



Delft University of Technology

Document Version

Final published version

Licence

CC BY-NC-ND

Citation (APA)

Kusters, A., Prajod, P., Cesar, P., & El Ali, A. (2026). More Human or More AI? Visualizing Human-AI Collaboration Disclosures in Journalistic News Production. In N. Oliver, D. A. Shamma, H. Candello, P. Cesar, P. Lopes, A. Bozzon, T. Kosch, V. Liao, X. Ma, & More Editors (Eds.), *CHI 2026 - Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* Article 729 Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3772318.3791288>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

More Human or More AI? Visualizing Human-AI Collaboration Disclosures in Journalistic News Production

Amber Kusters
 Centrum Wiskunde & Informatica
 Amsterdam, The Netherlands
 Utrecht University
 Utrecht, The Netherlands
 amber.kusters@cwi.nl

Pooja Prajod
 Centrum Wiskunde & Informatica
 Amsterdam, The Netherlands
 pooja@cwi.nl

Pablo Cesar
 Centrum Wiskunde & Informatica
 Amsterdam, The Netherlands
 Delft University of Technology
 Delft, The Netherlands
 p.s.cesar@cwi.nl

Abdallah El Ali
 Centrum Wiskunde & Informatica
 Amsterdam, The Netherlands
 Utrecht University
 Utrecht, The Netherlands
 aea@cwi.nl

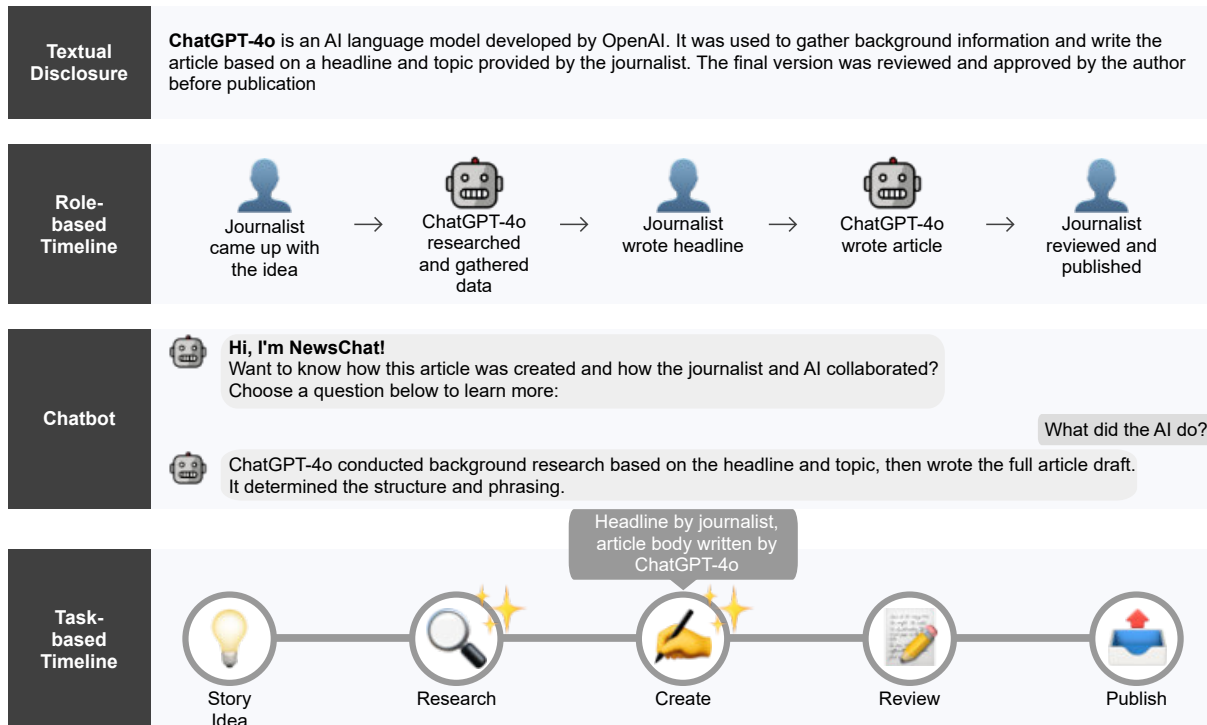


Figure 1: Our human-AI collaboration disclosure visualizations, with an example for a primarily AI-written article. From top to bottom: (1) Textual Disclosure, (2) Role-based Timeline, (3) Chatbot interface, and (4) Task-based Timeline.

Abstract

Within journalistic editorial processes, disclosing AI usage is currently limited to simplistic labels, which misses the nuance of how humans and AI collaborated on a news article. Through co-design sessions (N=10), we elicited 69 disclosure designs and implemented four prototypes that visually disclose human-AI collaboration in journalism. We then ran a within-subjects lab study (N=32) to examine how disclosure visualizations (Textual, Role-based Timeline, Task-based Timeline, Chatbot) and collaboration ratios (Primarily



Human vs. Primarily AI) influenced visualization perceptions, gaze patterns, and post-experience responses. We found that textual disclosures were least effective in communicating human-AI collaboration, whereas Chatbot offered the most in-depth information. Furthermore, while role-based timelines amplified AI contribution in primarily human articles, task-based timeline shifted perceptions toward human involvement in primarily AI articles. We contribute Human-AI collaboration disclosure visualizations and their evaluation, and cautionary considerations on how visualizations can alter perceptions of AI's actual role during news article creation.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**.

Keywords

AI disclosure, human-AI collaboration, visualization, journalism, news production

ACM Reference Format:

Amber Kusters, Pooja Prajod, Pablo Cesar, and Abdallah El Ali. 2026. More Human or More AI? Visualizing Human-AI Collaboration Disclosures in Journalistic News Production. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3772318.3791288>

1 Introduction

Artificial intelligence (AI) technologies such as large language models (LLMs) are increasingly being used for a variety of tasks in the newsroom [76]. Generative AI (GenAI) tools like ChatGPT, Gemini, and Midjourney can produce high-quality multimodal outputs based on user inputs [32, 66]. These developments have the potential to affect both users and the media ecosystem, essentially blurring the lines between fiction and reality as users consume news media content. Already in 2024, model output quality has advanced to a point where human observers can no longer reliably distinguish between generative AI (GenAI) output and that produced by humans [43]. This further intensified with recent work on Turing tests demonstrating near-chance human detection rates of state-of-the-art LLMs, where some instances were judged to be more human than real human interlocutors [60]. These tools are often used in collaborative workflows where GenAI assists journalists rather than replacing them entirely. Prior work has shown that journalists use GenAI for tasks such as news generation, validation, distribution, and moderation [23], and primarily for content production, particularly when it comes to text-based outputs [27]. This is also in line with the global public perception of AI usage in journalism, which includes text editing, translation, and data analysis, but also with some believing that the entire article has been written by AI [34].

