through a deeper comprehension of practical problems [158]. We argue that, bearing this design nature in mind, our design ideas could spur diverse AI products by tailoring them into fact-checker unique cultural contexts.

5.3 Limitations and Improvements for Future Matchmaking for AI

In this section, we discuss limitations of *Matchmaking for AI*. For example, our protocol requires considerable time. We designed a three-step process, aiming to understand participants work practices, have them understand AI, and brainstorm ideas based on this understanding. However, during the final brainstorming session, we had time to develop around two ideas only. We may not have captured all useful ideas due to this time constraint. We posit that if co-design participants have more cross-domain knowledge, more time could be allocated to the brainstorming session to design more workable ideas.

In our matchmaking step, we could have communicated upfront that the goal of the brainstorming session was to design novel ideas that were both attuned to their needs and not already addressed in prior NLP literature. This indicates that co-designers should have helped participants understand what was and wasn't novel. Not doing so may have led to some design ideas that were less novel (see ideas in Appendix C). Similarly, we may have allowed too much co-design time to be spent discussing ideas that were not novel.

¹⁵<https://www.poynter.org/ifcn-covid-19-misinformation/>

¹⁶<https://www.claimreviewproject.com/>

There was also a missed opportunity for participants to review and prioritize these ideas as part of the matchmaking process, since participants were not interviewed again afterward. Moreover, some ideas received more time and attention than others, especially when multiple participants mentioned them. This led to an uneven level of design details for the ideas. Future work could take a more structured approach to alleviate this issue by a) gathering as many ideas as possible in the beginning, b) presenting all ideas to participants for importance rating, and then c) conducting the existing matchmaking Step (3).

Running remote co-design workshops enabled us to engage with fact-checkers from around the world, and to more easily recruit busy professionals with limited availability. However, unlike traditional workshops wherein a group of participants brainstorm ideas together through sketching, mapping, or other design activities, the remote setting made it more challenging for participants to engage interpersonally and express their creativity. In future work, it would thus be interesting to explore *Matchmaking for AI* in group and/or in-person settings, which might also benefit researchers in analyzing participant ideas with a regional focus.

While we successfully developed a set of workable NLP-based ideas, our explicit focus on NLP reflects our own research bias. Due to this, we knowingly steered participants more toward ideas related to NLP-based approaches. This may have led, for example, to omitting ideas related to multi-modal content, which is important to today's fact-checking landscape. In addition, we have employed proactive strategies to closely align the study with journalistic values. However, we recognize that despite our best efforts, there may be gaps between our research and practical challenges in applying AI to journalism. We observed the presence of algorithmic aversion in participants, especially around the ethical concerns of incorporating AI solutions into practice (described in Section 5.2). For example, participants were averse to using AI solutions for tasks that might heavily involve ethical and moral judgments. We emphasize that there is a need for an ongoing dialogue between AI researchers and journalism professionals to address the complexities of applying AI in fact-checking while maintaining human agency.

Inspired by recent advances in LLMs, we envision that many future AI Probes could be implemented via LLMs. During our co-design studies, we improvised using GPT-3.5 as an additional probe for one of our final participants (P21). LLMs support a variety of useful NLP tasks with minimal burden for modeling and data acquisition [16, 19]. During our co-design workshops, P21 and facilitators used GPT to generate a misinformation template (Idea 1, Section 4.1) by specifying a news event in a “zero-shot” setup [16]. Experimenting with GPT was a starting point for this participant to describe more concrete design elements of a feasible AI solution without knowing extensive technical details. This suggests that LLMs could naturally serve as AI Probes that help non-AI domain experts envision new AI applications in their own contexts. For this reason, we believe that adopting LLMs is a promising direction to co-develop feasible AI solutions in future co-design research.

6 Conclusion

To better understand the needs of professional fact-checkers and opportunities for AI-based assistive tooling, we investigated a novel form of co-design that brought together a tripartite group of fact-checkers, designers, and AI researchers. In particular, we proposed *Matchmaking for AI*, extending Bly and Churchill [18]’s earlier *Matchmaking* concept to an AI co-design process. This co-design method seeks to accelerate the effective translation of state-of-the-art research into practice, to better support and benefit stakeholder or practitioners lacking AI expertise. Also, co-designed ideas from our workshops exemplify how fact-checker needs can or should be addressed by AI technology, and we uncover new needs not identified in prior empirical studies. Meanwhile, co-designed ideas offer a “north star” to guide future NLP-based fact-checking research toward technology development with

greater potential for adoption by professional fact-checkers. Additionally, by providing detailed descriptions of the matchmaking process, as well as our reflections on its efficacy, limitations and future improvements, we illustrate the broader potential of our co-design approach for other domains beyond fact-checking.

Acknowledgements

We appreciate the valuable writing feedback provided by the reviewers and the research insights from our external collaborators, including Tamar Wilner, Scott Hale, and David Corney. We also thank our participants, without whom our research would not be possible. This research was partially supported by the following: the Knight Foundation; the National Science Foundation IIS-1939606, DGE-2125858 grants; Good Systems¹⁷, a UT Austin Grand Challenge for developing responsible AI technologies; and UT Austin's School of Information. The statements made herein are solely the opinions of the authors and do not reflect the views of the sponsoring agencies.

References

- [1] Bill Adair. 2020. Squash report card: Improvements during State of the Union ... and how humans will make our AI smarter - Duke Reporters' Lab. <https://reporterslab.org/squash-report-card-improvements-during-state-of-the-union-and-how-humans-will-make-our-ai-smarter/>
- [2] Firoj Alam, Stefano Cresci, Tanmoy Chakraborty, Fabrizio Silvestri, Dimiter Dimitrov, Giovanni Da San Martino, Shaden Shaar, Hamed Firooz, and Preslav Nakov. 2022. A Survey on Multimodal Disinformation Detection. In *Proceedings of the 29th International Conference on Computational Linguistics*, Nicoletta Calzolari, Chu-Ren Huang, Hansaem Kim, James Pustejovsky, Leo Wanner, Key-Sun Choi, Pum-Mo Ryu, Hsin-Hsi Chen, Lucia Donatelli, Heng Ji, Sadao Kurohashi, Patrizia Paggio, Nianwen Xue, Seokhwan Kim, Younggyun Hahm, Zhong He, Tony Kyungil Lee, Enrico Santus, Francis Bond, and Seung-Hoon Na (Eds.). International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 6625–6643.
- [3] Firoj Alam, Fahim Dalvi, Shaden Shaar, Nadir Durrani, Hamdy Mubarak, Alex Nikolov, Giovanni Da San Martino, Ahmed Abdelali, Hassan Sajjad, Kareem Darwish, and Preslav Nakov. 2021. Fighting the COVID-19 Infodemic in Social Media: A Holistic Perspective and a Call to Arms. *Proceedings of the International AAAI Conference on Web and Social Media* 15, 1 (22 May 2021), 913–922. <https://doi.org/10.1609/icwsm.v15i1.18114>
- [4] Firoj Alam, Shaden Shaar, Fahim Dalvi, Hassan Sajjad, Alex Nikolov, Hamdy Mubarak, Giovanni Da San Martino, Ahmed Abdelali, Nadir Durrani, Kareem Darwish, Abdulaziz Al-Homaid, Wajdi Zaghouani, Tommaso Caselli, Gijs Danoe, Friso Stolk, Britt Bruntink, and Preslav Nakov. 2021. Fighting the COVID-19 Infodemic: Modeling the Perspective of Journalists, Fact-Checkers, Social Media Platforms, Policy Makers, and the Society. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-Tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 611–649. <https://doi.org/10.18653/v1/2021.findings-emnlp.56>
- [5] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (7 Nov. 2019), 1–30. <https://doi.org/10.1145/3359301>
- [6] Erik Arakelyan, Arnav Arora, and Isabelle Augenstein. 2023. Topic-Guided Sampling For Data-Efficient Multi-Domain Stance Detection. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 13448–13464. <https://doi.org/10.18653/v1/2023.acl-long.752>
- [7] Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of One, Many: Using Language Models to Simulate Human Samples. *Political analysis: an annual publication of the Methodology Section of the American Political Science Association* 31, 3 (July 2023), 337–351. <https://doi.org/10.1017/pan.2023.2>
- [8] Phoebe Arnold. 2020. The challenges of online fact checking: how technology can (and can't) help - Full Fact. <https://fullfact.org/blog/2020/dec/the-challenges-of-online-fact-checking-how-technology-can-and-cant-help/>. (Accessed on 09/11/2021).
- [9] Minja Axelsson, Raquel Oliveira, Mattia Racca, and Ville Kyrki. 2021. Social Robot Co-Design Canvases: A Participatory Design Framework. *J. Hum.-Robot Interact.* 11, 1 (18 Oct. 2021), 1–39. <https://doi.org/10.1145/3472225>

