

Small-world Networks of Agents Brainstorm AI Risks

Ke Zhou^{1,2} Edyta Bogucka¹, Daniele Quercia^{1,3}

¹Nokia Bell Labs, UK

²University of Nottingham, UK

³Politecnico di Torino, Italy

ke.zhou@nokia-bell-labs.com, edyta.bogucka@nokia-bell-labs.com, quercia@cantab.net

Abstract

The ideation phase of participatory AI risk assessment often starts with a blank slate or a limited list of predefined risks, making it difficult to surface indirect or systemic harms. To address this limitation, we propose a three-stage tool to support the ideation phase of AI risk assessment. The tool complements participatory AI, rather than replacing it, and helps focus later engagement with affected communities. First, it dynamically discovers stakeholders depending on the given AI use and recursively expanding outward, allowing overlooked or indirect stakeholders to emerge. Second, it simulates these stakeholders with LLMs (creating in-silico stakeholders), connecting them into a network of a given topology, and having them ideate about risks. Third, it prioritizes risks using network centrality measures (the centrality of in-silico stakeholders in the network). In an initial evaluation, we found that betweenness centrality run through agents connected in a small-world network works best as it elevates risks raised by stakeholders who bridge disconnected groups, surfacing novel, systemic harms that traditional methods often miss. On an AI chatbot companion use case, this approach increased the novelty of the identified risks by approximately 1.1 points over single LLM brainstorming, and by 0.5 points over agentic LLM brainstorming, measured on a normalized five-point Likert scale, without reducing the plausibility or severity of the identified risks. To test whether our framework helps a human-led ideation session using the Futures-Wheel approach, we divided 11 teams of non-western young chatbot users into two types: control (team) and treatment (team) in a participatory AI risk assessment. The control teams started from a list of risks generated by the 45 AI practitioners in the initial evaluation; the treatment teams started from a list generated by our framework. The treatment teams identified more risks overall, and more systemic, human-computer interaction, and environmental risks in particular, categories that practitioners typically find harder to spot.

Project — <https://social-dynamics.net/ai-risks/small-world>

1 Introduction


As AI systems are increasingly integrated into critical societal infrastructure, from healthcare diagnostics to companion chatbots for vulnerable populations, the imperative for comprehensive impact and risk assessment has transitioned

from a theoretical ideal to a governance necessity. A critical but often under-examined stage in this process is *risk ideation*: the initial step in which participants attempt to enumerate as many relevant risks as possible before prioritization and mitigation. Because all downstream analysis depends on this initial pool, the quality and breadth of ideation largely determine the effectiveness of the entire assessment. In practice, however, ideation frequently begins from a blank slate or a limited, organizer-defined list of risks. This creates a structural bottleneck. While frameworks like the EU AI Act (European Commission 2026) emphasize the need to catch “systemic risks”, participants tend to anchor on familiar, high-frequency risks (e.g., privacy or security), while overlooking indirect, systemic, or stakeholder-specific harms (Lancaster et al. 2024; Xia et al. 2023; Fröhling et al. 2026). This limitation is not merely a matter of scale, but of epistemic diversity: risks that emerge from the interaction between stakeholders, or that disproportionately affect marginalized groups, are systematically underrepresented (Birhane et al. 2022; Lancaster et al. 2024).


Traditional human workshops offer depth but are unscalable and constrained by the “bounded rationality” of small groups of participants (Diehl and Stroebe 1987). Conversely, automated tools that use Large Language Models (LLMs) to generate risk lists, such as Farsight (Wang et al. 2024), ExploreGen (Herdal et al. 2024), and AHA! (Bućinca et al. 2023), act as single expert oracles. They tend to converge on “consensus” risks: those that are plausible and severe, but often generic. Consequently, they struggle to explore the long tail of socio-technical failures that only emerge through the friction of conflicting perspectives, i.e., risks that often carry systemic consequences (Fröhling et al. 2026) (e.g., developing an unhealthy over-reliance on the chatbot that hinders real-life social relationship development). In short, current methods are good at confirming what we already know (high likelihood), but less effective at discovering what we do not (high novelty). As a result, they fail to adequately support the human ideation process.

To address this, we propose an ideation tool designed specifically to overcome the “blank slate” bottleneck in AI risk assessment. The tool seeds brainstorming workshops with a network of LLM agents, surfacing more diverse and systemic risks than human sessions seeded with practitioner-generated risks. We make three main contributions:

A FRAMEWORK FOR GENERATING RISKS



B EVALUATION OF THE FRAMEWORK FOR GENERATING RISKS



C EVALUATION OF THE FRAMEWORK FOR SUPPORTING IDEATION IN PARTICIPATORY AI RISK ASSESSMENT




Figure 1: **Overview of the proposed framework and its two evaluations.** (A) Our three-stage framework for generating risks. It involves identifying diverse stakeholders dynamically (Stage 1); building a stakeholder of in-silico agents network and generating risks by orchestrating the communication through a network with a given topology (Stage 2); and prioritizing risks using three centrality metrics: degree, betweenness, and closeness (Stage 3). (B) Evaluation of the framework’s risk generation quality by comparing system-generated risks with practitioners-generated risks from workshops. (C) Evaluation of the framework as a support tool for participatory AI risk assessment. Participants are assigned to either a control condition (report with practitioners-generated risks) or a treatment condition (report with network-generated risks), and use the provided reports as starting points for Futures Wheel ideation.

(1) *A Three-Stage Framework for Networked Stakeholder Simulation (Section 3):* We propose a framework that dynamically discovers stakeholders via recursive snowballing, instantiates them as in-silico LLM agents, and connects them into a network topology through which they communicate to brainstorm risks. Risks are then prioritized using network centrality measures applied to the resulting stakeholder graph. We show small-world topologies combined with betweenness centrality are uniquely effective at surfacing novel, systemic harms that flat ideation methods miss.

(2) *Evaluation of Risk Generation Quality (Section 4):* We evaluate our framework against single-LLM and multi-agent baselines using a dataset of 244 risks generated by 45 AI practitioners in blank-slate brainstorming workshops. Our approach improves novelty by approximately 1.1 points over single-LLM brainstorming, and 0.5 points over agentic LLM baselines, without reducing plausibility or severity.

(3) *Evaluation of Support for Human-Led Participatory*

AI Risk Assessment (Section 5): We conduct a controlled Futures Wheel study (Glenn 2009) with 40 non-Western young chatbot users. Treatment teams seeded with our framework’s risks identified substantially more risks overall, and significantly more systemic, human-computer interaction, and socioeconomic harms than control teams seeded with practitioner-generated risks.

2 Related Work

2.1 AI Risk Identification and Assessment

AI risk identification and AI risk assessment are distinct but related processes: AI risk identification uncovers potential harms, while risk assessment quantifies them for prioritization and decision-making. AI risk identification has evolved from narrowly defined technical safety evaluations toward broader socio-technical analyses. Early approaches surfaced technical risks at the model level, such as robustness fail-

ures and security vulnerabilities, and quantified them using metrics including error rates, perturbation robustness, and adversarial vulnerability (Schmitz et al. 2025; Steimers and Schneider 2022). As AI systems became embedded in social and institutional contexts, research attention shifted toward identifying and quantifying downstream harms.

AI risk taxonomies have played a central role in this shift. Influential work by Weidinger et al. (2022), alongside subsequent syntheses (Uuk et al. 2024; Zhang et al. 2025), provides comprehensive catalogues of risks spanning technical failures, human–AI interaction risks, and long-term systemic societal risks. These taxonomies offer shared vocabulary and conceptual clarity, but they are largely retrospective: they systematize known risks rather than supporting the anticipation of novel, context-specific risks in emerging applications. Moreover, they typically abstract away from the perspectives of distinct stakeholders.