As GenAI use in journalism grows, ethical concerns regarding transparency are rising, centering around trust and misinformation [35, 82], as well as privacy and accountability [95]. This is further pronounced considering that current algorithmic techniques for detecting text (co-)produced by GenAI systems (e.g., ChatGPT) remain unreliable [85], similarly with cryptographic approaches that aim to encode provenance in digital media, for example in images (e.g.,

Coalition for Content Provenance and Authenticity [16]) or audio (e.g., SynthID [26]). Even if such detection techniques were flawless, how provenance signals should be disclosed or communicated to end users remains a key challenge [28]. Recent studies found that most readers (sampled from six countries) expressed a desire for disclosure when exposed to journalistic tasks produced using 'AI with some human oversight' [34]. Moreover, regulatory frameworks such as the AI Disclosure Act of 2023 in the United States [104] and Article 50 of the EU AI Act [33] articulate requirements on disclosing AI usage, which is intended to mitigate these risks. However, the regulation remains vague with terms such as "authentic" and "context of use" which leaves room for interpretation, creating challenges in design and implementation [28]. Additionally, regulation may be too limited, exempting disclosure after human review even though audiences still expect to be informed of AI usage [81], in addition to expecting standardized measures and sourced references [74, 111]. While different types of disclosures have been researched, from certification labels [88], to detailed AI Usage Cards [112], or showing more details progressively [97], it remains a core challenge of how to not only disclose that AI was used, but also to capture the nuance of how humans and AI collaborated. Human-AI collaboration refers to a dynamic and adaptive partnership in which humans and GenAI systems work together towards a shared goal [37, 51, 86]. Representing and disclosing how much of the content was created by humans versus GenAI is important as GenAI usage can alter journalists' sense of authorship [55], and varies across writing stages and tasks. Yet, prior work has not explored how to visually represent these collaboration ratios across editorial stages within the journalistic workflow.

In this paper, we adopt a mixed-methods approach (Fig. 2) to design and evaluate disclosure visualizations of human-AI collaboration in journalistic news production. In part one, we ask: **(RQ1)** How can we visually represent the ratios of human-AI collaboration output in news production? To answer this, we conducted co-design sessions with designers and HCI experts (N=10) to generate disclosure design ideas. This generated 69 design concepts, exploring a wide design space of static and interactive representations. Based on usability and information visualization principles, a diverse selection was made to turn ideas into four prototypes: (a) Textual Disclosure, (b) Role-based Timeline, (c) Menu-based Chatbot, and (d) Task-based Timeline. This led to Part 2, where we ask: **(RQ2)** How do different disclosure visualizations communicate human-AI collaboration to readers? To answer this, we ran a controlled, within-subjects lab study (N=32) combining questionnaires, eye tracking, and interviews. We examined how visualization type and human-AI collaboration dynamics influenced perceived clarity, perceived information, perceptions of AI and journalist contributions, and visual attention patterns when engaging with news articles and disclosure visualizations.

Our key findings show that (a) all prototypes effectively communicated human-AI collaboration output ratios, but for varied purposes; (b) textual disclosures were least effective in communicating human-AI collaboration; (c) Role-based and Task-based Timelines provided clearer overviews of editorial steps, while the Chatbot offered the most in-depth information; and (d) role-based timelines amplified AI contribution in primarily human-written




Figure 2: The two part study approach with contributions outlined in (bold) blue.

articles, whereas task-based timeline shifted perceptions toward human involvement in primarily AI articles.

Our work offers two primary contributions: **(1)** We introduce and compare four disclosure visualization prototypes¹ for representing human-AI collaboration in news production. **(2)** We provide empirically-backed insights into how these prototypes affect reader perceptions, attention, and understanding, offering cautionary considerations on how visualization design can alter perceptions of AI’s actual role during news article creation.

2 Related Work

2.1 Characterizing human-AI collaboration in journalism

To conceptualize the relationship between humans and AI systems, prior research describes different concepts of human-AI interaction. AI assistance refers to AI providing support, emphasizing human control and responsibility [117]. AI automation reflects varying levels of system independence in generating content with more or less editorial oversight by journalists [100]. Other concepts focus on performance, human-AI augmentation describes AI system outperforming a human alone, whereas human-AI synergy implies the AI system outperforms either human or AI alone [108]. Moreover, human-AI co-creation is a dynamic partnership which leverages strengths towards a shared goal, emphasizing iterative, mutual learning and integration of insights [61], it has also been described as a strategy to amplify creativity and productivity by inspiring novel ideas [59]. While human-AI cooperation frames AI as an independent entity that understands human preferences, working toward shared goals alongside humans [6, 67], with a focus on aligning mental models, and communication [59]. In the work of Breckner et al. [13] cooperation is a part of collaboration, where the result does not exceed the combined individual inputs. In contrast, human-AI collaboration captures a broader, symbiotic, and adaptive relationship in which humans contribute creativity,

intuition, and contextual knowledge, while AI provides computational power, data-driven insights, and efficiency, working together towards common goals [86]. Fragiadakis et al. [37] extend this with a framework of tasks, goals, interaction, and task allocation. Tasks shape the collaboration type, goals can be individual or collective, communication and feedback enhance mutual understanding [113], and allocation shifts dynamically to utilize complementary strengths. Holter and El-Assady [51] similarly emphasize agency, interaction, and adaptation, stressing the need for flexible, balanced partnerships that foster trust, efficiency, and innovation [86]. In journalism, Formosa et al. [36] highlight how humans and AI jointly create content, focusing on continuity (how much AI-generated content is used) and responsibility (how contributions are credited). These considerations intersect with ethical concerns around trust, transparency, and authorship [11], particularly as increased AI involvement reduces the perception of humans as sole authors [36]. In this study, we adopt a definition of human-AI collaboration as a dynamic and adaptive partnership in which both contribute their unique strengths toward shared goals. This definition emphasizes iterative feedback, mutual learning, trust, and ethical transparency, making it especially suited to capture the multifaceted interplay of human and AI contributions in news production.

Human-AI collaboration in journalism occurs across multiple stages in the news production workflow, from story ideation to publishing [14, 91]. At each stage, the role of GenAI varies, affecting the ratio of human versus AI contribution. For example, while creating a headline may be a limited AI contribution, writing the first draft often shift a larger share of authorship to AI. In the story ideation phase, GenAI can support idea generation or content recommendation [40]. In research and information gathering, it can aggregate and analyze data or fact check [40, 91]. During content creation, GenAI is used for writing drafts, generating summaries and headlines, translating text or producing images or audio [40, 91]. In editing and reviewing, it can do style correction to match the news organization style and verify the content, while in publishing it can personalize feeds and customize content [91]. These tasks illustrate how collaboration ratios shift across the stages, causing the amount of contribution by GenAI in the final output to differ.