¹⁷<https://goodsystems.utexas.edu>

- [10] Oana Balalau, Simon Ebel, Théo Galizzi, Ioana Manolescu, Quentin Massonnat, Antoine Deiana, Emilie Gautreau, Antoine Krempf, Thomas Pontillon, Gérald Roux, and Others. 2022. Fact-checking Multidimensional Statistic Claims in French. In *TTO 2022-Truth and Trust Online*.
- [11] Krisztian Balog, Yi Fang, Maarten de Rijke, Pavel Serdyukov, and Luo Si. 2012. Expertise Retrieval. *Foundations and Trends® in Information Retrieval* 6, 2–3 (2012), 127–256. <https://doi.org/10.1561/15000000024>
- [12] Alberto Barrón-Cedeño, Firoj Alam, Tommaso Caselli, Giovanni Da San Martino, Tamer Elsayed, Andrea Galassi, Fatima Haouari, Federico Ruggeri, Julia Maria Struß, Rabindra Nath Nandi, Gullal S Cheema, Dilshod Azizov, and Preslav Nakov. 2023. The CLEF-2023 CheckThat! Lab: Checkworthiness, Subjectivity, Political Bias, Factuality, and Authority. In *Advances in Information Retrieval*. Springer Nature Switzerland, 506–517. https://doi.org/10.1007/978-3-031-28241-6_59
- [13] Irma Becerra-Fernandez. 2006. Searching for experts on the Web: A review of contemporary expertise locator systems. *ACM Transactions on Internet Technology* 6, 4 (1 Nov. 2006), 333–355. <https://doi.org/10.1145/1183463.1183464>
- [14] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M Vardoulakis. 2020. A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '20*). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376718>
- [15] Andrew Beers, Melinda McClure Haughey, Ahmer Arif, and Kate Starbird. 2020. Examining the digital toolsets of journalists reporting on disinformation. In *Proceedings of the 2020 Computation+ Journalism Symposium* (Boston, MA, USA)(*C+ J 2020*).
- [16] Iz Beltagy, Arman Cohan, Robert Logan, IV, Sewon Min, and Sameer Singh. 2022. Zero- and Few-Shot NLP with Pretrained Language Models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, Luciana Benotti, Naoaki Okazaki, Yves Scherrer, and Marcos Zampieri (Eds.). Association for Computational Linguistics, Dublin, Ireland, 32–37. <https://doi.org/10.18653/v1/2022.acl-tutorials.6>
- [17] Kasturi Bhattacharjee, Miguel Ballesteros, Rishita Anubhai, Smaranda Muresan, Jie Ma, Faisal Ladhak, and Yaser Al-Onaizan. 2020. To BERT or Not to BERT: Comparing Task-specific and Task-agnostic Semi-Supervised Approaches for Sequence Tagging. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 7927–7934. <https://doi.org/10.18653/v1/2020.emnlp-main.636>
- [18] Sara Bly and Elizabeth F Churchill. 1999. Design through matchmaking: technology in search of users. *Interactions* 6, 2 (1 March 1999), 23–31. <https://doi.org/10.1145/296165.296174>
- [19] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [20] Carrie J Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. “Hello AI”: Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of the ACM on human-computer interaction* 3, CSCW (7 Nov. 2019), 1–24. <https://doi.org/10.1145/3359206>
- [21] Juan Cao, Peng Qi, Qiang Sheng, Tianyun Yang, Junbo Guo, and Jintao Li. 2020. Exploring the Role of Visual Content in Fake News Detection. In *Disinformation, Misinformation, and Fake News in Social Media: Emerging Research Challenges and Opportunities*, Kai Shu, Suhang Wang, Dongwon Lee, and Huan Liu (Eds.). Springer International Publishing, Cham, 141–161. https://doi.org/10.1007/978-3-030-42699-6_8
- [22] Santiago Castro, Naihao Deng, Pingxuan Huang, Mihai Burzo, and Rada Mihalcea. 2022. In-the-Wild Video Question Answering. In *Proceedings of the 29th International Conference on Computational Linguistics*. 5613–5635.
- [23] Christopher Chapp, Paul Roback, Kendra Johnson-Tesch, Adrian Rossing, and Jack Werner. 2019. Going Vague: Ambiguity and Avoidance in Online Political Messaging. *Social science computer review* 37, 5 (1 Oct. 2019), 591–610. <https://doi.org/10.1177/0894439318791168>
- [24] Jifan Chen, Grace Kim, Aniruddh Sriram, Greg Durrett, and Eunsol Choi. 2024. Complex Claim Verification with Evidence Retrieved in the Wild. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 3569–3587. <https://doi.org/10.18653/v1/2024.naacl-long.196>
- [25] Jifan Chen, Aniruddh Sriram, Eunsol Choi, and Greg Durrett. 2022. Generating Literal and Implied Subquestions to Fact-check Complex Claims. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 3495–3516. <https://doi.org/10.18653/v1/2022.emnlp-main.229>
- [26] Wei-Fan Chen, Khalid Al Khatib, Henning Wachsmuth, and Benno Stein. 2020. Analyzing Political Bias and Unfairness in News Articles at Different Levels of Granularity. In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, David Bamman, Dirk Hovy, David Jurgens, Brendan O'Connor, and