Participatory approaches attempt to address these limitations by incorporating diverse human perspectives into risk identification (Delgado et al. 2023; Lee et al. 2019). Methods such as the *Judgment Call* card game (Ballard, Chappell, and Kennedy 2019), participatory scenario planning (Kieslich, Helberger, and Diakopoulos 2026), and end users auditing (Deng et al. 2025) leverage structured interaction to surface values, concerns, and experiential knowledge that may otherwise remain implicit. While effective in eliciting rich insights, these approaches are resource-intensive, difficult to scale, and constrained by practical barriers to participation. In particular, assembling representative groups of end users, regulators, and marginalized stakeholders is often infeasible for organizations conducting AI risk assessments. In addition, current participatory AI literature (Kallina, Bohné, and Singh 2025; Birhane et al. 2022; Corbett, Denton, and Erete 2023; Delgado et al. 2023) overlooks the “blank slate” bottleneck in risk ideation, where forcing participants to brainstorm from scratch exhausts their cognitive bandwidth on obvious issues (Hohma et al. 2023).

Current AI risk assessment practices are strongly influenced by regulatory frameworks (European Commission 2026), and predominantly focus on evaluating identified risks in terms of likelihood and severity, an approach inherited from risk assessment traditions in domains such as safety engineering and finance. These dimensions are also prominent evaluation criteria in prior research (Wang et al. 2024). By contrast, novelty and diversity of perspectives have rarely been treated as explicit assessment objectives (Herdel et al. 2024), and there remains a lack of robust measures for how they should be defined and operationalized.

2.2 AI-Augmented Brainstorming for Risk Assessment

The use of AI systems to support brainstorming and ideation has gained increasing attention (Wang et al. 2025; Ashkinaze et al. 2025b; Choi et al. 2024). Empirical studies show that LLMs can substantially increase the quantity of generated ideas (Bouschery, Blazevic, and Piller 2024; Yu-Han and Chun-Ching 2023). However, these gains are frequently accompanied by reductions in diversity and originality, with models exhibiting fixation on stereotypical or

high-frequency patterns (Wadinambiarachchi et al. 2024; Meincke, Nave, and Terwiesch 2025; AlDahoul, Rahwan, and Zaki 2025; Bai et al. 2025). This tendency is particularly salient in risk assessment contexts. For example, when prompted to identify “AI risks”, LLMs often default to generic concerns such as privacy or bias, overlooking subtle, domain-specific, or interactional harms (Lancaster et al. 2024; Xia et al. 2023; Fröhling et al. 2026).

To scale risk discovery, the industry has increasingly turned to AI red teaming and AI-augmented brainstorming. Red teaming leverages adversarial testing to probe models for vulnerabilities and policy violations (Ge et al. 2024; Ganguli et al. 2022). While automated red teaming is highly effective at identifying technical exploits (e.g., jailbreaks), it fundamentally relies on predefined threat models and often struggles to anticipate nuanced, socio-technical harms that only emerge through complex stakeholder interactions. To broaden the scope of risk identification beyond red teaming, several tools leverage LLMs to assist AI risk assessment directly. Systems such as Farsight (Wang et al. 2024), ExploreGen (Herdel et al. 2024), and AHA! (Buçinca et al. 2023) generate or visualize candidate risks based on system descriptions. While these tools improve efficiency, they typically treat the LLM as a single expert oracle. As a result, they do not explicitly model interaction, disagreement, or perspective-taking among stakeholders, mechanisms that are known to underpin diversity in human group brainstorming (Meincke, Nave, and Terwiesch 2025).

2.3 Stakeholder Analysis and Network Science

Recent work has explored the use of LLMs to simulate aspects of human behavior and social interaction, often described as *in-silico* social science. The Plurals framework (Ashkinaze et al. 2025a) demonstrates that LLM agents initialized with distinct roles or personas can participate in structured deliberative processes, producing outputs that differ meaningfully from single-agent prompting. Similarly, Park et al.’s *Generative Agents* (Park et al. 2023) show that agents endowed with memory and social context can exhibit emergent social dynamics. Within AI risk identification, however, prior applications of multi-agent LLM systems have largely relied on unstructured debate or flat ideation structures. Fukumura and Ito (2025) explored agent-supported brainstorming, and found that benefits were primarily limited to general or non-specialized topics. These findings suggest that role differentiation alone may be insufficient; the structure of interaction itself plays a critical role in shaping outcomes.

Stakeholder theory posits that organizational outcomes depend on the relationships among stakeholders, not solely on individual attributes (Parmar et al. 2010). Network science provides formal tools for analyzing these relationships. Rowley (Rowley 1997) argues that stakeholder influence is a function of network position such as centrality or brokerage rather than intrinsic importance alone. In adjacent domains, network measures (Borgatti, Jones, and Everett 1998; Watts and Strogatz 1998) have been used to identify influential actors, information bottlenecks, and structural vulnerabilities. However, these methods have rarely been applied to in-silico

brainstorming or assessment processes. Existing approaches typically aggregate risks by frequency or expert judgment, implicitly assuming all contributors are equally informative.

Research Gap. The initial ideation phase often starts from a “blank slate” or a narrow set of predefined risks. This forces participants to spend time listing well-known concerns before they can address broader, systemic risks. To address this gap, we introduce a tool that helps generate early ideas about AI risks by simulating discussions among in-silico stakeholders. The tool structures these interactions as networks, allowing us to explore how communication patterns shape the risks they identify. This approach surfaces risks from a wider range of perspectives, especially from stakeholders who “bridge” otherwise separate groups.

3 Our Framework

Our framework operationalizes AI risk brainstorming through network-mediated ideation, a process where in-silico agents representing diverse stakeholders generate risks within a network topology. Rather than treating risk identification as an isolated ideation task, we conceptualize it as a collaborative activity shaped by who participates, how they are connected, and whose perspectives bridge across roles.

Figure 1 provides an overview of the methodology, consisting of three stages: (1) Discovering stakeholders dynamically (Section 3.1); (2) Building a stakeholder network (Section 3.2); and (3) Prioritizing risks using network centrality (Section 3.3). Formally, given an AI use case, our goal is to produce a set of top- K risks $\mathcal{R} = \{r_1, r_2, \dots, r_K\}$ on a ranked list, where each risk r_k is associated with one or more stakeholders.

3.1 Stage 1: Identifying Stakeholders Dynamically

Creating An Initial Set of Stakeholders. We begin by identifying an initial set of stakeholders relevant to the AI use case U : $\mathcal{S}_0 = \{s_1, s_2, \dots, s_m\}$, where each s_i corresponds to an abstract role rather than a specific individual.

We define a *typical stakeholder set* grounded in responsible AI practice and governance literature (San Vito et al. 2025). This set includes five roles: *Developers*, responsible for building and maintaining the AI system; *Designers*, shaping user interaction and experience; *Decision-makers*, responsible for deployment and organizational strategy; *End users*, who directly interact with or are affected by the system; and *Regulators*, responsible for oversight and compliance. These roles align with prior work on AI auditing and accountability (San Vito et al. 2025), which emphasizes the inclusion of decision subjects, domain experts, and regulatory actors in assessment processes. For AI companion chatbots in particular, these categories capture both technical and socio-emotional dimensions of risk.

We also define a *nuanced stakeholder set* based on NIST standard framework for AI risk assessments (AI 2023), including AI developers, machine learning engineers, data scientists, data collectors and curators, system architects, model trainers, and technical auditors; product and project

managers, system owners, organizational leadership, enterprise risk managers, legal counsel, compliance officers, procurement teams, and IT deployment and operations staff; clinicians, judges, HR professionals, and analysts; external regulators, supervisory authorities, policymakers, and legislators; and academic researchers, ethics experts, journalists and environmental stakeholders.