2.2 AI disclosure representation and communication

AI disclosure can be represented in many forms, ranging from annotations and metadata to visible labels [16, 38]. In journalism, bylines are a traditional form of metadata used to attribute authorship, and could similarly be extended to GenAI systems [58, 90]. Labeling is the most common strategy to reduce risks of GenAI, covering content with visible warnings to alert users [115], such as watermarks [71], filter tags, or platform warnings [90]. Labels can reduce misinformation belief and increase trust in newsroom processes [15, 120], but their effectiveness depends on whether they pursue process-based goals (explaining how content was created) or impact-based goals (warning about potential consequences) [115]. This research builds on process-based approaches, which communicate the human-AI collaboration ratio, while keeping a neutral stance towards the consequences of GenAI usage in news media. Showing human involvement alongside AI contributions can improve transparency

¹GitHub: <https://github.com/cwi-dis/CHI2026-AIDisclosureVisualizations>

and prevent assumptions that AI alone produced the content [120]. Oversimplified “AI-generated” labels decrease belief in news headlines and willingness to share them because readers assume no human input was involved [2]. Additionally, effective disclosure should specify how, when, and where AI was used in the creation process [15]. Readers consistently express a desire for disclosure when news is produced primarily by AI with some human oversight [34] and want to know about AI involvement regardless of its impact on public opinion or perceptions of its accuracy, and trustworthiness [81]. Even though Article 50 of the EU AI Act exempts disclosure when content has been human-reviewed [33], readers still want to know the origin of the content [81]. To address such shortcomings, our work aligns with the design space proposed by Gamage et al. [38], but applies it to visually disclose the nuanced collaboration between humans and AI in journalism. We focus on iconography, and level of detail, while excluding sentiment as our approach is process-based [115], and positioning as our study focuses on design rather than placement.

Social media platforms showcase a range of AI disclosure strategies. Meta attaches “AI Info” labels with expandable details [8], TikTok uses “AI-generated” and “creator labeled as AI-generated” [101], LinkedIn employs Content Credentials metadata [24], and YouTube applies broader warnings such as “Altered or synthetic content” [41]. Indeed, AI disclosure representations have had various approaches, focusing on design, context, and impact on user understanding, trust, and engagement. Label wording strongly influences perception, terms such as “AI-generated,” “Generated with an AI tool,” or “AI Manipulated” are consistently associated with AI involvement [31]. Users prefer pictograms and explicit details about AI usage [94]. Alternative formulations have been tested, including “Content Generated by AI” [65], “Written by artificial intelligence” (translated from Dutch) [81], and variants such as “AI-Generated” or “Altered” [116]. More elaborate formats include “AI Usage Cards” for standardized reporting in scientific research [112] and certification labels that improve trust in low- and high-stakes contexts [88]. Additionally, progressive disclosure can balance transparency and cognitive load by showing basic AI involvement upfront and revealing more detail upon user request [97]. The cognitive load of the user should be taken into account, showing which elements involved GenAI, but too much technical details will be overwhelming [15]. Therefore, an effective label should enable viewers to recognize human-AI collaboration at first glance.

2.3 Information visualization techniques for AI disclosures

Information visualization is the interactive, computer-generated graphical representation of data, traditionally used to help humans understand and analyze complex datasets [18, 72]. Its goal is to transform information into intuitive and meaningful representations that provide insight [18]. Design is inevitable to this process: visualizations must balance utility, reliability, and attractiveness, while accounting for perceptual and cognitive constraints that shape user interpretation [72]. For disclosing human-AI collaboration, the main challenge is how to represent ratios clearly and effectively. Proportional visualization is common in media [93], using familiar techniques such as pie and donut charts, stacked bar

charts, tree maps, waffle charts, and numerical or textual formats [7]. Pie and donut charts are widely used despite difficulties in judging areas [96], together with the stacked bar chart, the most common visualization techniques in media [93]. Tree maps extend to hierarchical data, waffle charts represent proportions as grids, and textual notations such as percentages (“75% AI, 25% Human”) are representations for displaying proportions [25]. Although these techniques are static and simple, they provide a useful starting point. Embellishments, though sometimes seen as distracting, can in fact attract attention and support long-term memorization [25].

Design principles from HCI, such as Nielsen’s heuristics, are prominently used in interaction design. We focus on three of these heuristics for disclosure representations: *aesthetic and minimalist design*, *recognition rather than recall*, and *match between system and the real world* [75]. These map onto three core principles for disclosure design: simplicity, cognitive load management, and user-centered approach. These principles will ensure that users can accurately interpret AI involvement, avoid misperception, and interact with information at different levels of detail. AI disclosure visualizations are primarily informative rather than interactive, which makes many of Nielsen’s other heuristics less relevant, as they address complex user interactions and error recovery rather than the goal of creating clear, self-explanatory representations. Creating AI disclosure representations that are clear and free from unnecessary information is crucial for usability, aligning with Nielsen’s heuristic of aesthetic and minimalist design [75] and the design principle of simplicity in interaction [84]. Clarity is essential [109], since misleading or ambiguous representations can distort perceptions of AI’s role. When human and AI contributions are not clearly distinguished, users may overestimate or underestimate AI’s role [2]. Managing cognitive load is equally important, as users should immediately understand information without relying on memory [75]. Progressive disclosure provides one solution, offering detail on demand [30] which helps avoiding overload. Representations must also match the real world [75], using familiar formats. Additionally, interactive visualizations should not only be functional, but also intuitive and engaging [54]. These heuristics align with the specific goal of creating human-AI disclosure visualizations in journalism, ensuring that the representations we develop emphasize clarity and user friendliness while taking cognitive load into account.

3 Part 1: Designing human-AI collaboration representations

3.1 Co-design sessions

Co-design sessions were conducted to generate design ideas for human-AI collaboration representations in news media. These sessions leverage the collective creativity of participants with different ideas and knowledge to come up with a design solution together [87, 98]. The aim was to explore diverse static and interactive representations, where static designs present information at once, while interactive designs progressively reveal information [97]. Participants received a sensitizing booklet beforehand, followed by the co-design session.

3.1.1 Procedure. Sensitizing Booklet. Each participant completed a sensitizing booklet (~20 min) in Miro² (Fig. 3) which introduced concepts such as GenAI in journalism, human-AI collaboration, and AI disclosures. While also reflecting on their experiences and increasing confidence to contribute during the sessions [107]. This was completed before the co-design session, and was submitted together with a signed informed consent form.