- Svitlana Volkova (Eds.). Association for Computational Linguistics, Online, 149–154. <https://doi.org/10.18653/v1/2020.nlpccs-1.16>
- [27] Eunsol Choi, Jennimaria Palomaki, Matthew Lamm, Tom Kwiatkowski, Dipanjan Das, and Michael Collins. 2021. Decontextualization: Making sentences stand-alone. *Transactions of the Association for Computational Linguistics* 9 (2021), 447–461. https://doi.org/10.1162/tacl_a_00377
- [28] Elizabeth F Churchill. 2020. HCI and UX as translational research. *Interactions* 27, 5 (1 Sept. 2020), 22–23. <https://doi.org/10.1145/3417108>
- [29] Lucas Colusso, Cynthia L Bennett, Gary Hsieh, and Sean A Munson. 2017. Translational Resources: Reducing the Gap Between Academic Research and HCI Practice. In *Proceedings of the 2017 Conference on Designing Interactive Systems* (Edinburgh, United Kingdom) (DIS '17). Association for Computing Machinery, New York, NY, USA, 957–968. <https://doi.org/10.1145/3064663.3064667>
- [30] Lucas Colusso, Ridley Jones, Sean A Munson, and Gary Hsieh. 2019. A Translational Science Model for HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19, Paper 1). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3290605.3300231>
- [31] David Corney. 2021. Towards a common definition of claim matching. <https://fullfact.org/blog/2021/oct/towards-common-definition-claim-matching/>. Accessed: 2022-9-12.
- [32] Anubrata Das, Houjiang Liu, Venelin Kovatchev, and Matthew Lease. 2023. The state of human-centered NLP technology for fact-checking. *Information processing & management* 60, 2 (1 March 2023), 103219. <https://doi.org/10.1016/j.ipm.2022.103219>
- [33] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond “Add Diverse Stakeholders and Stir”. (1 Nov. 2021). arXiv:2111.01122 [cs.AI]
- [34] Gianluca Demartini, Stefano Mizzaro, and Damiano Spina. 2020. Human-in-the-loop Artificial Intelligence for Fighting Online Misinformation: Challenges and Opportunities. *IEEE Data Eng. Bull.* 43, 3 (2020), 65–74.
- [35] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of experimental psychology. General* 144, 1 (Feb. 2015), 114–126. <https://doi.org/10.1037/xge0000033>
- [36] Michael Dobbs. 2012. *The rise of political fact-checking: How Reagan inspired a journalistic movement: A reporter's eye view*. Technical Report.
- [37] Ensheng Dong, Hongru Du, and Lauren Gardner. 2020. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet infectious diseases* 20, 5 (May 2020), 533–534. [https://doi.org/10.1016/S1473-3099\(20\)30120-1](https://doi.org/10.1016/S1473-3099(20)30120-1)
- [38] Karen M Douglas. 2021. COVID-19 conspiracy theories. *Group processes & intergroup relations: GPIR* 24, 2 (1 Feb. 2021), 270–275.
- [39] Karen M Douglas, Joseph E Uscinski, Robbie M Sutton, Aleksandra Cichocka, Turkay Nefes, Chee Siang Ang, and Farzin Deravi. 2019. Understanding conspiracy theories. *Political psychology* 40, S1 (Feb. 2019), 3–35. <https://doi.org/10.1111/pops.12568>
- [40] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 278–288. <https://doi.org/10.1145/3025453.3025739>
- [41] Angela Fan, Aleksandra Piktus, Fabio Petroni, Guillaume Wenzek, Marzieh Saeidi, Andreas Vlachos, Antoine Bordes, and Sebastian Riedel. 2020. Generating Fact Checking Briefs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 7147–7161. <https://doi.org/10.18653/v1/2020.emnlp-main.580>
- [42] Daniel Funke. 2017. In a step toward automation, Full Fact has built a live fact-checking prototype. <https://www.poynter.org/fact-checking/2017/in-a-step-toward-automation-full-fact-has-built-a-live-fact-checking-prototype/>. Accessed: 2023-6-15.
- [43] Anastasia Giachanou, Paolo Rosso, and Fabio Crestani. 2019. Leveraging Emotional Signals for Credibility Detection. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR '19). Association for Computing Machinery, New York, NY, USA, 877–880. <https://doi.org/10.1145/3331184.3331285>
- [44] D Graves. 2018. Understanding the promise and limits of automated fact-checking. (2018). <https://doi.org/10.60625/risj-nqnx-bg89>
- [45] Lucas Graves. 2017. Anatomy of a Fact Check: Objective Practice and the Contested Epistemology of Fact Checking. *Communication, Culture and Critique* 10, 3 (27 Oct. 2017), 518–537. <https://doi.org/10.1111/cccr.12163>
- [46] Lucas Graves. 2018. Boundaries Not Drawn. *Journalism Studies* 19, 5 (4 April 2018), 613–631. <https://doi.org/10.1080/1461670X.2016.1196602>

- [47] Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics* 10 (2022), 178–206.
- [48] Soumyajit Gupta, Sooyong Lee, Maria De-Arteaga, and Matthew Lease. 2023. Same Same, But Different: Conditional Multi-Task Learning for Demographic-Specific Toxicity Detection. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW '23). Association for Computing Machinery, New York, NY, USA, 3689–3700. <https://doi.org/10.1145/3543507.3583290>
- [49] Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2022. A Survey on Stance Detection for Mis- and Disinformation Identification. In *Findings of the Association for Computational Linguistics: NAACL 2022*, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 1259–1277. <https://doi.org/10.18653/v1/2022.findings-naacl.94>
- [50] Christina Harrington and Tawanna R Dillahunt. 2021. Eliciting Tech Futures Among Black Young Adults: A Case Study of Remote Speculative Co-Design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21, Article 397). Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3411764.3445723>
- [51] Naemul Hassan, Gensheng Zhang, Fatma Arslan, Josue Caraballo, Damian Jimenez, Siddhant Gawsane, Shohedul Hasan, Minumol Joseph, Aaditya Kulkarni, Anil Kumar Nayak, Vikas Sable, Chengkai Li, and Mark Tremayne. 2017. ClaimBuster: the first-ever end-to-end fact-checking system. *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases* 10, 12 (1 Aug. 2017), 1945–1948. <https://doi.org/10.14778/3137765.3137815>
- [52] Chris Hays, Zachary Schutzman, Manish Raghavan, Erin Walk, and Philipp Zimmer. 2023. Simplistic Collection and Labeling Practices Limit the Utility of Benchmark Datasets for Twitter Bot Detection. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW '23). Association for Computing Machinery, New York, NY, USA, 3660–3669. <https://doi.org/10.1145/3543507.3583214>
- [53] Tuan-Anh Hoang and Ee-Peng Lim. 2012. Virality and Susceptibility in Information Diffusions. *Proceedings of the International AAAI Conference on Web and Social Media* 6, 1 (2012), 146–153. <https://doi.org/10.1609/icwsm.v6i1.14245>
- [54] Robert R Hoffman, Nigel R Shadbolt, A Mike Burton, and Gary Klein. 1995. Eliciting Knowledge from Experts: A Methodological Analysis. *Organizational behavior and human decision processes* 62, 2 (1 May 1995), 129–158. <https://doi.org/10.1006/obhd.1995.1039>
- [55] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) (CHI '99). Association for Computing Machinery, New York, NY, USA, 159–166. <https://doi.org/10.1145/302979.303030>
- [56] Yoyo Tsung-Yu Hou and Malte F Jung. 2021. Who is the Expert? Reconciling Algorithm Aversion and Algorithm Appreciation in AI-Supported Decision Making. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (18 Oct. 2021), 1–25. <https://doi.org/10.1145/3479864>
- [57] Kung-Hsiang Huang, Kathleen McKeown, Preslav Nakov, Yejin Choi, and Heng Ji. 2023. Faking Fake News for Real Fake News Detection: Propaganda-Loaded Training Data Generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 14571–14589. <https://doi.org/10.18653/v1/2023.acl-long.815>
- [58] Caroline Jack. 2017. Lexicon of lies: Terms for problematic information. *Data & Society* 3, 22 (2017), 1094–1096.
- [59] Shan Jiang and Christo Wilson. 2021. Structurizing Misinformation Stories via Rationalizing Fact-Checks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 617–631. <https://doi.org/10.18653/v1/2021.acl-long.51>
- [60] Xiaoneng Jin, Mark Evans, Hua Dong, and Anqi Yao. 2021. Design Heuristics for Artificial Intelligence: Inspirational Design Stimuli for Supporting UX Designers in Generating AI-Powered Ideas. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '21, Article 219). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3411763.3451727>
- [61] Yohan Jo, Seojin Bang, Emaad Manzoor, Eduard Hovy, and Chris Reed. 2020. Detecting Attackable Sentences in Arguments. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Online). Association for Computational Linguistics, Stroudsburg, PA, USA, 1–23. <https://doi.org/10.18653/v1/2020.emnlp-main.1>
- [62] Purna Juneja and Tanushree Mitra. 2022. Human and Technological Infrastructures of Fact-checking. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (11 Nov. 2022), 1–36. <https://doi.org/10.1145/3555143>
- [63] Ehsan Kamaloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. 2023. HAGRID: A Human-LLM Collaborative Dataset for Generative Information-Seeking with Attribution. (July 2023). arXiv:2307.16883 [cs.CL]
- [64] Georgios Karagiannis, Mohammed Saeed, Paolo Papotti, and Immanuel Trummer. 2020. Scrutinizer: fact checking statistical claims. *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases* 13, 12 (1 Aug.