Making The Set of Stakeholders Comprehensive Via Snowballing. To capture stakeholders omitted by the initial abstraction, we perform iterative stakeholder expansion by employing an LLM-driven snowballing technique for recommending more relevant stakeholders recursively. At iteration t , each stakeholder $s_i \in \mathcal{S}_t$ proposes a set of additional stakeholders $\Delta\mathcal{S}_i^{(t)} = \{s_{i1}, s_{i2}, \dots\}$, where we query the LLM: “Given the AI use case [Description] and your stakeholder role as [Stakeholder], list what other stakeholders (no personal names) are significantly impacted by or have influence over this system?” The candidate set for iteration $t + 1$ is: $\mathcal{S}_{t+1} = \mathcal{S}_t \cup_i \Delta\mathcal{S}_i^{(t)}$.

To prevent unbounded growth, we deduplicate roles using semantic similarity. Let $\phi(s)$ denote an embedding of a stakeholder description. A newly proposed stakeholder s' is merged with an existing stakeholder s if: $\cos(\phi(s'), \phi(s)) \geq \tau$. Expansion terminates when either there are no novel stakeholders are added ($\mathcal{S}_{t+1} = \mathcal{S}_t$), or a maximum number of iterations $T_{\max}$ is reached.

3.2 Stage 2: Building a Stakeholder Network and Generating Risks

Connecting Stakeholders In A Network Within A Given Topology. We define a stakeholder interaction network as a directed weighted graph: $G = (V, E, W)$, where: $V = \mathcal{S}_T$ is the final stakeholder set, $E \subseteq V \times V$ represents recommendation of stakeholders indicating influences, $W : E \rightarrow \mathbb{R}^+$ assigns edge weights. An edge (s_i, s_j) exists, if stakeholder s_j is recommended by stakeholder s_i , and therefore deemed to maintain influence over stakeholder s_j during risk ideation. Using this stakeholder set, we generate multiple instantiations of G with different network topologies:

Hierarchical Network: It is constructed by assigning stakeholders to levels $\ell(s) \in \mathbb{N}$ based on organizational authority. Edges are constrained such that: $\ell(s_i) < \ell(s_j) \Rightarrow (s_i \rightarrow s_j)$. This reflects top-down flow typical of formal governance structures.

Small-World Network: It is designed to balance dense local interaction among closely related stakeholder roles with a limited number of long-range connections across the network. We construct a small-world-inspired directed acyclic graph by preserving strong high-weight edges within stakeholder clusters identified through the recommendation, while enforcing acyclicity. Formally, an edge is retained between stakeholders $s_i, s_j \in V$, if one recommends the other, and the edge does not form circles. This structure supports locally coherent ideation, while enabling risks to propagate across stakeholder communities through bridging roles.

Random Network: In the random baseline, edges are sampled uniformly at random from $V \times V$, subject to acyclicity and fixed edge count $|E|$.

Pruning Network Graph and Removing Cycles. Given the directed interaction structures, we transform G into a directed weighted graph G' by performing backbone extraction, and graph cycle removal procedures.

Backbone Extraction. The initial stakeholder interaction graph $G = (V, E, W)$ may contain a dense set of edges reflecting weak, redundant, or noisy relationships. We apply a disparity filter (Serrano, Boguná, and Vespignani 2009) to remove statistically insignificant edges, retaining only edges whose normalized weight exceeds a significance threshold δ . We set $\delta = 0.05$ following prior work on backbone extraction in weighted networks (Serrano, Boguná, and Vespignani 2009). Applying the disparity filter yields a pruned graph: $G_b = (V, E_b, W_b)$, where $E_b \subseteq E$ contains only statistically significant edges. This backbone graph retains the majority of total edge weight, while substantially improving interpretability and computational efficiency.

Graph Cycle Removal. The network-mediated ideation structures are required to be directed acyclic graphs (DAGs) in order to support multi-round, feed-forward information flow. However, the backbone graph G_b may still contain cycles, reflecting reciprocal or mutual influence among stakeholders. To enforce acyclicity, we therefore transform G_b into a DAG G' using a greedy maximum-weight acyclic subgraph construction. Given a directed weighted graph $G_b = (V, E_b, W_b)$, we seek a directed acyclic subgraph $G' = (V, E', W')$ such that: $E' \subseteq E_b$, and the total retained weight $\sum_{(i,j) \in E'} w_{ij}$ is maximized. This problem is NP-hard in general. We therefore adopt an efficient greedy approximation: sorting remaining edges by descending weight, and greedily adding them to G' unless doing so would create a cycle. This algorithm prioritizes preserving high-weight influence relationships while eliminating weaker reciprocal ties that would violate acyclicity. As a result, G' approximates the most influential pathways within the stakeholder system. To assess robustness, we verified that alternative cycle-breaking strategies (e.g., random edge removal) produce qualitatively similar results in downstream risk diversity and ranking, though with slightly higher variance.

Generating Risks By Orchestrating The Communication Among Stakeholders. Each stakeholder $s_i \in V$ is instantiated as an in-silico agent a_i , parameterized by: $a_i = \langle s_i, U, \mathcal{I}_i^{(t)} \rangle$, where $\mathcal{I}_i^{(t)}$ denotes information available at round t for the use case U . Stakeholder agents generate risks via conditional language modeling: $r_{ik}^{(t)} \sim p_\theta(r \mid s_i, U, \mathcal{I}_i^{(t)})$, where θ denotes the underlying LLM parameters. At each round t , stakeholder agent a_i receives inputs from its immediate predecessors: $\mathcal{I}_i^{(t)} = \bigcup_{(s_j \rightarrow s_i) \in E'} \mathcal{R}_j^{(t-1)}$, where $\mathcal{R}_j^{(t-1)}$ is the set of risks generated by stakeholder agent a_j in the previous round. This induces a structured dependency between stakeholder agents' outputs, operationalizing network-mediated ideation.

3.3 Stage 3: Prioritizing Risks Using Network Centrality

Let $G' = (V, E')$ be the final stakeholder graph on which we compute three network measures. A core hypothesis of the

work is that *who* identifies a risk matters as much as *what* the risk is. We derive three primary rankings for prioritizing risks based on these network centrality measures:

Degree Centrality Ranking. Risks are ranked by the sum of the Degree Centrality of the stakeholders who identified them: $Score_D(R) = \sum_{s \in Contributors(R)} C_{deg}(s)$. This favors risks identified by “popular” or most obvious stakeholders (e.g., Developers, Product Managers).

Betweenness Centrality Ranking. Risks are ranked by the sum of the Betweenness Centrality of the stakeholders: $Score_B(R) = \sum_{s \in Contributors(R)} C_{bet}(s)$. Betweenness centrality (C_{bet}) measures how often a node acts as a bridge along the shortest path between two other nodes. In our context, high-betweenness stakeholders are hypothesized to see risks that arise from stakeholders that bridge disconnected groups (e.g., “Ethics & Compliance Officer” sitting between “Regulatory Body” and “Developer”).

Closeness Centrality Ranking. Risks are ranked by the sum of the Closeness Centrality of the stakeholders who identified them: $Score_C(R) = \sum_{s \in Contributors(R)} C_{clo}(s)$. Closeness centrality (C_{clo}) captures how close a stakeholder is, on average, to all other stakeholders in the network, based on shortest path distances. This ranking strategy prioritizes risks raised by stakeholders who are well-positioned to rapidly integrate information across the network, favoring broadly informed perspectives rather than highly popular (degree) or structurally bridging (betweenness) roles.