Co-design Session. Three co-design sessions were then conducted (90 min each) in March 2025: two online via Zoom and Miro³ with drawing tablets, and one in-person at our institute using drawing materials (Fig. 4). In all sessions, designers worked in pairs, combining more experienced with less experienced participants. The sessions started with an explanation of the concepts and reflection on the booklet, followed by two ideation rounds using scenarios of varying AI contributions in the journalistic workflow, in between these rounds they were shown requirements and Information Visualization techniques to inspire, but not restrict them. During the ideation rounds they received ideation templates (Fig. 5) which looked like a wireframe of a news website with a scenario, after these rounds there was a short feedback round to foster collaboration between the teams.

3.1.2 Participants. Ten designers were recruited through snowball sampling. Three were (UX) design researchers, two were HCI experts, and five were design master’s students. The participants that joined the online session were required to have access to a drawing tablet. To focus on communication design rather than newsroom workflow, journalists were deliberately excluded. Each participant joined a single session and was compensated with a \$/€20 gift card.

3.1.3 Analysis and findings. All sketches (digital and scanned) (Fig. 6) were collected, transcribed into descriptions, and inductively coded [12]. Duplicate ideas were merged, resulting in 69 unique concepts: 29 static and 40 interactive. Static designs fell into three categories: (1) labels/icons/stamps, (2) textual explanations, and (3) more elaborate visualizations such as bar or gauge-like indicators. Interactive designs were grouped into six categories: (1) sliders/adjustable, (2) click/toggle, (3) hover/reveal, (4) pop-ups, (5) scroll/floating elements, and (6) chatbot-style disclosures. Placement varied widely, with ideas situated near headlines, images, article bodies, or underneath the article. A full overview of the results can be found in Supplementary Materials A. These ideas provided a wide design space for visualizing human-AI collaboration ratio’s from which prototypes for evaluation were selected.

3.2 Design selection

From the 69 design ideas generated, we selected four prototypes to use in the controlled user study. The set was purposefully limited to four prototypes to balance the need for diverse designs with the practical constraints of a lab-based user study, including experiment duration, multiple measurements, and participant attention. The aim was to capture diversity in modality (visual vs. textual) and interactivity (static vs. interactive), building on the design space proposed by Gamage et al. [38]. Modality and interactivity map

²Anonymized Miro board of sensitizing booklet: https://miro.com/app/board/uXjVlhVnry8=?share_link_id=270962526928

³Anonymized Miro board of co-design session: https://miro.com/app/board/uXjVLuPkmM=?share_link_id=330092544705

Criterion	Description
Human-AI collaboration	Must represent both human and AI contributions, not only AI usage.
Feasibility (technical)	Can be implemented as a working prototype for the study.
Feasibility (journalistic)	Realistic to integrate into journalistic workflow.
Simplicity	Clear [83], free from unnecessary elements, and immediately understandable [75].
Cognitive load management	Users should be able to grasp the AI-human ratio without needing to remember prior information or interpret complex details [75].
User-centered approach	The visualization should align with user expectations and avoid unnecessary technical complexity [75].
Low interaction count	Minimizes number of required interactions, which aligns with simplicity and reduces cognitive load.

Table 1: Criteria used for evaluation and selection of design ideas for prototyping.

onto usability principles of using familiar formats and progressive disclosure, ensuring clarity, feasibility, and communicative value (Sec. 2.3).

Selection criteria (Table 1) were grounded in HCI and Information Visualization principles [75, 83] and practical constraints of lab-based testing. Designs had to represent both human and AI contributions, be technically and journalistically feasible, minimize cognitive load, and avoid unnecessary complexity. This process was conducted collaboratively by two authors.

Based on these criteria, four designs were selected: (1) **Textual Disclosure (TD)** – a label-like textual explanation of GenAI model and its contributions, (2) **Role-based Timeline (RT)** – icons of person and robot to show division of work, (3) **Chatbot (C)** – a menu-based chatbot with predefined questions, and (4) **Task-based Timeline (TT)** – a workflow timeline with interactive hover element. These four design ideas were turned into working prototypes for user testing.

3.3 Prototype development

The four selected design ideas were turned into fully functional prototypes (Fig. 1). Each was developed in JavaScript, JSON, and CSS for integration in a React-based web application. All visualizations are similar in size for consistency and easy comparison.

The **TD** describes the used AI model, its contributions, and the journalist’s contributions. The **RT** visualization shows a linear progression of the roles between the journalist and AI, including icons as visual elements with a textual explanation underneath. The AI is displayed as a robot, this was also often done by the designers in the design sessions. The **C** disclosure is menu-based with predefined questions about the AI model, and human and AI contributions. The **TT** shows the five stages of news article creation, visualized with icons. Human-AI collaboration is indicated with a sparkle symbol, a neutral label commonly seen in AI context [38]. When users hover over the icons, additional information about the contributions is

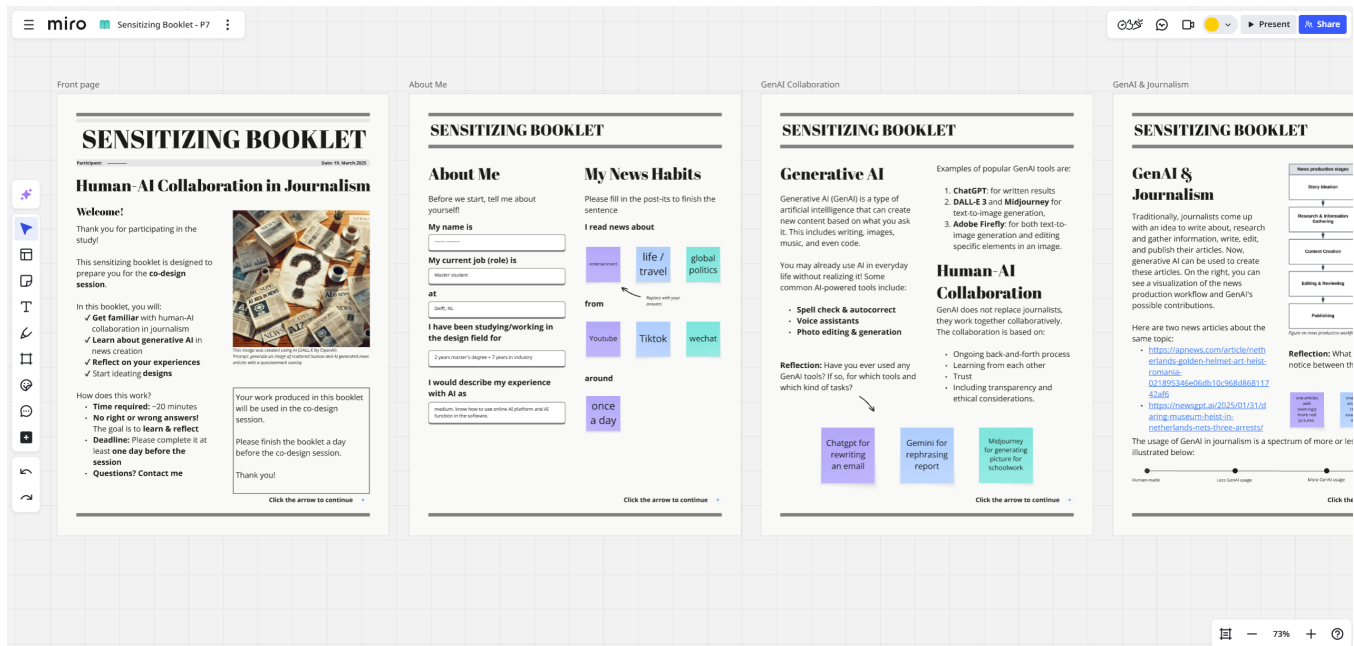


Figure 3: Screenshot of a filled out sensitizing booklet in Miro, as part of our preparation for the co-design session.