- 2020), 2965–2968. <https://doi.org/10.14778/3415478.3415520>
- [65] Yavuz Selim Kartal, Busra Guvenen, and Mucahid Kutlu. 2020. Too Many Claims to Fact-Check: Prioritizing Political Claims Based on Check-Worthiness. (April 2020). arXiv:2004.08166 [cs.CL]
- [66] Ashkan Kazemi, Kiran Garimella, Devin Gaffney, and Scott Hale. 2021. Claim Matching Beyond English to Scale Global Fact-Checking. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 4504–4517. <https://doi.org/10.18653/v1/2021.acl-long.347>
- [67] Ashkan Kazemi, Zehua Li, Verónica Pérez-Rosas, and Rada Mihalcea. 2021. Extractive and Abstractive Explanations for Fact-Checking and Evaluation of News. In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*, Anna Feldman, Giovanni Da San Martino, Chris Leberknight, and Preslav Nakov (Eds.). Association for Computational Linguistics, Online, 45–50. <https://doi.org/10.18653/v1/2021.nlp4if-1.7>
- [68] Finn Kensing and Joan Greenbaum. 2012. Heritage: Having a say. In *Routledge international handbook of participatory design*. Routledge, 41–56.
- [69] Hyun Suk Kim, Yoo Ji Suh, Eun-Mee Kim, Eunryung Chong, Hwajung Hong, Boyoung Song, Yena Ko, and Ji Soo Choi. 2022. Fact-Checking and Audience Engagement: A Study of Content Analysis and Audience Behavioral Data of Fact-Checking Coverage from News Media. *Digital Journalism* 10, 5 (28 May 2022), 781–800. <https://doi.org/10.1080/21670811.2021.2006073>
- [70] Junsol Kim and Byungkyu Lee. 2023. AI-Augmented Surveys: Leveraging Large Language Models and Surveys for Opinion Prediction. (16 May 2023). arXiv:2305.09620 [cs.CL]
- [71] Minjung Kim, Saebyeol Kim, Jinwoo Kim, Tae-Jin Song, and Yuyoung Kim. 2024. Do stakeholder needs differ? - Designing stakeholder-tailored Explainable Artificial Intelligence (XAI) interfaces. *International journal of human-computer studies* 181 (2024), 103160. <https://doi.org/10.1016/j.ijhcs.2023.103160>
- [72] Philipp Koehn and Rebecca Knowles. 2017. Six Challenges for Neural Machine Translation. In *Proceedings of the First Workshop on Neural Machine Translation*. Association for Computational Linguistics, Vancouver, 28–39. <https://doi.org/10.18653/v1/W17-3204>
- [73] Ajay Kohli and Saurabh Jha. 2018. Why CAD Failed in Mammography. *Journal of the American College of Radiology: JACR* 15, 3 Pt B (March 2018), 535–537. <https://doi.org/10.1016/j.jacr.2017.12.029>
- [74] Tomoko Komatsu, Marisela Gutierrez Lopez, Stephann Makri, Colin Porlezza, Glenda Cooper, Andrew MacFarlane, and Sondess Missaoui. 2020. AI should embody our values: Investigating journalistic values to inform AI technology design. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) (NordCHI '20, Article 11). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3419249.3420105>
- [75] Lev Konstantinovskiy, Oliver Price, Mevan Babakar, and Arkaitz Zubiaga. 2021. Toward Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection. *Digital Threats* 2, 2 (15 April 2021), 1–16. <https://doi.org/10.1145/3412869>
- [76] Tzu-Sheng Kuo, Hong Shen, Jisoo Geum, Nev Jones, Jason I Hong, Haiyi Zhu, and Kenneth Holstein. 2023. Understanding Frontline Workers' and Unhoused Individuals' Perspectives on AI Used in Homeless Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23, Article 860). Association for Computing Machinery, New York, NY, USA, 1–17. <https://doi.org/10.1145/3544548.3580882>
- [77] Vivian Lai, Chacha Chen, Q Vera Liao, Alison Smith-Renner, and Chenhao Tan. 2021. Towards a Science of Human-AI Decision Making: A Survey of Empirical Studies. (Dec. 2021). arXiv:2112.11471 [cs.AI]
- [78] Matt Lease. 2020. Designing human-AI partnerships to combat misinformation.
- [79] John D Lee and Katrina A See. 2004. Trust in automation: designing for appropriate reliance. *Human factors* 46, 1 (2004), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- [80] Mina Lee, Percy Liang, and Qian Yang. 2022. CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22, Article 388). Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3491102.3502030>
- [81] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (1 Jan. 2018), 2053951718756684. <https://doi.org/10.1177/2053951718756684>
- [82] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D Procaccia. 2019. WeBuildAI: Participatory Framework for Algorithmic Governance. *Proc. ACM Hum.-Comput. Interact* 3, CSCW (7 Nov. 2019), 1–35. <https://doi.org/10.1145/3359283>
- [83] Piyawat Lertvittayakumjorn and Francesca Toni. 2021. Explanation-based human debugging of NLP models: A survey. *Transactions of the Association for Computational Linguistics* 9 (30 Dec. 2021), 1508–1528. <https://doi.org/10.1162/tacl.2021.9.0.00000>

tacl_a_00440

- [84] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [85] Tongliang Li, Lei Fang, Jian-Guang Lou, Zhoujun Li, and Dongmei Zhang. 2021. AnaSearch: Extract, Retrieve and Visualize Structured Results from Unstructured Text for Analytical Queries. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (Virtual Event, Israel) (WSDM '21). Association for Computing Machinery, New York, NY, USA, 906–909. <https://doi.org/10.1145/3437963.3441694>
- [86] Shuhao Ma, Paulo Bala, Valentina Nisi, John Zimmerman, and Nuno Jardim Nunes. 2023. Uncovering Gig Worker-Centered Design Opportunities in Food Delivery Work. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) (DIS '23). Association for Computing Machinery, New York, NY, USA, 688–701. <https://doi.org/10.1145/3563657.3596123>
- [87] Shuai Ma, Ying Lei, Xinru Wang, Chengbo Zheng, Chuhan Shi, Ming Yin, and Xiaojuan Ma. 2023. Who Should I Trust: AI or Myself? Leveraging Human and AI Correctness Likelihood to Promote Appropriate Trust in AI-Assisted Decision-Making. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23, Article 759). Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3544548.3581058>
- [88] Alexandre Maros, Jussara M Almeida, and Marisa Vasconcelos. 2021. A Study of Misinformation in Audio Messages Shared in WhatsApp Groups. In *Disinformation in Open Online Media*. Springer International Publishing, 85–100. https://doi.org/10.1007/978-3-030-87031-7_6
- [89] Nicholas Micallef, Vivienne Armacost, Nasir Memon, and Sameer Patil. 2022. True or False: Studying the Work Practices of Professional Fact-Checkers. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1 (7 April 2022), 1–44. <https://doi.org/10.1145/3512974>
- [90] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) (FAT* '19). Association for Computing Machinery, New York, NY, USA, 220–229. <https://doi.org/10.1145/3287560.3287596>
- [91] Preslav Nakov, Alberto Barrón-Cedeño, Giovanni Da San Martino, Firoj Alam, Julia Maria Struß, Thomas Mandl, Rubén Míguez, Tommaso Caselli, Mucahid Kutlu, Wajdi Zaghouani, Chengkai Li, Shaden Shaar, Gautam Kishore Shahi, Hamdy Mubarak, Alex Nikolov, Nikolay Babulkov, Yavuz Selim Kartal, and Javier Beltrán. 2022. The CLEF-2022 CheckThat! Lab on Fighting the COVID-19 Infodemic and Fake News Detection. In *Advances in Information Retrieval*. Springer International Publishing, 416–428. https://doi.org/10.1007/978-3-030-99739-7_52
- [92] Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. Automated Fact-Checking for Assisting Human Fact-Checkers. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, 4551–4558. <https://doi.org/10.24963/ijcai.2021/619>
- [93] An T Nguyen, Aditya Kharosekar, Saumyaa Krishnan, Siddhesh Krishnan, Elizabeth Tate, Byron C Wallace, and Matthew Lease. 2018. Believe it or not: Designing a Human-AI Partnership for Mixed-Initiative Fact-Checking. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 189–199. <https://doi.org/10.1145/3242587.3242666>
- [94] Thanh Tam Nguyen, Matthias Weidlich, Hongzhi Yin, Bolong Zheng, Quang Huy Nguyen, and Quoc Viet Hung Nguyen. 2020. FactCatch: Incremental Pay-as-You-Go Fact Checking with Minimal User Effort. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 2165–2168. <https://doi.org/10.1145/3397271.3401408>
- [95] Jian-Yun Nie. 2022. *Cross-Language Information Retrieval*. Springer Nature.
- [96] Donald A Norman. 2010. The research-practice gap: the need for translational developers. *Interactions* 17, 4 (1 July 2010), 9–12. <https://doi.org/10.1145/1806491.1806494>
- [97] Yikang Pan, Liangming Pan, Wenhui Chen, Preslav Nakov, Min-Yen Kan, and William Wang. 2023. On the Risk of Misinformation Pollution with Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 1389–1403. <https://doi.org/10.18653/v1/2023.findings-emnlp.97>
- [98] Michael Quinn Patton. 2014. *Qualitative Research & Evaluation Methods: Integrating Theory and Practice*. SAGE Publications.
- [99] Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. KILT: a Benchmark for Knowledge Intensive Language Tasks. In *Proceedings of the 2021 Conference of the North American*

- Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 2523–2544. <https://doi.org/10.18653/v1/2021.naacl-main.200>
- [100] Massimo Poesio, Juntao Yu, Silviu Paun, Abdulrahman Aloraini, Pengcheng Lu, Janosch Haber, and Derya Cokal. 2023. Computational Models of Anaphora. *Annual review of applied linguistics* 9, Volume 9, 2023 (17 Jan. 2023), 561–587. <https://doi.org/10.1146/annurev-linguistics-031120-111653>
 - [101] Ethan Porter and Thomas J Wood. 2021. The global effectiveness of fact-checking: Evidence from simultaneous experiments in Argentina, Nigeria, South Africa, and the United Kingdom. *Proceedings of the National Academy of Sciences of the United States of America* 118, 37 (14 Sept. 2021). <https://doi.org/10.1073/pnas.2104235118>
 - [102] Rob Procter, Miguel Arana Catania, Yulan He, Maria Liakata, Arkaitz Zubiaga, Elena Kochkina, and Runcong Zhao. 2023. Some Observations on Fact-Checking Work with Implications for Computational Support. In *Workshop Proceedings of the 17th International AAAI Conference on Web and Social Media* (Limassol, Cyprus). Association for the Advancement of Artificial Intelligence. <https://doi.org/10.36190/2023.28>
 - [103] Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Martha Palmer, Rebecca Hwa, and Sebastian Riedel (Eds.). Association for Computational Linguistics, Copenhagen, Denmark, 2931–2937. <https://doi.org/10.18653/v1/D17-1317>
 - [104] Adrian Rauchfleisch and Jonas Kaiser. 2020. The False positive problem of automatic bot detection in social science research. *PLoS one* 15, 10 (1 Oct. 2020), e0241045. <https://doi.org/10.1371/journal.pone.0241045>
 - [105] Sravana Reddy, Yongze Yu, Aasish Pappu, Aswin Sivaraman, Rezvaneh Rezapour, and Rosie Jones. 2021. Detecting Extraneous Content in Podcasts. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, Paola Merlo, Jorg Tiedemann, and Reut Tsarfay (Eds.). Association for Computational Linguistics, Online, 1166–1173. <https://doi.org/10.18653/v1/2021.eacl-main.99>
 - [106] Craig T Robertson, Rachel R Mourão, and Esther Thorson. 2020. Who Uses Fact-Checking Sites? The Impact of Demographics, Political Antecedents, and Media Use on Fact-Checking Site Awareness, Attitudes, and Behavior. *The International Journal of Press/Politics* 25, 2 (1 April 2020), 217–237. <https://doi.org/10.1177/1940161219898055>
 - [107] Daniel Russo, Serra Sinem Tekiroğlu, and Marco Guerini. 2023. Benchmarking the generation of fact checking explanations. *Transactions of the Association for Computational Linguistics* 11 (20 Oct. 2023), 1250–1264. https://doi.org/10.1162/tacl_a_00601
 - [108] Hohyon Ryu, Matthew Lease, and Nicholas Woodward. 2012. Finding and exploring memes in social media. In *Proceedings of the 23rd ACM conference on Hypertext and social media* (Milwaukee, Wisconsin, USA) (HT '12). Association for Computing Machinery, New York, NY, USA, 295–304. <https://doi.org/10.1145/2309996.2310044>
 - [109] Arkadiy Saakyan, Tuhin Chakrabarty, and Smaranda Muresan. 2021. COVID-Fact: Fact Extraction and Verification of Real-World Claims on COVID-19 Pandemic. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 2116–2129. <https://doi.org/10.18653/v1/2021.acl-long.165>
 - [110] Mohammed Saeed and Paolo Papotti. 2021. Fact-checking statistical claims with tables. <https://www.eurecom.fr/publication/6630/download/data-publi-6630.pdf>. Accessed: 2023-6-14.
 - [111] Elizabeth B-N Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *CoDesign* 4, 1 (1 March 2008), 5–18. <https://doi.org/10.1080/15710880701875068>
 - [112] Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose Opinions Do Language Models Reflect?. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 29971–30004.
 - [113] Mourad Sarrouiti, Asma Ben Abacha, Yassine Mrabet, and Dina Demner-Fushman. 2021. Evidence-based Fact-Checking of Health-related Claims. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-Tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 3499–3512. <https://doi.org/10.18653/v1/2021.findings-emnlp.297>
 - [114] Connie Moon Sehat, Ryan Li, Peipei Nie, Tarunima Prabhakar, and Amy X Zhang. 2024. Misinformation as a harm: Structured approaches for fact-checking prioritization. *Proceedings of the ACM on human-computer interaction* 8, CSCW1 (17 April 2024), 1–36. <https://doi.org/10.1145/3641010>
 - [115] Shaden Shaar, Nikolay Babulov, Giovanni Da San Martino, and Preslav Nakov. 2020. That is a Known Lie: Detecting Previously Fact-Checked Claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 3607–3618. <https://doi.org/10.18653/v1/2020.acl-main.332>
 - [116] Shaden Shaar, Maram Hasanain, Baya Hamdan, Zien Sheikh Ali, Fatima Haouari, Alex Nikolov, Mucahid Kutlu, Yavuz Selim Kartal, Firoj Alam, and Giovanni Da San Martino. 2021. Overview of the CLEF-2021 CheckThat! lab task