In summary, these three ranking strategies reflect distinct theories of influence in collective risk assessment: degree centrality emphasizes visibility and consensus, betweenness centrality highlights cross-role mediation and systemic insight, and closeness centrality prioritizes globally informed perspectives.

Next, we empirically evaluate our proposed framework on an AI chatbot companion use case. We selected this specific use case because it represents a high-stakes, socio-emotional application characterized by wide demographic use and timely regulatory attention. We evaluate our framework in two ways. First, we evaluate our framework with blank-slate of risks by AI practitioners (Section 4). Second, we directly test our framework’s effectiveness in supporting human-led ideation in Participatory Risk Assessments through a controlled study using the Futures Wheel technique (Glenn 2009) (Section 5), in which non-Western young chatbot users who started from risks generated by our framework are compared to those starting from practitioners-generated risks in our first evaluation (i.e., by AI practitioners).

4 Evaluation of Our Framework for Generating Risks

To validate the efficacy of our framework, we evaluate whether networked simulation improves risk generation compared to flat baselines without network structures, using a dataset of practitioners-generated risks as a benchmark.

4.1 Establishing Brainstorming Workshops for Generating Practitioners-based Risks

To evaluate our framework, we drew on a dataset from 8 human-based brainstorming workshops focused on identifying risks associated with AI chatbot companions (Govers, Šćepanović, and Quercia 2026). Rather than representing an idealized participatory engagement, these workshops were designed to mimic the typical, early-stage internal risk workshops commonly conducted within public institutions and private companies. We recruited 45 AI practitioners across 3 locations via mailing lists, online ads, and EventBrite. All participants reported active personal use of AI chatbot companions, grounding their risk identification in everyday AI chatbot interaction. Those workshops utilized blank slate brainstorming (i.e., with no risks provided), each lasting 1 hour and 45 minutes with 4 to 7 participants. The study received institutional ethics approval and participants were compensated above minimum wage. Participants (detailed demographics shown in Appendix) were relatively young (majority aged 25–34), highly educated and all held university degrees (51% Master’s, 20% PhD). Many participants came from institutions and companies (37% were from the company, 43% researchers, and 20% students). To capture a rich socio-technical perspective, the team specifically included: (1) *AI professionals*, comprising software developers, NLP experts, AI monitoring/auditing experts, and designers, who contributed insights into technical and operational failure modes; (2) *Applied AI researchers*, who developed AI algorithms in a company and evaluated their limitations; (3) *Researchers in Responsible AI and psychology*, who specialized in the socio-emotional and downstream impacts of companions; and (4) *Civil society representatives*, who advocated for protection of vulnerable groups. Self-reported AI expertise was high, with 65% of participants reporting knowledgeable or expert levels. All workshops were balanced to include company workers, researchers, and students. Across all sessions, participants generated 244 distinct risks, which form the basis of our evaluation.

4.2 Defining the Multi-Dimensional Rubric

We adopted a multi-dimensional rubric (Table 1) to evaluate the quality of practitioners-generated risks for chatbot companion within those brainstorming workshops. Synthesizing criteria from the NIST AI Risk Management Framework and other prior AI risk assessment frameworks (AI 2023; Diakopoulos and Johnson 2021; Kieslich, Helberger, and Diakopoulos 2026; Dean et al. 2006; Lewis and Sauro 2009; Brooke et al. 1996), the rubric assesses plausibility, probability, severity, uniqueness, novelty, and usability. To explicitly distinguish between evaluation constructs, our rubric separates Plausibility (theoretical feasibility), Likelihood (realistic probability of occurrence in deployment), and Uniqueness (specificity across domains), a correlation matrix provided in Appendix demonstrates that these dimensions capture distinct evaluative signals rather than redundant measures. Five external AI-domain experts recruited via Prolific rated the risks using Likert scales, following attention and comprehension screening. They are AI ethicists

Dimension	Definition
Plausibility	Assesses whether the described risk could realistically arise in the given AI use context (1–5 Likert (Diakopoulos and Johnson 2021)).
Likelihood	Estimates the likelihood that a risk will occur (1–7 Likert, based on NIST AI risk assessment measures (AI 2023) and foresight (Fröhling et al. 2026)).
Uniqueness	Evaluates whether the risk is specific to the particular AI use case, rather than broadly applicable across AI systems or domains (1–3 Likert (Kieslich, Helberger, and Diakopoulos 2026)).
Novelty	Rates the degree of creativity and non-obviousness of the risk, ranging from mundane or checklist-like to imaginative and original (1–3 Likert (Dean et al. 2006; Fröhling et al. 2026),).
Usability	Measures how clear, understandable, and practically engaging the risk (1–5 Likert (Lewis and Sauro 2009; Brooke et al. 1996)).
Severity	Assesses the seriousness of potential harm and its downstream social, economic, or institutional consequences (1–5 Likert (AI 2023; Fröhling et al. 2026)).

Table 1: Evaluation rubric used for annotating AI risks. Each dimension corresponds to explicit annotation questions posed to evaluators, and is grounded in prior literature.

who frequently develop assessments for AI. For consistency and ease of comparisons, all the ratings were normalized to five-point Likert scale. These expert annotations provide a quantitative basis for comparing risk quality across brainstorming methods and with system-generated outputs.

4.3 Establishing Single-/Multi-Agent Baselines

We compare our approach against two baselines commonly used in AI risk assessment:

Single-Turn LLM Expert Prompting: A single LLM is prompted to generate and rank risks from an expert perspective, conditioned on a description of the AI use case system and a list of stakeholders. This single LLM approach reflects current practice in automated risk generation tools, such as Farsight (Wang et al. 2024), ExploreGen (Herdal et al. 2024), and AHA! (Bućinca et al. 2023).

Multi-Turn Plurals without Network Awareness: We use standard Plurals framework (Ashkinaze et al. 2025a) to simulate multiple in-silico stakeholder agents interacting in a flat ensemble structure, without explicit network structure (topology) or network-based ranking. This baseline isolates the contribution of network structure and measures, and represents current agentic LLM brainstorming method.

To evaluate those two baselines and our framework, we mapped all system-generated risks to practitioners-generated risks using a two-stage process combining embedding-based similarity and manual validation. Each of the 244 practitioners-curated risks was embedded and matched to the five most similar system outputs, producing over 1,000 candidate pairs. We manually validated semantic matches, allowing one-to-many mappings. In total, 108 human risks (44%) matched at least one system-generated risk. These matched risks were used for our evaluation.




Figure 2: Our analysis utilizes two metrics: “Combined Score” aggregates Plausibility, Likelihood, Uniqueness, Usability and Severity while “Novelty” reflects risk diversity. All measures were normalized to the five-point Likert scale, and the top performing systems are bolded. **A:** Comparison of Network Topologies (using Snowballing Typical Roles and Degree Centrality): Small-World networks outperform both structured hierarchies and unstructured random networks. They significantly outperform the Baselines, demonstrating the value of structured interaction. **B:** Impact of Centrality Measures on Risk Quality (Small-World Network + Snowballing Typical Roles): Degree Centrality yields the highest quality (Combined Score) by capturing consensus; Betweenness Centrality is most effective at surfacing Novel risks by prioritizing risks from bridging stakeholders.

4.4 Results: Network Prioritization Drives Novelty

We evaluate our framework against static and unstructured baselines using two metrics: a *Combined Score* (an aggregate of plausibility, likelihood, uniqueness, usability, and severity), and a *Novelty* score (specifically isolating risk diversity and non-obviousness). All scores are normalized to a five-point Likert scale. We explicitly report *Novelty* score to capture non-obvious risks, often disproportionately affecting minority or marginalized populations, that are systematically obscured by consensus-driven aggregates.