Figure 4: In-person co-design session setup.



Figure 5: Example ideation template for co-design session

revealed. Each visualization went through multiple design iterations. The eight developed prototypes (4 disclosure visualizations x 2 collaboration ratios) are included in Supplementary Material A, the four disclosure visualizations with high AI contributions are displayed in Fig. 1.

4 Part 2: Evaluating human-AI collaboration disclosure visualizations

To assess the effectiveness of the selected disclosure designs, we conducted a controlled lab-based user study. The aim was to investigate how users perceive, interpret, and interact with the different

disclosure visualizations of human-AI collaboration in online news articles. A mixed methods study was conducted, combining quantitative data from eye tracking, and questionnaires, with qualitative data from semi-structured interviews.

Eye tracking was included to understand how users visually engage with the disclosure visualizations. The self-reported questionnaires capture how participants perceived the disclosures and semi-structured interviews reveal deeper insights into their understanding, while gaze measures offer complementary insights on attention and information processing [1, 3, 50]. This mixed methods approach allows us to explore whether participants understood




Figure 6: Two examples of the results of the co-design sessions, drawn on the ideation template. The example at the top was created online in Miro, the example underneath was digitized from the in-person session.

the disclosures and how they visually navigate them, offering a comprehensive understanding on the communicative effectiveness in the news context.

4.1 Study design

Our study follows a 4 (IV1: disclosure visualizations: TD, RT, C, TT) x 2 (IV2: human-AI collaboration ratio: Primarily Human vs. Primarily AI) within-subject study design. The trials were created using two news articles. For each article, a primarily human and primarily AI contribution version was created (1-Human, 1-AI, 2-Human, 2-AI). Each of these four article versions was paired with each of the four disclosure visualizations (TD, RT, C, TT), so every disclosure visualization appeared with every article version. Therefore, each participant completed 16 trials (4 article versions x 4 visualizations), with disclosure visualizations and article version order counterbalanced using a 4x4 Latin square to minimize order

effects. All sessions were conducted between May and June, 2025 at our institute. The sessions took approximately one hour, and no tasks prior to the experiment were asked of the participants.

4.1.1 Measures. In this study, quantitative measures consisted of the following dependent variables: (a) Perceived human-AI collaboration, adapted from IBM's AI Attribution Toolkit [48], capturing the communicated proportion of human vs. AI contributions. (b) Perceived AI role per task, based on Altay et al. [2], using 7-point Likert ratings for AI involvement in tasks such as topic selection, writing, editing, fact-checking, and headline creation. (c) Perceived clarity and value, 7-point Likert-scale adapted from Wittenberg et al. [116], including items such as clear, helpful, easy to understand, useful, and informative. (d) Perceived information, 7-point Likert ratings of whether the visualization conveyed overview, in-depth information, specific workflow steps, or who contributed what. (e) Trial duration, logged automatically from the moment a disclosure appeared until the questionnaire was submitted. Although differences were expected between static and interactive visualizations, this measure provided insight into the time needed for each disclosure type, which may influence real-world adoption. (f) Eye-tracking metrics, recorded within predefined Areas of Interest (AOIs): total gaze duration, mean fixation duration, fixation count, saccade count, and mean saccade length [50].

These measures assess whether readers understood who contributed what to the article, such understanding improves transparency, can reduce misinformation belief and increase trust in newsroom processes [15, 120]. If a disclosure fails to communicate human-AI collaboration clearly, it cannot fulfill its transparency purpose, and its effect on authorship [55]. Combined with measures of clarity, information type, and visual engagement of the disclosures, they create a comprehensive assessment of disclosure effectiveness. Therefore, these measures are not manipulation checks but core outcomes that indicate whether a disclosure visualization communicates the information needed for meaningful transparency in journalism. Qualitative data were gathered after the trials. Consensually audio-recorded semi-structured interview was conducted, focusing on participants' overall interpretation of disclosures, design preferences, and real-world application. The list of interview questions can be found in the Supplementary Material B. Our study followed strict guidelines from our institute's ethics and data protection committee.

4.2 Study procedure

Our study procedure is shown in Fig. 7, and lasted approximately 60 minutes. Participants first read and signed the informed consent form, and filled in a short pre-task questionnaire about demographics, news consumption habits, and their AI experience. Vision screening questions ensured participants did not have impairments that could interfere with eye tracking.

A verbal overview of the study was provided, explaining the structure of the trials and how to navigate the web interface. Participants were informed that they would be reading news articles and evaluating disclosure visualizations, but were not told the specific disclosure visualizations in advance to maintain a natural first impressions. The Tobii Pro Fusion eye tracker was calibrated at the start of each session.




Figure 7: User study procedure. The four Articles refer to the four article versions (1-Human, 1-AI, 2-Human, 2-AI), counterbalanced between participants using a Latin square.

Following calibration, the main experiment task started, consisting of 16 trials (4 article versions \times 4 visualizations). Each trial followed the same order: (1) the news article appeared without disclosure, (2) the disclosure visualization was displayed beneath the article, and (3) a questionnaire appeared on the right side of the screen. Participants continued through these steps by pressing a button on the interface. This sequential order ensured clear mapping of eye tracking data to pre-defined Areas of Interest (AOIs) [49].

After completing the 16 trials, participants filled out a short questionnaire evaluating the four disclosure visualizations on qualities such as perceived clarity, value and informativeness. They then took part in a semi-structured interview, where they reflected on their preferences, interpretation of the disclosures, and possible use in real-world news contexts. All interviews were audio-recorded for later transcription and thematic analysis. Participants were compensated with a \$/€10 gift card.