- 1 on check-worthiness estimation in tweets and political debates. In *CLEF (Working Notes)*. gcris.etu.edu.tr.
- [117] S Shaar, Maram Hasanain, Bayan Hamdan, Zien Sheikh Ali, Fatima Haouari, Alex Nikolov, Mucahid Kutlu, Yavuz Selim Kartal, Firoj Alam, Giovanni Da San Martino, Alberto Barrón-Cedeño, Rubén Míguez, Javier Beltrán, Tamer Elsayed, and Preslav Nakov. 2021. Overview of the CLEF-2021 CheckThat! Lab task 1 on check-worthiness estimation in tweets and political debates. *Conference and Labs of the Evaluation Forum* (2021), 369–392.
- [118] Gautam Kishore Shahi, Julia Maria Struß, and Thomas Mandl. 2021. Overview of the CLEF-2021 CheckThat! Lab: Task 3 on Fake News Detection.. In *CLEF (Working Notes)*. 406–423.
- [119] Chengcheng Shao, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. 2016. Hoaxy: A Platform for Tracking Online Misinformation. In *Proceedings of the 25th International Conference Companion on World Wide Web* (Montréal, Québec, Canada) (*WWW '16 Companion*). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 745–750. <https://doi.org/10.1145/2872518.2890098>
- [120] J Sheridan. 2021. Should you trust media bias charts. URL: <https://www.poynter.org/factchecking/media-literacy/2021/should-you-trust-media-bias-charts> (2021).
- [121] Ben Shneiderman. 2020. Design Lessons From AI's Two Grand Goals: Human Emulation and Useful Applications. *IEEE Transactions on Technology and Society* 1, 2 (June 2020), 73–82. <https://doi.org/10.1109/TTS.2020.2992669>
- [122] Prakhar Singh, Anubrata Das, Junyi Jessy Li, and Matthew Lease. 2021. The Case for Claim Difficulty Assessment in Automatic Fact Checking. (Sept. 2021). arXiv:2109.09689 [cs.CL]
- [123] Zoe Skinner, Stacey Brown, and Greg Walsh. 2020. Children of Color's Perceptions of Fairness in AI: An Exploration of Equitable and Inclusive Co-Design. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI EA '20*). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3382901>
- [124] Santhosh Srinivasan and Craig Fagan. 2019. Rating the disinformation risks of news domains: the Global Disinformation Index. In *TTO 2019*. truthandtrustonline.com.
- [125] Christopher Starke, Janine Baleis, Birte Keller, and Frank Marcinkowski. 2022. Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society* 9, 2 (1 July 2022), 2053951722115189. <https://doi.org/10.1177/2053951722115189>
- [126] Marc Steen. 2013. Co-design as a process of joint inquiry and imagination. *Design issues* 29, 2 (April 2013), 16–28. https://doi.org/10.1162/desi_a_00207
- [127] Mark Stencel, Erica Ryan, and Joel Luther. 2022. *Fact-checkers extend their global reach with 391 outlets, but growth has slowed*. Technical Report. Duke Reporters' Lab.
- [128] Elias Storms, Oscar Alvarado, and Luciana Monteiro-Krebs. 2022. 'Transparency is Meant for Control' and Vice Versa: Learning from Co-designing and Evaluating Algorithmic News Recommenders. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (11 Nov. 2022), 1–24. <https://doi.org/10.1145/3555130>
- [129] Victor Suarez-Lledo and Javier Alvarez-Galvez. 2021. Prevalence of Health Misinformation on Social Media: Systematic Review. *Journal of medical Internet research* 23, 1 (20 Jan. 2021), e17187. <https://doi.org/10.2196/17187>
- [130] Hariharan Subramonyam, Jane Im, Colleen Seifert, and Eytan Adar. 2022. Solving Separation-of-Concerns Problems in Collaborative Design of Human-AI Systems through Leaky Abstractions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22, Article 481*). Association for Computing Machinery, New York, NY, USA, 1–21. <https://doi.org/10.1145/3491102.3517537>
- [131] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. 2021. Towards A Process Model for Co-Creating AI Experiences. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA) (*DIS '21*). Association for Computing Machinery, New York, NY, USA, 1529–1543. <https://doi.org/10.1145/3461778.3462012>
- [132] Tara Qian Sun and Rony Medaglia. 2019. Mapping the challenges of Artificial Intelligence in the public sector: Evidence from public healthcare. *Government information quarterly* 36, 2 (1 April 2019), 368–383. <https://doi.org/10.1016/j.giq.2018.09.008>
- [133] Briony Swire-Thompson, Nicholas Miklaucic, John P Wihbey, David Lazer, and Joseph DeGutis. 2022. The backfire effect after correcting misinformation is strongly associated with reliability. *Journal of experimental psychology. General* 151, 7 (July 2022), 1655–1665. <https://doi.org/10.1037/xge0001131>
- [134] The Poynter Institute. 2021. Global Fact 8 Pre-recorded Segment 4. <https://www.youtube.com/watch?v=gOhPKDaeQxI&t=770s>
- [135] Bart van Dijk and John Zimmerman. 2019. Discovering Users for Technical Innovations through Systematic Matchmaking. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI EA '19, Paper LBW2517*). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3290607.3312843>
- [136] Slavena Vasileva, Pepa Atanasova, Lluís Màrquez, Alberto Barrón-Cedeño, and Preslav Nakov. 2019. It Takes Nine to Smell a Rat: Neural Multi-Task Learning for Check-Worthiness Prediction. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, Ruslan Mitkov and Galia Angelova

- (Eds.). INCOMA Ltd., Varna, Bulgaria, 1229–1239. https://doi.org/10.26615/978-954-452-056-4_141
- [137] Froukje Sleswijk Visser, Pieter Jan Stappers, Remko van der Lugt, and Elizabeth B-N Sanders. 2005. Contextmapping: experiences from practice. *CoDesign* 1, 2 (1 April 2005), 119–149. <https://doi.org/10.1080/15710880500135987>
- [138] Nguyen Vo and Kyumin Lee. 2019. Learning from Fact-checkers: Analysis and Generation of Fact-checking Language. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (*SIGIR'19*). Association for Computing Machinery, New York, NY, USA, 335–344. <https://doi.org/10.1145/3331184.3331248>
- [139] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 7534–7550. <https://doi.org/10.18653/v1/2020.emnlp-main.609>
- [140] Zihan Wang, Karthikeyan K, Stephen Mayhew, and Dan Roth. 2020. Extending Multilingual BERT to Low-Resource Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 2649–2656. <https://doi.org/10.18653/v1/2020.findings-emnlp.240>
- [141] Zijie J Wang, Dongjin Choi, Shenyu Xu, and Diyi Yang. 2021. Putting Humans in the Natural Language Processing Loop: A Survey. In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, Su Lin Blodgett, Michael Madaio, Brendan O'Connor, Hanna Wallach, and Qian Yang (Eds.). Association for Computational Linguistics, Online, 47–52.
- [142] Claire Wardle. [n. d.]. Misinformation Has Created a New World Disorder. *Scientific American* ([n. d.]). <https://doi.org/10.1038/scientificamerican0919-88>
- [143] Brian E Weeks. 2015. Emotions, partisanship, and misperceptions: How anger and anxiety moderate the effect of partisan bias on susceptibility to political misinformation. *The Journal of communication* 65, 4 (Aug. 2015), 699–719. <https://doi.org/10.1111/jcom.12164>
- [144] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned Language Models Are Zero-Shot Learners. (3 Sept. 2021). arXiv:2109.01652 [cs.CL]
- [145] Chao-Yuan Wu and Philipp Krähenbühl. 2021. Towards Long-Form Video Understanding. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1884–1894. <https://doi.org/10.1109/CVPR46437.2021.00192>
- [146] 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* (New Orleans, LA, USA) (*CHI '22, Article 385*). Association for Computing Machinery, New York, NY, USA, 1–22. <https://doi.org/10.1145/3491102.3517582>
- [147] Zeqiu Wu, Yi Luan, Hannah Rashkin, David Reitter, Hannaneh Hajishirzi, Mari Ostendorf, and Gaurav Singh Tomar. 2022. CONQRR: Conversational Query Rewriting for Retrieval with Reinforcement Learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 10000–10014. <https://doi.org/10.18653/v1/2022.emnlp-main.679>
- [148] Zhiping Xiao, Weiping Song, Haoyan Xu, Zhicheng Ren, and Yizhou Sun. 2020. TIMME: Twitter Ideology-detection via Multi-task Multi-relational Embedding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Virtual Event, CA, USA) (*KDD '20*). Association for Computing Machinery, New York, NY, USA, 2258–2268. <https://doi.org/10.1145/3394486.3403275>
- [149] 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 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI '19, Paper 185*). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300415>
- [150] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI '19, Paper 238*). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300468>
- [151] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supriya Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, and John Zimmerman. 2023. Creating Design Resources to Scaffold the Ideation of AI Concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) (*DIS '23*). Association for Computing Machinery, New York, NY, USA, 2326–2346. <https://doi.org/10.1145/3563657.3596058>
- [152] Angie Zhang, Alexander Boltz, Chun Wei Wang, and Min Kyung Lee. 2022. Algorithmic Management Reimagined For Workers and By Workers: Centering Worker Well-Being in Gig Work. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22, Article 14*). Association for Computing Machinery, New York, NY, USA, 1–20. <https://doi.org/10.1145/3491102.3501866>

- [153] Qiang Zhang, Jonathan Cook, and Emine Yilmaz. 2021. Detecting and Forecasting Misinformation via Temporal and Geometric Propagation Patterns. In *Advances in Information Retrieval*. Springer International Publishing, 455–462. https://doi.org/10.1007/978-3-030-72240-1_48
- [154] Shi Zhi, Yicheng Sun, Jiayi Liu, Chao Zhang, and Jiawei Han. 2017. ClaimVerif: A Real-time Claim Verification System Using the Web and Fact Databases. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (Singapore, Singapore) (*CIKM '17, Vol. Part F1318*). Association for Computing Machinery, New York, NY, USA, 2555–2558. <https://doi.org/10.1145/3132847.3133182>
- [155] Xinyi Zhou and Reza Zafarani. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. *ACM Comput. Surv.* 53, 5 (28 Sept. 2020), 1–40. <https://doi.org/10.1145/3395046>
- [156] Zhibin Zhou, Lingyun Sun, Yuyang Zhang, Xuanhui Liu, and Qing Gong. 2020. ML Lifecycle Canvas: Designing Machine Learning-Empowered UX with Material Lifecycle Thinking. *Human-Computer Interaction* 35, 5-6 (1 Nov. 2020), 362–386. <https://doi.org/10.1080/07370024.2020.1736075>
- [157] Haiyi Zhu, Bowen Yu, Aaron Halfaker, and Loren Terveen. 2018. Value-Sensitive Algorithm Design: Method, Case Study, and Lessons. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (1 Nov. 2018), 1–23. <https://doi.org/10.1145/3274463>
- [158] John Zimmerman, Aaron Steinfeld, Anthony Tomic, and Oscar J. Romero. 2022. Recentring Reframing as an RtD Contribution: The Case of Pivoting from Accessible Web Tables to a Conversational Internet. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22, Article 541*). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3491102.3517789>
- [159] Douglas Zytco, Pamela J. Wisniewski, Shion Guha, Eric P. S. Baumer, and Min Kyung Lee. 2022. Participatory Design of AI Systems: Opportunities and Challenges Across Diverse Users, Relationships, and Application Domains. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI EA '22, Article 154*). Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3491101.3516506>