By moving beyond flat lists of stakeholders, we show that network topology, centrality measures, and discovery strategies fundamentally shape the quality of generated risks (evaluation results on all system variants are provided in Appendix with a detailed discussion of the discovery strategies). We highlight our key findings below.

Small-World Topologies Optimize Ideation Efficiency. Network science literature (Watts and Strogatz 1998) posits that “small-world” networks, characterized by high local clustering and short average path lengths, are uniquely efficient at information diffusion. We find that this structural advantage holds for simulated AI risk ideation, significantly outperforming unstructured interactions. Figure 2A compares three network topologies using the same strategy for creating the set of stakeholders:

Small-World network achieves the best balance (Combined 4.05, Novelty 3.18). The high local clusters allows distinct stakeholder groups (e.g., a “Security Cluster” vs. a “Patient Advocacy Cluster”) to develop coherent, deep arguments locally before propagating them.

Random network performs competitively on the Combined Score but suffer on Novelty (3.05). We found stakeholders lack the local context to refine specific critiques, leading to more generic outputs.

Hierarchical network performs poorly on Novelty (2.95). Designed to mimic corporate reporting lines, these structures appear to introduce information filtering. “Middle manager” agents may inadvertently smooth over radical sig-

nals from leaf nodes (end-users) before they reach the central nodes, effectively suppressing diverse risk identification.

The structural impact is evident when comparing Small-World results to the *Single-turn LLM* (Combined 3.29) and *Multi-turn Plurals* (Combined 3.72). The Single-turn baseline lacks the iterative refinement of ideas, while the Multi-turn baseline lacks the community structure to support specialized discourse. The Small-World topology provides a +0.76 improvement in Combined Score over the standard LLM baseline, demonstrating that *how* agents are connected is as critical as the agents themselves.

Novelty is Driven by Bridging Stakeholders (Using Betweenness Centrality). A core challenge in AI auditing is moving beyond “known knowns” (high-consensus risks) to uncover “unknown unknowns”. Standard methods often converge on high-likelihood but generic risks (DeVos et al. 2022; Lancaster et al. 2024). We find that the choice of network centrality measure used to prioritize risks alters the nature of the output, revealing a trade-off between consensus and discovery (Figure 2B) Ranking risks by *Degree Centrality* (popularity) yields the highest Combined Scores (4.05). However, qualitative analysis suggests these risks often represent consensus views, broadly recognized issues like data privacy leaks or standard bias. While these are valid, they are often already known to development teams. In contrast, prioritizing risks via *Betweenness Centrality* results in the highest *Novelty* scores (3.48). Betweenness centrality identifies stakeholders who act as bridges between otherwise disconnected sub-communities. Unlike high-degree nodes that reinforce dominant narratives, bridging nodes are uniquely positioned to translate context across epistemic boundaries.

To illustrate this mechanism, we analyzed the risks generated by the top 10 stakeholders ranked by Betweenness Centrality (Figure 3). Some of those top 10 stakeholders also overlap with those in top 10 of Degree Centrality measures such as IT Manager and Product Manager. The data reveals a stark divergence in Novelty depending on the *type* of domain the stakeholder bridges. Stakeholders who also has high Degree Centrality tended to identify valid but low-

A TOP 10 STAKEHOLDERS BY BETWEENNESS CENTRALITY MEASURE

STAKE-HOLDER	RISK STATEMENT	NOVELTY SCORE
Parent	Chatbot overreliance could hinder relationship development	5
Mental Health Researcher	Failure to replicate evidence-based therapeutic techniques could worsen mental health	4
Policy Maker	Failure to identify crisis situations or escalate them could endanger user safety during emergencies	4
Clinical Pharmacist	Misinterpreting medication instructions provided by chatbot could lead to non-adherence	4
Nurse	Providing inaccurate health responses could cause confusion and adverse patient outcomes	3
IT Manager	Poor 24/7 scalability could cause crashes and disrupt support	2
End User	Failure to recognize emotional cues may make interactions feel impersonal during distress	4
Product Manager	Infrastructure incompatibility could delay deployment and reduce operational effectiveness	1
User Experience Designer	Failure to incorporate clinical insights may result in a product that ignores emotional needs	4
Regulatory Body	Failure to adhere to transparency and bias standards triggers regulatory scrutiny	3

B STAKEHOLDER NETWORK




Figure 3: **A:** Risks and their novelty scores generated by the top 10 stakeholders ranked by Betweenness Centrality. **B:** The stakeholder network constructed through snowballing, within which in-silico LLM agents are connected and communicate to brainstorm risks. Many stakeholders who bridge diverse domains with top Betweenness Centrality such as Parent, Mental Health Researchers, Policy Makers and Clinical Pharmacist are the primary identifiers of highly novel risks (scores 4–5) such as psychological dependency and user safety during crisis.

novelty risks. For example, the *Product Manager* identified the risk of “*Incompatibility with existing infrastructure*” (Novelty=1), and the *IT Manager* identified scalability risks regarding “*24/7 volume causing crashes*” (Novelty=2). Their outputs reflect standard operational concerns common in software deployment. In contrast, stakeholders who are mainly bridging sub-communities with extremely high Betweenness Centrality but relatively low Degree Centrality provided the highest novelty risks:

The *Parent* (Novelty 5), acting as a bridge between the domestic sphere and the educational/technical system, identified the risk of “*Developing an unhealthy over-reliance on the chatbot*”, hindering real-life social development. This is a profound socio-psychological harm that purely technical agents missed, although this harm is documented in prior literature (e.g., loneliness-through-technology (Wilson 2018)). The *Mental Health Researcher* and *Clinical Pharmacist* (both Novelty 4) bridged clinical standards and AI interaction, identifying risks such as “*Inability to effectively replicate evidence-based therapeutic techniques*” and “*Misinterpreting medication instructions*”, respectively.

The *Policy Maker* (Novelty 4), bridging public safety and user interaction, highlighted the systemic failure to “*identify*

crisis situations or escalate them”, a critical safety gap.

This analysis confirms that novelty is not randomly distributed. It is structurally concentrated in stakeholders who navigate the boundaries between distinct communities. By prioritizing these bridging nodes via Betweenness Centrality, the system identifies high-impact, non-obvious risks that are invisible to the commonly identified risks.

5 Evaluation of Our Framework Supporting Participatory AI Risk Assessment

In this section, we focus on testing whether our framework actually helps a human-led ideation session. The goal of this evaluation was to test whether seeding a participatory session with network-generated risks leads participants to identify more risks, particularly systemic and downstream risks, than seeding the session with practitioner-generated risks.

5.1 Participatory AI Risk Assessment with Futures Wheel

We conducted a user study using the Futures Wheel method (Glenn 2009). It is a structured brainstorming technique for identifying first-, second-, and third-order consequences

of a central topic, enabling participants to explore cascading and interconnected effects. Starting from an AI chatbot companion, participants iteratively answer the question “If this occurs, what happens next?”, generating primary consequences in the first layer, followed by secondary and tertiary consequences in the second and third layers. This process encourages participants to move beyond immediate risks toward more systemic and long-term impacts.

We selected the Futures Wheel as our method of participatory AI risk assessment for three reasons. First, it offers a simple, structured format for team ideation that enables participants to generate a large number of risks without requiring domain expertise (Glenn 2009). Second, it pushes participants to think beyond obvious risks for AI chatbots by exploring their second- and third-order consequences (Bengston 2016). Third, it has been successfully used in large-scale participatory AI risk assessment studies (with over 250 participants across six Global Majority and Minority countries (Hohendanner et al. 2025)), making it appropriate for our intended participants.