4.2.1 Participants. 32 participants⁴ (19 female, 12 male, 1 non-binary) were recruited through multiple platforms and snowball sampling. Ages ranged from 22 to 72. Educational backgrounds were varied (completed high school, bachelor, master, and doctorate), and news consumption habits ranged from daily to rarely. AI literacy, measured by an adapted⁵ subset of items from the modular 11-point MAIIS scale [17], averaged 5.69 (SD = 1.58). The MAIIS items were selected from four dimensions: AI understanding, detection, ethical awareness, and persuasion literacy. These align most closely to how readers interpret and evaluate AI involvement in news production (the questionnaires can be found in Supplementary Material B). Cronbach’s $\alpha = .87$ showed good internal consistency for the selected items [99]. Average self-reported ChatGPT experience on a 1–7 scale (7 = very experienced) was 4.22 (SD = 1.96). The only requirement was that participants did not have significant visual impairments.

4.3 Hardware and software setup

4.3.1 Hardware. Participants were seated at a desk with a 25-inch monitor at a resolution of 2560 \times 1440 pixels, with a Tobii Pro Fusion screen-based eye tracker⁶ mounted below the screen (Fig. 8). The system was connected to a laptop running Tobii Pro Lab, which was used for calibration, data collection, and later analysis.

4.3.2 Software. Stimuli were presented via a custom built React web application hosted locally, to present the stimuli and collect the questionnaire responses. The interface divided the screen into two panes: the left displayed the stimulus (article and visualization) in a realistic news website format, and the right contained the questionnaire (Fig. 8).

Two real low-stake news articles were selected from the MiRA-GeNews dataset [53], originally published by BBC⁷. Both articles were adapted into two versions to manipulate the human-AI collaboration ratio, adapted from the method of Gilardi et al. [39]:

- (1) Primarily human-contributions (Article 1-Human and Article 2-Human): The original article was used as input for ChatGPT-4o to generate a headline and change the style. This prompt mimicked the Gilardi et al. [39] prompt style: “Here is a newspaper article: [article text]. As an AP journalist, write a headline for this article and do small adjustments to the text to keep it in AP style.” This generated an article that was almost entirely written by the journalist, but still collaborating with GenAI.
- (2) Primarily AI-contributions (Article 1-AI and Article 2-AI): Only the article headline was provided to ChatGPT-4o, using an adapted prompt from Gilardi et al. [39]: “As an AP journalist, you must write a 400-word article on this topic: [article headline].” This generated an article that was almost entirely written by GenAI.

The full news articles used as stimuli can be found Supplementary Materials B. Each version was paired with the four prototypes (TD, RT, C, TT; see Fig. 1), creating 16 unique stimuli. The information represented in these prototypes were derived from these controlled

⁴For effect size $f = .25$ under $\alpha = .05$ and power $(1 - \beta) = .95$, with 16 repeated measurements within factors, one would need a minimum of 15 participants.

⁵We changed the term in Q5 from ‘devices’ to ‘technology’ to ensure the scale is meaningful to participants, since we do not focus on devices per se.

⁶<https://www.tobii.com/products/eye-trackers/screen-based/tobii-pro-fusion>

⁷<https://www.bbc.com/news/world-europe-48905867>, and <https://www.bbc.com/news/world-49760240>

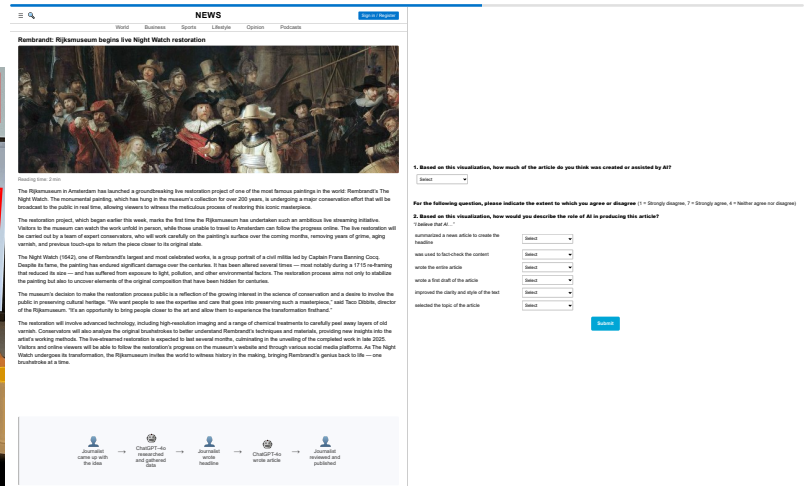
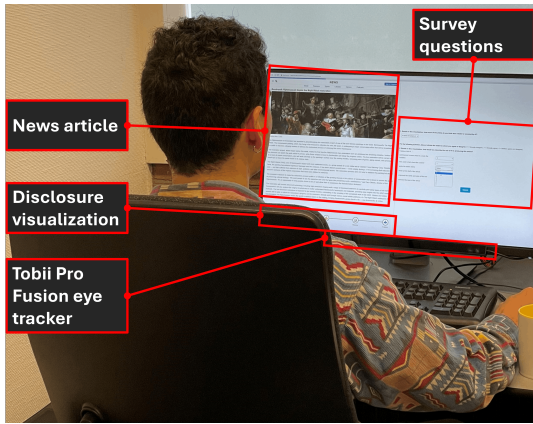


Figure 8: User study setup

prompt-based manipulations to create the article versions. Disclosure visualizations were consistently placed at the bottom of the article to control for placement effects.

4.4 Quantitative results

We report our analysis of time spent on task, perceived human-AI collaboration, AI role, perceived clarity and value, perceived information, and eye tracking metrics. The ordinal response variables (perceived human-AI collaboration, AI role, and clarity and informativeness ratings) were analyzed using cumulative link mixed-effects models (CLMMs) [20] to take the ordered nature of a Likert-scale into account. Continuous variables (time spent and eye tracking metrics) were analyzed using linear mixed-effects models (LMM) [5]. Both models included participants as random intercepts to model the repeated measures design [62], and self-reported AI literacy (scale 0-10) and ChatGPT experience (scale 1-7) as continuous covariates. Interactivity (static vs. dynamic visualization) was not included as a random effect, as two-level grouping factors do not provide enough information to reliably estimate random-effect variance [78]. These mixed-effects models allowed us to model individual differences while testing fixed effects of visualization type and collaboration condition where applicable. To control for potential inflation of Type I Error, False Discovery Rate (FDR) correction [47] was applied to adjust the p -value across the six perceived AI role features, the ten perceived clarity, value, and information features, and for all post hoc comparisons. Due to the exploratory nature of this study, FDR was chosen instead of more conservative Bonferroni or Tukey adjustments. Effect sizes for CLMMs are reported as odds ratios (OR) with 95% confidence intervals (CI95%) for pairwise contrasts. For LMMs, effect sizes are reported as partial η^2 for main effects and Cohen's d for pairwise contrasts [21].