A Fact-checking Criteria - Part 1

Criteria	Description (which participants voiced this criteria)
<i>Monitoring and selecting claims</i>	
News judgment	What news sources are good for finding claims to check (3, 7, 10, 18)
Newsworthiness	Breaking news that the public cares about (3, 5, 8) Trending social media posts (10, 11) Reported news by politicians or traditional media (3, 5, 11)
Fact-checkable	Whether it is a claim, not an opinion or prediction (All) Whether there's any publicly available evidence to validate (2, 3, 8, 9) If the claim or claim's terms are ambiguous, they should be well-defined (3, 4, 9, 10)
Predicted truth	Quickly evaluate whether the claim is true (2, 3, 9, 17, 19, 20) Prioritize claims that are most likely to be false (7, 9, 16, 17, 20) Primarily check false claims because organizations get paid from social media companies (9, 19) Also check claims that are likely to be true (18)
Virality	Trending claims on different social media platforms or reach threshold of engagements (1, 2, 5, 21) Claims reported by traditional media outlets (13) Claims repeated by public figures (8, 10, 13) Claims that are sensationalized, or misinterpreted, or contain hyperbolic language, grabbing public attention (1, 3, 8, 14)
Harmfulness	Claims specifically relating to medical issues (5, 6, 15, 19) Claims containing risks on finance, public welfare, people's understanding of society (19) Claims containing persuasive language that calls people to action (3) Claims targeting a specific region or group of people (5, 9, 13, 15, 16) Claims that are both viral and relevant to people's physical and social lives (1, 2, 4, 8, 11, 12, 19)
Personal Interest	Claims that are less viral but specifically target certain groups of people (5, 9, 11) Claim topics that fact-checkers are very familiar with (4, 6, 12, 13)
Organizational Interest	Politics (2, 3, 12) Region-specific (5, 9, 13, 15, 16) Global (20, 21) Topics previously with ample misinformation (4, 6, 8, 13) Claims relating to young audience (18)
Biases awareness	Intentionally try to find claims from different parties (5, 6, 8, 16) Maintaining neutrality is essentially impossible (3, 15, 20) Do not conduct themselves in any political affairs (3, 7, 19) Editorial review for political biases (5, 7) LGBTQ (11) Intentionally hire diverse groups of fact-checkers (3, 18)

Table 4. Fact-checking Criteria - Part 1

B Fact-checking Criteria - Part 2

Criteria	Description (which participants voiced this criteria)
<i>Investigating claims</i>	
Provenance	Identifying claims origins to verify claims (2, 3, 8, 13)
Reliability	More than one source for validation (15) Addition of multiple experts (10, 12, 13, 15) Checking news source bias (8, 13) Don't use secondary sources (5, 13) Sources from reputable international, non-profit, or non-partisan organizations (2, 4, 5, 13) Follow evidence pyramid for scientific facts (4, 6)
Thoroughness	Ask contextual questions to breakdown claims by internal fallacy (2, 5, 13) Define the exact wording of every term/concept (1, 2, 6, 8) Reconstruct a claim's origin (3, 6, 7, 13)
<i>Writing fact checks</i>	
Writing structure	Begins with facts, introduces claim & context, then evidence explaining the rating (2) Begins with the claim & context, then gives the verdict, provides facts explaining that verdict, concludes with a final verdict (3, 4, 6, 8, 9, 13, 15)
Readability	Simple language for readers to understand (4, 6, 8, 10, 11, 13) Definition of a concept should be well-explained (2, 8, 9)
Coherence	Clearly justify the provenance of the claim (3, 13) Readers have access to all reliable sources to replicate the check (1, 2, 8, 10, 12, 18) Enough evidence to validate the claim clearly (2, 4, 8, 9)
Collaborative or hierarchical editing	Peer review (5, 10, 11) 2 rounds of editorial review (2, 4, 6, 8, 9, 13, 15) 3 or more rounds of editorial review (19, 22)

Table 5. Fact-checking Criteria - Part 2

C Existing Needs and Designs

While our co-design process was primarily focused on generating novel ideas (Section 4), four additional ideas were generated that remain significant and merit further attention, despite earlier coverage in prior work [46, 89, 102]. For these ideas, we synthesize participant quotes and our reflections. We expect this material to further inform community understanding of this need, and to further encourage additional technical and design work to develop, implement, and deploy these ideas in practice.

Fact-checking stages	No.	Design ideas	New idea	New need	Related AI research
Monitoring claims	12	Tracing the Spread of Similar Claims across Social Media Platforms			[31, 66, 115]
	13	Matching Similar Multilingual Claims			[72, 140]
	14	Detecting and Checking Claims from Multi-modal Content			[2, 22, 105, 145]
Investigating facts	15	Recommending Claim-Specific Experts			[11, 13]

Table 6. Four additional ideas generated that remain significant and merit further attention, despite earlier coverage in prior work.

C.1 Tracing the Spread of Similar Claims across Social Media Platforms

We learned from our participants that as online misinformation spreads, it becomes warped, eventually growing to encompass a plethora of similar, yet nonidentical, claims. Our participants want to group these similar, related claims together to gain a more accurate view of “how these claims evolve [over time]” (P1) and “[to learn] the full scope of its virality on whatever platforms” (P5). They (P12, 15, 20) pointed out two key functions of tracing similar claims: 1) retrieving checked claims similar to an unchecked claim at hand; and 2) grouping similar unchecked claims across different platforms. For example, when deciding whether to pursue a new fact-check, our participants (P2, 9, 13-14, 18, 29) typically first investigate whether the claim has already been checked using rudimentary Google searches, to avoid overlap. However, participants may continue to pursue the claim despite the fact it is similar to the previously checked claims if certain criteria are met, such as: 1) it is highly relevant to their targeted readers (P2, 18); 2) it needs supplemental evidence to be fully fact-checked in a local context (P9, 13); or 3) its previous report does not meet high fact-checking standard (P18, 20)¹⁸.

This ‘provenance tracing’ tool would enable fact-checkers to not only investigate whether a claim has been checked before, but also identify its provenance, tracking its spread across multiple platforms. Regarding specific features, P14 would like to sort claims by similarity relative to the

¹⁸The high standards of a fact-check includes its high readability and coherence. More details are described in the Appendix B.

“primary” claim. P9 and P12 wanted to order claims by the time of posting, so as to be helpful in investigating a claim’s provenance.

Claim matching could be directly applied to retrieve previously fact-checked claims that are lexically or semantically similar to a new claim [31, 66, 115]. With the development of ClaimReview¹⁹ - a tagging system for existing fact-checks - it is now much easier to retrieve previously checked claims to enable this function. As misinformation often backfires, researchers can also utilize natural language inference (NLI) to auto-verify these claims based on the veracity of similar previously fact-checked claims. Fact-checkers then can conduct a quick sanity check on these automated fact-checks and prioritize more time checking other claims.

To aggregate similar unchecked claims, participants suggested that existing claim detection models could be combined with information retrieval [108, 119], each fine-tuned on a specific media type for that platform. A claim could be searched for across all of these models and the results could be aggregated. Additionally, although there have been recent developments of APIs from different media platforms (e.g., Meta²⁰ and TikTok²¹) and new cross-platform third-party retrieval tools (e.g., Tracking Exposed²²), participants requested more engineering work to combine these resources into a unified technological infrastructure for mass-monitoring checkable claims.