Participant Demographics and Recruitment. We are particularly interested in testing how our framework supports risk ideation among a crucial but under-represented demographic in AI risk literature: young, frequent chatbot users from predominantly non-Western (non-US, non-European) backgrounds (Png 2022; Birhane et al. 2022). Young users in the Global South and non-Western contexts are major consumers of AI companions (Scherr et al. 2025), yet their perspectives are rarely centered in participatory risk assessments (Png 2022; Birhane et al. 2022). We recruited 11 teams of participants, comprising a total of 40 individuals (with three to four people per team). Participants generally reported high AI chatbot usage: with 20 participants (50%) using AI chatbots multiple times per day, 6 (15%) once daily, 11 (27.5%) 2–6 times per week, and 3 (7.5%) once per week. Additionally, 13 participants (32.5%) reported having experienced technology-related harms, citing issues such as misinformation, privacy loss, social media addiction, surveillance, and over-reliance on technology. Ethnically, the cohort was diverse: 15 participants (37.5%) identified as Middle Eastern or North African, 10 (25%) as Asian, 8 (20%) as White, 4 (10%) as Black, 1 (2.5%) as Mixed, 1 (2.5%) as Central Asian, and 1 (2.5%) preferred not to say.

Experimental Conditions: Practitioners-Seeded (Control) vs. Framework-Seeded Ideation (Treatment). The teams were randomly assigned to two conditions: *Control* teams: five teams were provided with a report containing a set of risks brainstormed by AI practitioners in our first evaluation, i.e., blank-slate risk brainstorming workshops as described in Section 4.1; *Treatment* teams: six teams were instead provided with a report containing the same number of risks but generated by our framework with a small-world network topology (Section 3). All the risks were randomly sampled from the set of top prioritized risks from each approach to ensure coverage of the wide spectrum. All teams then conducted the participatory AI risk assessment ideation exercise, brainstorming risks across three layers of consequences in the Futures Wheel.




Figure 4: **Futures Wheel participatory ideation results.** Teams seeded with our networked simulation generated risks (Treatment) identified more second- and third-order risks (A) than teams seeded with practitioner-generated risks (Control). They also surfaced a broader range of socio-technical harms, particularly human–computer interaction and socioeconomic and environmental risks (B).

5.2 Results: Our Framework Identifies More and Deeper Systemic Risks

Overall, the treatment teams generated a substantially higher volume of risks compared to the control teams. On average, treatment teams brainstormed 21.1 risks per team, compared with 15.6 risks among control teams. Crucially, as shown in Figure 4A, this increase in volume was entirely driven by a greater capacity to uncover higher-order, systemic consequences, which tend to be harder to envision (Friedman and Hendry 2012). Both conditions yielded an identical average of 5.2 risks at Level 1 (immediate consequences). However, the treatment group significantly outperformed the control group in subsequent layers, identifying an average of 8.7 risks at Level 2 (compared to 5.2 in the control) and 7.2 risks at Level 3 (compared to 5.2 in the control). This demonstrates that our framework effectively prompts human participants to think deeper and anticipate cascading effects rather than exhausting their effort on obvious primary outcomes.

To understand which risks were identified, we classified the brainstormed risks using the MIT AI Risk taxonomy (Slattery et al. 2026) (Figure 4B). We found that the treatment teams generated more risks across almost all categories. Most notably, they surfaced substantially more risks in complex, downstream areas such as Human-Computer Interaction (8.2 vs. 6.0), Socioeconomic & Environmental harms (5.6 vs. 2.4), and AI System Safety Failures & Limitations (3.5 vs. 0.0). Interestingly, the control teams only produced more risks in the “Privacy & Security” category (1.5 vs. 1.0). This demonstrates practitioners-seeded baselines often anchor human ideation on well-known, conventional risks, whereas the networked stakeholder simulation encourages the exploration of non-obvious, systemic harms.

Ultimately, these results indicate that utilizing our framework to support the ideation stage of participatory AI risk

assessments is significantly more effective than relying on a baseline of practitioner-generated risks. It not only increases the total number of identified risks, but specifically enhances human capacity to discover the higher-order socio-technical harms that are typically the hardest to spot.

6 Conclusion and Discussion

We presented a framework that transforms AI risk assessment into a dynamic, multi-perspective exploration. By generating stakeholder networks and utilizing social network analysis measures like Betweenness Centrality, we demonstrated a capability to uncover high-impact risks that traditional methods miss. Our findings suggest that AI risk assessment benefits not only from who is involved, but from how stakeholder perspectives, even in-silico, are structured, connected, and weighted. By treating risk identification as a network-mediated ideation, we provide practitioners with a tool to help with a human-led ideation session in the participatory AI risk assessments.

Implications. This work has three key implications:

From flat independent assessments to network-mediated ideation. Current AI risk assessment practices, whether human-led brainstorming workshops or LLM-generated risk assessments, implicitly assume that risks are enumerable, independent, and equally visible to all evaluators. Risks emerge at the intersection of stakeholder roles (e.g., between legal compliance and user vulnerability, or between product incentives and mental health outcomes) are systematically underrepresented by flat generation methods. By modeling stakeholder interaction explicitly, our approach shifts the focus from coverage of known risk categories to discovery of emergent, relational harms.

Scalable pluralism and concrete usage pathways. Our framework is explicitly designed as a complementary pre-assessment tool, not a substitute for authentic participation in participatory risk assessment. Risk management professionals can integrate this tool into early design sprints to conduct structured gap analyses against regulatory taxonomies (e.g., the EU AI Act). By running the simulation prior to a human workshop, organizers can use the pre-generated materials, specifically areas of simulated disagreement and heavily populated systemic risks, as prompts for human ideation as we have done in our second evaluation. This prevents human participants from starting with a “blank slate” and focuses their valuable time contextualizing, and deepening complex risks. The tool allows organizers to explicitly balance AI practitioners’ operational insights with the broader, theoretically informed views of researchers. If the simulation reveals that a specific bridging stakeholder (e.g., a “Mental Health Advocate”) is central to identifying high-severity risks, organizers are provided an evidence-based justification to invest resources into recruiting humans holding those lived experiences for subsequent participatory stages. We explicitly caution against using our tool to bypass authentic inclusion with actual human participation for the full risk assessment. Most importantly, outputs from this tool must not be submitted to auditors or regulators as a substitute for authentic stakeholder evidence or testimony.

Prioritization as a normative choice. Our comparison of centrality network based ranking strategies highlights that risk prioritization is not value-neutral. Degree-based ranking favors widely recognized, consensus risks; Betweenness-based ranking elevates boundary-spanning concerns. This trade-off mirrors real-world governance tensions between addressing “known knowns” and anticipating low-frequency but high-impact harms (potentially to marginalized stakeholders). Making these prioritization logics explicit allows practitioners and regulators to align risk assessment outputs with their normative goals.

Limitations. Our work has five main limitations.

Firstly, our framework relies on LLM-simulated stakeholders to model structured information flow among network-conditioned agents, rather than authentic stakeholder deliberation (Habermas 1981). Because these in-silico representations cannot fully capture the lived experience, contextual knowledge, or normative authority of real affected communities, certain harms, particularly those rooted in cultural, emotional, or marginalized perspectives, may be misrepresented or missed.

Secondly, while our Futures Wheel study grounded the framework in a human-seeded task, it involved only one stakeholder type (young, non-Western users). Crucial domain experts and vulnerable groups, such as mental health clinicians and social workers, were not included. Systematically benchmarking against workshops involving both highly vulnerable stakeholders and these specialized domain experts remains a critical direction for future work.