Eye tracking data were collected using a Tobii Pro Fusion with Tobii Pro Lab, and only included when disclosure visualizations were visible, during both initial display and questionnaire answering. Only raw data inside the disclosure AOI (Fig. 9) were processed

with the Tobii I-VT (Fixation) filter, which classifies gaze movements slower than $30^\circ/s$ as fixations and faster as saccades, with additional parameters to remove short fixations and merge adjacent ones [79]. Data were considered valid if at least 80% of gaze points were valid [102], data points were included if a valid measure was available from at least one eye [52]. Validity was assessed per participant, leading to the exclusion of eight. Tobii defines good accuracy as below 0.8° and precision SD below 1.5° [102], but due to the large Areas of Interest we adjusted the average accuracy threshold to 1.6° , resulting in a final dataset of 20 participants. CLMM and LMM results for all response variables and gaze features on disclosure visualization main effect are shown in Table 2. Full contrast test tables are provided in Supplementary Material C.

Variable	χ^2	df	p
Clarity	17.77	3	0.003 **
Informativeness	13.60	3	0.018 *
Helpfulness	0.91	3	0.935
Easiness to understand	28.02	3	<.001 ***
Usefulness	2.27	3	0.890
Overview	33.78	3	<.001 ***
In-depth information	41.48	3	<.001 ***
Understanding of specific steps	26.43	3	<.001 ***
Who contributed what	5.03	3	0.637
Sense of AI vs. human input	4.25	3	0.642

Variable	F	Df	p	$\eta^2 p$
Gaze duration	31.42	3	<.001 ***	0.62
Fixation duration	45.96	3	<.001 ***	0.71
Fixation count	44.85	3	<.001 ***	0.70
Saccade count	34.44	3	<.001 ***	0.64
Saccade length	11.11	3	<.001 ***	0.37

Table 2: Cumulative link mixed model results for response variables and linear mixed model results for gaze features, testing main effect of disclosure visualization. P values were adjusted with FDR corrections.

Article

World Business Sports Lifestyle Opinion Podcasts Sign In Register

Rembrandt: Rijksmuseum begins live Night Watch restoration

Reading time: 2 min

The Rijksmuseum in Amsterdam has launched a groundbreaking live restoration project of one of the most famous paintings in the world: Rembrandt's *The Night Watch*. The monumental painting, which has hung in the museum's collection for over 200 years, is undergoing a major conservation effort that will be broadcast to the public in real time, allowing viewers to witness the meticulous process of restoring this iconic masterpiece.

The restoration project, which began earlier this week, marks the first time the Rijksmuseum has undertaken such an ambitious live streaming initiative. Visitors to the museum can watch the work unfold in person, while those unable to travel to Amsterdam can follow the progress online. The live restoration will be carried out by a team of expert conservators, who will work carefully on the painting's surface over the coming months, removing years of grime, aging varnish, and previous touch-ups to return the piece closer to its original state.

The *Night Watch* (1642), one of Rembrandt's largest and most celebrated works, is a group portrait of a civil militia led by Captain Frans Banning Cocq. Despite its fame, the painting has endured significant damage over the centuries. It has been altered several times — most notably during a 1715 re-framing that reduced its size — and has suffered from exposure to light, pollution, and other environmental factors. The restoration process aims not only to stabilize the painting but also to uncover elements of the original composition that have been hidden for centuries.

The museum's decision to make the restoration process public is a reflection of the growing interest in the science of conservation and a desire to involve the public in preserving cultural heritage. "We want people to see the expertise and care that goes into preserving such a masterpiece," said Taco Dibbits, director of the Rijksmuseum. "It's an opportunity to bring people closer to the art and allow them to experience the transformation firsthand."

The restoration will involve advanced technology, including high-resolution imaging and a range of chemical treatments to carefully peel away layers of old varnish. Conservators will also analyze the original brushstrokes to better understand Rembrandt's techniques and materials, providing new insights into the artist's working methods. The live-streamed restoration is expected to last several months, culminating in the unveiling of the completed work in late 2025. Visitors and online viewers will be able to follow the restoration's progress on the museum's website and through various social media platforms. As *The Night Watch* undergoes its transformation, the Rijksmuseum invites the world to witness history in the making, bringing Rembrandt's genius back to life — one brushstroke at a time.

Disclosure

1. Based on this visualization, how much of the article do you think was created or assisted by AI?

Select

For the following question, please indicate the extent to which you agree or disagree (1 = Strongly disagree, 7 = Strongly agree, 4 = Neither agree nor disagree)

2. Based on this visualization, how would you describe the role of AI in producing this article?

7 delivered (1/2 AI)

summarized a news article to create the headline Select

was used to factcheck the content Select

wrote the entire article Select

wrote a first draft of the article Select

improved the clarity and style of the text Select

selected the topic of the article Select

Submit

Survey

Figure 9: Heatmap of the gaze points of one participant on a stimulus. The red boxes indicate the predefined AOIs.

4.4.1 Mean time spent per trial. We first looked at the time participants spent during each trial (i.e. from the moment the disclosure was displayed until the survey was submitted for that disclosure), as an indicator of potential disengagement or abandonment. As expected, interactive disclosures require more time than static disclosures. Therefore, the mean trial duration is reported to contextualize readers engagement rather than as a performance outcome. On average, time spent per trial across disclosure visualization conditions (in seconds) were: Chatbot ($M=150, SD=126$), Task-based Timeline ($M=111, SD=110$), Textual Disclosure ($M=116, SD=123$), Role-based Timeline ($M=110, SD=108$). A linear mixed-effects model was used to account for repeated measures and individual differences in AI literacy and ChatGPT experience. Results revealed significant main effects of disclosure visualization, $F(3, 474) = 3.55, p = .014$. Post hoc contrasts showed that the Chatbot trials took significantly longer compared to all other conditions (vs. TD: $p=0.035, d=0.30$, vs. RT: $p=0.020, d=0.35$, vs. TT: $p=0.020, d=0.34$). Additionally, both covariates had a significant influence, higher AI literacy was predicted a longer time spent per trial ($\beta=11.64 \text{ sec}, p=.012$), whereas more ChatGPT experience predicted less time spent ($\beta=-9.52 \text{ sec}, p=.011$).

4.4.2 Perceived human-AI collaboration dynamics. How the human-AI collaboration was perceived compared to the true collaboration ratios is shown in Fig. 10. As expected, primarily AI-generated articles (1-AI and 2-AI) were mostly recognized as such, and primarily human-written articles (1-Human and 2-Human) were typically seen as "Primarily human-created." Across article

versions, there were responses for "Human-AI blend," indicating that some participants assumed creation with an even blend of human and AI contributions [48].