C.2 Matching Similar Multilingual Claims

Misinformation is often spread through multiple languages, reaching larger populations through new languages. For our participants, it is important to catch these claims quickly, however, this becomes difficult without advanced knowledge of multiple languages. Our participants in multilingual countries, such as the Philippines (P16), South Africa (P4, 6-8, 13), Nepal (P5, 10-11), India (P21), and Mexico (P17), specifically indicated the need to detect similar claims across multiple languages.

P21 requested cross-lingual information retrieval [95] to: “search [the claim] in Ukrainian and Russian at the same time, so that [they] get ...more results on the same topic.” P10 and P11 also echoed similar needs for detecting checkable claims in Nepali. Five of the eight (62.5%) participants who fact-checked multilingual claims mentioned that this should be an important feature for existing fact-checking tools. P17 reported that the claim detection tool he used (Facebook’s tool) was developed to only find claims in English and is useful about 15-20% of the time to find Spanish claims. This function could help our participants to “trace the impact” (P21) of a claim and gain a more holistic view of how the claim has spread (similar to Idea 4.3). Additionally, it would also “save [us] a ton of time” (P17) by cutting down manual human translation and conducting multiple searches in different languages. Participants (P5, 10-11) also reported that off-the-shelf translation tools lack adequate performance for low-resources languages (LRLs) [72]. For example, P10 said, when translating English claims to Nepali using Google Translate, Nepali claims are interpreted very literally and lack the language-specific context she needs. These findings support results from the previous studies [44, 92], suggesting the need for further engineering work in translating high-resources languages into LRLs.

To implement this idea for high-resource languages, automatic translation models perform relatively well and can be used. For LRLs, one approach to applying transfer learning to train LRLs from a large monolingual language corpus (e.g., English) [140]. More generally, additional dataset or model development may be needed to enable effective support across languages.

¹⁹<https://schema.org/ClaimReview>

²⁰<https://research.facebook.com/blog/2021/3/new-analytics-api-for-researchers-studying-facebook-page-data/>

²¹<https://newsroom.tiktok.com/en-sg/an-update-on-our-platform-api-for-researchers-sg>

²²<https://tracking.exposed/>

C.3 Detecting and Checking Claims from Multi-modal Content

With an increased percentage of misinformation originating from media beyond text, participants want AI to help them better check claims from audio, photo, and video content. This idea has two primary aspects: 1) detecting checkable claims from multi-modal content; and 2) checking these claims by tracing their provenance²³. We learned that identifying checkable claims from audio or video content is more challenging than from text. P18, who primarily checks claims on TikTok, mentioned video feeds on TikTok are curated based on user interests, and cannot be centrally monitored by fact-checkers. Additionally, due to the segmentation of existing monitoring tools (e.g., Facebook’s tool only monitors Meta-owned platforms, FullFact Alpha monitors news outlets, and TweetDeck monitors Tweets), our participants (P5, 9, 12, 14, 20-21) are forced to use multiple tools. They voiced that creating a single tool that sources claims across platforms would save time.

Because, as participants told us, claims extracted from videos or photos are often intentional misinterpretations, they usually begin their fact-checking process by tracing the provenance: finding the original photos or videos. For example, P3 said that “a video [claim] is taken out of context, and you have to watch the whole video to understand [and to check it].” To combat this, they imagined improving existing reverse video or image search with more advanced features, namely, retrieving the original video and mapping the claim to its original timestamp. Note that implementing such tools would require a combination of methods from computer vision, multimodal NLP, and information retrieval (see [21] for an overview).

To identify this multi-modal misinformation, our participants proposed that AI could first transcribe it to text, then extract claims from the resulting transcripts using NLP. Some recent work [2] has approached claim detection by combining image data, such as a screenshot from a video or an individual picture, with text data. However, directly identifying claims from videos and audio remains under-explored. Researchers may need to construct new datasets that pairs textual claims with corresponding audio and video content, both short-form (TikTok and Instagram Reels) and long-form (YouTube). Recent work in video question answering [22], long-form video understanding [145], and detecting content patterns in podcasts [105] may be explored in order to develop new methods for this task. Additionally, mass-monitoring multi-modal content also requires infrastructure work, such as cross-referencing information from different social media platforms.

C.4 Recommending Claim-Specific Experts

A tool that can recommend “experts” for verification of findings related to a claim is helpful in the fact-checking process (P1-2, 5, 8-10, 15). Organizations tend to maintain a database of experts. Our participants mentioned that they either need to manually search through such a database or identify new experts. Sometimes it is difficult to fact-check a claim without access to an expert “there are many claims that we find it very difficult to fact check [since we don’t have access to experts]” (P10). Implementing a focused information retrieval technique like *expert retrieval* [11, 13] can be investigated in the context of fact-checking to address this issue.

D Fact-checker Needs and Co-Designed Ideas influenced by Regions and Organizational Interests

By mapping participant criteria (Appendix A and B) to their demographics and organizations, we uncovered potential regional patterns that appear to influence participant needs and the co-creation of ideas. Nonetheless, these findings are preliminary and necessitate further systematic validation by subsequent research. We report them below.

²³this idea refers to two fact-checking tasks: monitoring claims and investigating facts.

First, we found that participants from Philippines, South Africa, Nepal, and New Zealand, predominantly concentrated on verifying local misinformation, indicating a strong need for personalization shared across different ideas. For example, participants (P5, 9, 11) emphasized the significance of scrutinizing regionally targeted misinformation. Because this misinformation could inflict direct harm upon local communities even if its virality is limited. On the other hand, participants working in international fact-checking organizations (P20, 21), although geographically distributed, engaged in global fake news by collaborating with international colleagues where the claimants or subjects of fake news are located. This variance in regional focus, whether local or global, indicates a common interest in tailoring claim filtering and selection (Idea 5, Section 4.5) and in examining the demographic composition of the readership for published fact-checks (Idea 9, Section 4.9).

Also, the difference in regional political discourse and cultural context tends to shape what and how fact-checkers check. For example, P3 reported that because politicians in the country have learned to modify their tactics to circumvent fact-checking scrutiny, sometimes they also check vague claims from politicians. This situation requires them to use more nuanced ratings to assess the truthfulness of such claims. Consequently, she expressed a strong interest in finding and providing context for ambiguous claims (Idea 4, Section 4.4). This idea is beneficial not only for highlighting inaccuracies in claims and pinpointing arguments requiring further investigation, but it also acts as a tool to aid novice fact-checkers in cultivating critical analytical skills, a vital competency for professional fact-checkers.

Furthermore, regional difference highlights a language challenge in the fact-checking process. Participants from multilingual countries such as Mexico, Philippines, South Africa, Nepal, and India emphasized multilingual support. This need is particularly relevant in tracking the spread of claims (Idea 12, Section C.1) and matching new claims with previously fact-checked claims (Idea 13, Section C.2).

Beyond regional differences, the affiliations of different fact-checking organizations influence the selection criteria for fact-checkers. While a majority of participants focused on debunking false news, few also check potentially true stories. Reported by editor participants (P7, 19), checking fake news is financially beneficial for their organizations, particularly in collaboration with mainstream social media platforms like Facebook, helping them identify online hoaxes, fake news, or propaganda. Conversely, P18 (who is primarily engaged in video fact-checking for cable news) mentioned a preference for choosing stories based on viewer interest, regardless of their perceived veracity. These differences highlight the need for multifaceted features in tools designed for claim filtering and selection (Idea 5, Section 4.5). Additionally, there is a recognized need to refine existing annotation schema and techniques for claim detection to adapt to diverse contexts of claims-making.

The editorial process for writing fact-checks also varies among organizations. Smaller entities with limited resources, like those represented by participants (P5, 10, and 11), rely on peer-review to edit their fact-check reports. In contrast, larger and more established organizations often implement a more rigorous process, involving two or three rounds of editing. Despite these differences, participants from both types of organizations were interested in using AI assistance to critique and edit fact-check reports (Idea 11, Section 4.11). This use of AI is seen as a way to compensate for limited editorial resources or to enhance the efficiency and productivity of the editing process.

Received July 2023; revised January 2024; accepted March 2024