Thirdly, the quality and diversity of generated risks depend on the underlying language model and the architectural choices made by the system designers. This introduces a structural paradox: while the tool is intended to mitigate the narrow biases of traditional human brainstorming, we (as researchers) defined the network topologies and prompt structures, which might inevitably encode our own biases into the system. Compounding this issue, because LLMs often encode representational biases from historically skewed training data (Bai et al. 2025; Gupta et al. 2023; Shujaa, Dorleón, and Tang 2025), simulated stakeholders run the risk of producing algorithmic caricatures rather than authentic reflections of marginalized communities (Vecchione, Levy, and Barocas 2021). Procedurally, it is imperative that our framework is strictly utilized as a pre-assessment tool to seed human ideation, not as a substitute for authentic human participation. A vital direction for future work is to systematically benchmark these stereotyping effects and validate our simulations directly against workshops involving highly vulnerable, directly affected stakeholders to ensure the tool safely bridges, rather than widens, representational gaps.

Fourthly, the constructed stakeholder networks encode simplifying assumptions about influence flow (e.g., acyclicity), which may not fully reflect the reciprocal and evolving nature of real-world stakeholder dynamics.

Finally, our evaluation focuses on a single use case (AI chatbot companions), and further studies are needed to assess the generalizability of the approach across domains with different stakeholder structures and risk profiles.

Ethics Statement

All human participants, including the AI practitioners in the baseline study and the annotators recruited via Prolific, provided informed consent prior to participation. No personally identifiable information was collected, and all data was anonymized prior to analysis. Participants in the baseline and Futures Wheel evaluation phases were compensated exceeding the local minimum wage. For the expert verification tasks on Prolific, stringent screening criteria were applied to ensure domain expertise.

References

- AI, N. 2023. Artificial Intelligence Risk Management framework (AI RMF 1.0). URL: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai>, 100–1.
- AIDahoul, N.; Rahwan, T.; and Zaki, Y. 2025. AI-Generated Faces Influence Gender Stereotypes and Racial Homogenization. *Scientific Reports*, 15(1).
- Ashkinaze, J.; Fry, E.; Edara, N.; Gilbert, E.; and Budak, C. 2025a. Plurals: A System for Guiding LLMs via Simulated Social Ensembles. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 1–21.
- Ashkinaze, J.; Mendelsohn, J.; Qiwei, L.; Budak, C.; and Gilbert, E. 2025b. How AI Ideas Affect the Creativity, Diversity, and Evolution of Human Ideas: Evidence From a Large, Dynamic Experiment. In *Proceedings of the ACM Collective Intelligence Conference (CI)*, 198–213.
- Bai, X.; Wang, A.; Sucholutsky, I.; and Griffiths, T. L. 2025. Explicitly Unbiased Large Language Models Still Form Biased Associations. *Proceedings of the National Academy of Sciences*, 122(8): e2416228122.
- Ballard, S.; Chappell, K. M.; and Kennedy, K. 2019. Judgment Call the Game: Using Value Sensitive Design and Design Fiction to Surface Ethical Concerns Related to Technology. In *Proceedings of the ACM Conference in Designing Interactive Systems Conference (DIS)*, 421–433. ACM. ISBN 9781450358507.
- Bengston, D. N. 2016. The Futures Wheel: A Method for Exploring the Implications of Social–Ecological Change. *Society & natural resources*, 29(3): 374–379.
- Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022. Power to the People? Opportunities and Challenges for Participatory AI. In *Proceedings of the ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 1–8.
- Borgatti, S. P.; Jones, C.; and Everett, M. G. 1998. Network Measures of Social Capital. *Connections*, 21(2): 27–36.
- Bouschery, S. G.; Blazevec, V.; and Piller, F. T. 2024. Artificial Intelligence-Augmented Brainstorming: How Humans and AI Beat Humans Alone. SSRN Working Paper, posted April 3, 2024.
- Brooke, J.; et al. 1996. SUS: A Quick and Dirty Usability Scale. *Usability Evaluation in Industry*, 189(194): 4–7.
- Buçinca, Z.; Pham, C. M.; Jakesch, M.; Ribeiro, M. T.; Olteanu, A.; and Amershi, S. 2023. AHA!: Facilitating AI Impact Assessment by Generating Examples of Harms. arXiv:2306.03280.
- Choi, D.; Hong, S.; Park, J.; Chung, J. J. Y.; and Kim, J. 2024. Creativeconnect: Supporting Reference Recombination for Graphic Design Ideation With Generative AI. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 1–25.
- Corbett, E.; Denton, R.; and Erete, S. 2023. Power and Public Participation in AI. In *Proceedings of the ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 1–13.
- Dean, D. L.; Hender, J. M.; Rodgers, T. L.; and Santanen, E. L. 2006. Identifying Quality, Novel, and Creative Ideas: Constructs and Scales for Idea Evaluation. *Journal of the Association for Information Systems*, 7(10): 646–699.
- Delgado, F.; Yang, S.; Madaio, M.; and Yang, Q. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 1–23.
- Deng, W. H.; Claire, W.; Han, H. Z.; Hong, J. I.; Holstein, K.; and Eslami, M. 2025. WeAudit: Scaffolding User Auditors and AI Practitioners in Auditing Generative AI. *Proc. ACM Hum.-Comput. Interact.*, 9(7).
- DeVos, A.; Dhabalia, A.; Shen, H.; Holstein, K.; and Eslami, M. 2022. Toward User-Driven Algorithm Auditing: Investigating Users’ Strategies for Uncovering Harmful Algorithmic Behavior. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 1–19.
- Diakopoulos, N.; and Johnson, D. 2021. Anticipating and Addressing the Ethical Implications of Deepfakes in the Context of Elections. *New Media & Society*, 23(7): 2072–2098.
- Diehl, M.; and Stroebe, W. 1987. Productivity Loss in Brainstorming Groups: Toward the Solution of a Riddle. *Journal of Personality and Social Psychology*, 53(3): 497–509.
- European Commission. 2026. Consolidated text: Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance).
- Friedman, B.; and Hendry, D. 2012. The Envisioning Cards: A Toolkit for Catalyzing Humanistic and Technical Imaginations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 1145–1148.
- Fröhlings, L.; Giaconia, A.; Bogucka, E. P.; and Quercia, D. 2026. Agent-Supported Foresight for AI Systemic Risks: AI Agents for Breadth, Experts for Judgment. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM. ISBN 9798400722783.
- Fukumura, K.; and Ito, T. 2025. Can LLM-Powered Multi-Agent Systems Augment Human Creativity? Evidence from Brainstorming Tasks. In *Proceedings of the ACM Collective Intelligence Conference (CI)*, 20–29. ACM. ISBN 9798400714894.

- Ganguli, D.; Lovitt, L.; Kernion, J.; Askell, A.; Bai, Y.; Kadavath, S.; Mann, B.; Perez, E.; Schiefer, N.; Ndousse, K.; et al. 2022. Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned. *arXiv preprint arXiv:2209.07858*.
- Ge, S.; Zhou, C.; Hou, R.; Khabsa, M.; Wang, Y.-C.; Wang, Q.; Han, J.; and Mao, Y. 2024. MART: Improving LLM Safety with Multi-round Automatic Red-Teaming. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (ACL)*, 1927–1937.
- Glenn, J. C. 2009. The Futures Wheel. In Glenn, J. C.; and Gordon, T. J., eds., *Futures Research Methodology, Version 3.0*, 245–263. The Millennium Project.
- Govers, J.; Šćepanović, S.; and Quercia, D. 2026. When and How AI Should Assist Brainstorming for AI Impact Assessment. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency*, 2745–2780.
- Gupta, S.; Shrivastava, V.; Deshpande, A.; Kalyan, A.; Clark, P.; Sabharwal, A.; and Khot, T. 2023. Bias Runs Deep: Implicit Reasoning Biases in Persona-Assigned LLMs. *arXiv preprint arXiv:2311.04892*.
- Habermas, J. 1981. *The Theory of Communicative Action*. Beacon press.
- Herdel, V.; Šćepanović, S.; Bogucka, E.; and Quercia, D. 2024. ExploreGen: Large Language Models for Envisioning the Uses and Risks of AI Technologies. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, volume 7, 584–596.
- Hohendanner, M.; Ullstein, C.; Onyekwelu, B. A.; Katirai, A.; Kuribayashi, J.; Babalola, O.; Ema, A.; and Grossklags, J. 2025. Initiating the Global AI Dialogues: Laypeople Perspectives on the Future Role of genAI in Society from Nigeria, Germany and Japan. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM. ISBN 9798400713941.
- Hohma, E.; Boch, A.; Trauth, R.; and Lütge, C. 2023. Investigating Accountability for Artificial Intelligence Through Risk Governance: A Workshop-Based Exploratory Study. *Frontiers in Psychology*, 14: 1073686.
- Kallina, E.; Bohné, T.; and Singh, J. 2025. Stakeholder Participation for Responsible AI Development: Disconnects Between Guidance and Current Practice. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 1060–1079.
- Kieslich, K.; Helberger, N.; and Diakopoulos, N. 2026. Scenario-based Sociotechnical Rnvisioning (SSE): An Approach to Enhance Systemic Risk Assessments. *AI and Ethics*, 6(3).
- Lancaster, C. M.; Schulenberg, K.; Flathmann, C.; McNeese, N. J.; and Freeman, G. 2024. “It’s Everybody’s Role to Speak Up... But Not Everyone Will”: Understanding AI Professionals’ Perceptions of Accountability for AI Bias Mitigation. *ACM Journal on Responsible Computing*, 1(1): 1–30.
- Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; et al. 2019. WeBuildAI: Participatory Framework for Algorithmic Governance. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–35.
- Lewis, J. R.; and Sauro, J. 2009. The Factor Structure of the System Usability Scale. In Kurosu, M., ed., *Human Centered Design*, 94–103. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-02806-9.
- Meincke, L.; Nave, G.; and Terwiesch, C. 2025. ChatGPT Decreases Idea Diversity in Brainstorming. *Nature Human Behaviour*, 9(6): 1107–1109.
- Park, J. S.; O’Brien, J.; Cai, C. J.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*, 1–22.
- Parmar, B. L.; Freeman, R. E.; Harrison, J. S.; Wicks, A. C.; Purnell, L.; and De Colle, S. 2010. Stakeholder Theory: The State of the Art. *Academy of Management Annals*, 4(1): 403–445.
- Png, M.-T. 2022. At the Tensions of South and North: Critical Roles of Global South Stakeholders in AI Governance. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 1434–1445.
- Rowley, T. J. 1997. Moving Beyond Dyadic Ties: A Network Theory of Stakeholder Influences. *Academy of Management Review*, 22(4): 887–910.
- San Vito, P. D. C.; Stumpf, S.; Hyde-Vaamonde, C.; and Thuermer, G. 2025. Ensuring Artificial Intelligence is Safe and Trustworthy: The Need for Participatory Auditing. In *ACM Conference on Human Factors in Computing Systems (CHI): Sociotechnical AI Governance: Opportunities and Challenges for HCI*.
- Scherr, S.; Cao, B.; Jiang, L. C.; and Kobayashi, T. 2025. Explaining the Use of AI Chatbots as Context Alignment: Motivations Behind the Use of AI Chatbots Across Contexts and Culture. *Computers in Human Behavior*, 172: 108738.
- Schmitz, A.; Mock, M.; Gorge, R.; Cremers, A. B.; and Poretschkin, M. 2025. A Global Scale Comparison of Risk Aggregation in AI Assessment Frameworks. *AI and Ethics*, 5(2): 1407–1432.
- Serrano, M. Á.; Boguná, M.; and Vespignani, A. 2009. Extracting the Multiscale Backbone of Complex Weighted Networks. *Proceedings of the National Academy of Sciences*, 106(16): 6483–6488.
- Shujaa, S.; Dorleon, G.; and Tang, A. 2025. How LLMs Handle Cultural Bias: Reactions to Asian Minority Historical Narratives. In *International Conference on Asian Digital Libraries*, 39–52. Springer.
- Slattery, P.; Saeri, A. K.; Grundy, E. A.; Graham, J.; Noetel, M.; Uuk, R.; Dao, J.; Pour, S.; Casper, S.; and Thompson, N. 2026. The AI Risk Repository: A Meta-Review, Database, and Taxonomy of Risks from Artificial Intelligence. *Patterns*, 7(5): 101517.

- Steimers, A.; and Schneider, M. 2022. Sources of Risk of AI Systems. *International Journal of Environmental Research and Public Health*, 19(6): 3641.
- Uuk, R.; Gutierrez, C. I.; Guppy, D.; Lauwaert, L.; Kasirzadeh, A.; Velasco, L.; Slattery, P.; and Prunkl, C. 2024. A Taxonomy of Systemic Risks From General-Purpose AI. *arXiv:2412.07780*.
- Vecchione, B.; Levy, K.; and Barocas, S. 2021. Algorithmic Auditing and Social Justice: Lessons From the History of Audit Studies. In *Proceedings of the ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 1–9.
- Wadinambiarachchi, S.; Kelly, R. M.; Pareek, S.; Zhou, Q.; and Velloso, E. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM. ISBN 9798400703300.
- Wang, W.-F.; Lu, C.-T.; Ponsa i Campanyà, N.; Chen, B.-Y.; and Chen, M. Y. 2025. Aldeation: Designing a Human-AI Collaborative Ideation System for Concept Designers. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 1–28.
- Wang, Z. J.; Kulkarni, C.; Wilcox, L.; Terry, M.; and Madaio, M. 2024. Farsight: Fostering Responsible AI Awareness During AI Application Prototyping. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 1–40.
- Watts, D. J.; and Strogatz, S. H. 1998. Collective Dynamics of ‘Small-World’ Networks. *Nature*, 393(6684): 440–442.
- Weidinger, L.; Uesato, J.; Rauh, M.; Griffin, C.; Huang, P.-S.; Mellor, J.; Glaese, A.; Cheng, M.; Balle, B.; Kasirzadeh, A.; et al. 2022. Taxonomy of Risks Posed by Language Models. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 214–229.
- Wilson, C. 2018. Is It Love or Loneliness? Exploring the Impact of Everyday Digital Technology Use on the Wellbeing of Older Adults. *Ageing & Society*, 38(7): 1307–1331.
- Xia, B.; Lu, Q.; Perera, H.; Zhu, L.; Xing, Z.; Liu, Y.; and Whittle, J. 2023. Towards Concrete and Connected AI Risk Assessment (C2AIRA): A Systematic Mapping Study. In *IEEE/ACM International Conference on AI Engineering–Software Engineering for AI (CAIN)*, 104–116. IEEE Computer Society.
- Yu-Han, C.; and Chun-Ching, C. 2023. Investigating the Impact of Generative Artificial Intelligence on Brainstorming: A Preliminary Study. In *International Conference on Consumer Electronics - Taiwan (ICCE-Taiwan)*, 193–194.
- Zhang, R.; Li, H.; Meng, H.; Zhan, J.; Gan, H.; and Lee, Y.-C. 2025. The Dark Side of AI Companionship: A Taxonomy of Harmful Algorithmic Behaviors in Human-AI Relationships. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 1–17.