We grouped the responses into three categories (Human, Blend, AI) to ensure sufficient counts per category and analyzed with a cumulative link mixed model (CLMM) to account for the ordinal categories and repeated measures. Pairwise contrasts revealed context-dependent shifts. For primarily human articles, the Role-based Timeline significantly increased perceptions of AI contribution relative to the Textual Disclosure ($\beta = -1.30, SE = 0.40, z = -3.26, p = .002, OR(95\%CI) = 0.27(0.08, 0.95)$), Chatbot ($\beta = 1.62, SE = 0.42, z = 3.82, p < .001, OR(95\%CI) = 5.06(1.35, 19.02)$), and Task-based Timeline ($\beta = 1.25, SE = 0.39, z = 3.19, p = .002, OR(95\%CI) = 3.49(1.02, 11.92)$). For primarily AI articles, the Task-based Timeline nudged perceptions toward human involvement, showing significant differences with Textual Disclosure ($\beta = 1.05, SE = 0.44, z = 2.42, p = .022, OR(95\%CI) = 2.86(0.73, 11.14)$), Chatbot ($\beta = 0.95, SE = 0.43, z = 2.21, p = .036, OR(95\%CI) = 2.58(0.68, 9.79)$), but not compared to Role-based Timeline. Fig. 11 illustrates the response distributions, where these shifts appear most clearly for the Role-based Timeline (increasing Blend/AI responses in primarily human texts) and the Task-based Timeline (increasing Human responses in primarily AI texts). These findings validate that these disclosure visualizations can effectively convey nuanced impressions of human-AI collaboration, highlighting their

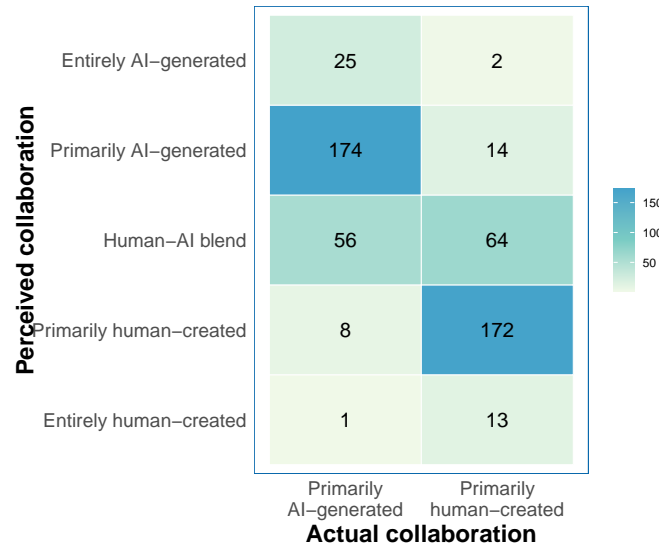


Figure 10: Confusion matrix of human-AI collaboration per article

potential to communicate underlying dynamics beyond simple labels. Yet, they may also introduce bias perceptions, shifting toward more or less AI involvement depending on the context.

4.4.3 Perceived AI role per article. CLMMs revealed that the perceived role of AI can differ by disclosure visualization and that AI literacy and ChatGPT can also shape this (Fig. 12). Significant main effects on collaboration were expected due to the difference in responses for Primarily Human or AI. As seen in Table 3, fact-checking showed a main effect of visualization ($p < .001$). Post-hoc tests revealed a main effect on visualization type. Across both contexts (Primarily Human and Primarily AI), Role-based Timeline made the perception of the role of AI in fact-checking more than Textual Disclosure ($\beta = -0.84$, $SE = 0.25$, $z = -3.30$, $p = .002$, $OR(95\%CI) = 0.43(0.22, 0.84)$), Chatbot ($\beta = 1.01$, $SE = 0.26$, $z = 3.95$, $p < .001$, $OR(95\%CI) = 2.75(1.40, 5.40)$) and Task-based Timeline ($\beta = 0.91$, $SE = 0.25$, $z = 3.62$, $p < .001$, $OR(95\%CI) = 2.48(1.28, 4.80)$). For headline creation, higher AI literacy predicted lower perceived AI involvement ($OR(95\%CI) = 0.78(0.66, 0.91)$, $p = .006$). Similarly, for topic selection, higher AI literacy decreased perceived AI involvement ($OR(95\%CI) = 0.53(0.33, 0.85)$, $p = .022$), but more ChatGPT experience ($OR(95\%CI) = 1.76(1.19, 2.60)$, $p = .021$) increased the perceived contribution of AI in topic selection. These findings suggest that disclosure visualizations can shape how readers attribute AI contributions to specific tasks (particularly fact-checking), and that the reader's AI literacy and ChatGPT experience plays a role in shaping these perceptions.

4.4.4 Perceived clarity. All visualizations were perceived clearly, with means above the midpoint. Fig. 13 shows the distributions of all response variables, with significance bars from CLMMs. There was a significant main effect of disclosure visualization on perceived clarity, $p = .003$. Post hoc comparisons showed that both timelines

Response Variable	Factor	χ^2	df	p
Headline	Visualization	4.03	3	0.259
	Collaboration	256.56	1	< .001 ***
	Visualization × Collaboration	5.97	3	0.141
	<i>AI literacy</i>	9.32	1	0.006 **
	<i>ChatGPT experience</i>	2.94	1	0.141
Fact-check	Visualization	20.16	3	< .001 ***
	Collaboration	91.30	1	< .001 ***
	Visualization × Collaboration	5.79	3	0.204
	<i>AI literacy</i>	0.20	1	0.651
	<i>ChatGPT experience</i>	0.21	1	0.651
Entire article	Visualization	7.11	3	0.114
	Collaboration	262.05	1	< .001 ***
	Visualization × Collaboration	2.81	3	0.423
	<i>AI literacy</i>	3.43	1	0.114
	<i>ChatGPT experience</i>	1.58	1	0.261
First draft	Visualization	0.76	3	0.860
	Collaboration	247.03	1	< .001 ***
	Visualization × Collaboration	8.56	3	0.089
	<i>AI literacy</i>	0.04	1	0.860
	<i>ChatGPT experience</i>	0.09	1	0.860
Style/clarity	Visualization	3.99	3	0.328
	Collaboration	188.34	1	< .001 ***
	Visualization × Collaboration	4.77	3	0.316
	<i>AI literacy</i>	2.39	1	0.305
	<i>ChatGPT experience</i>	0.28	1	0.598
Topic	Visualization	4.04	3	0.278
	Collaboration	1.18	1	0.278
	Visualization × Collaboration	6.53	3	0.148
	<i>AI literacy</i>	6.88	1	0.022 *
	<i>ChatGPT experience</i>	8.16	1	0.021 *

Table 3: Cumulative link mixed model results for perceived AI role with per-outcome FDR correction. p -values corrected within each outcome. Covariates are in italics.

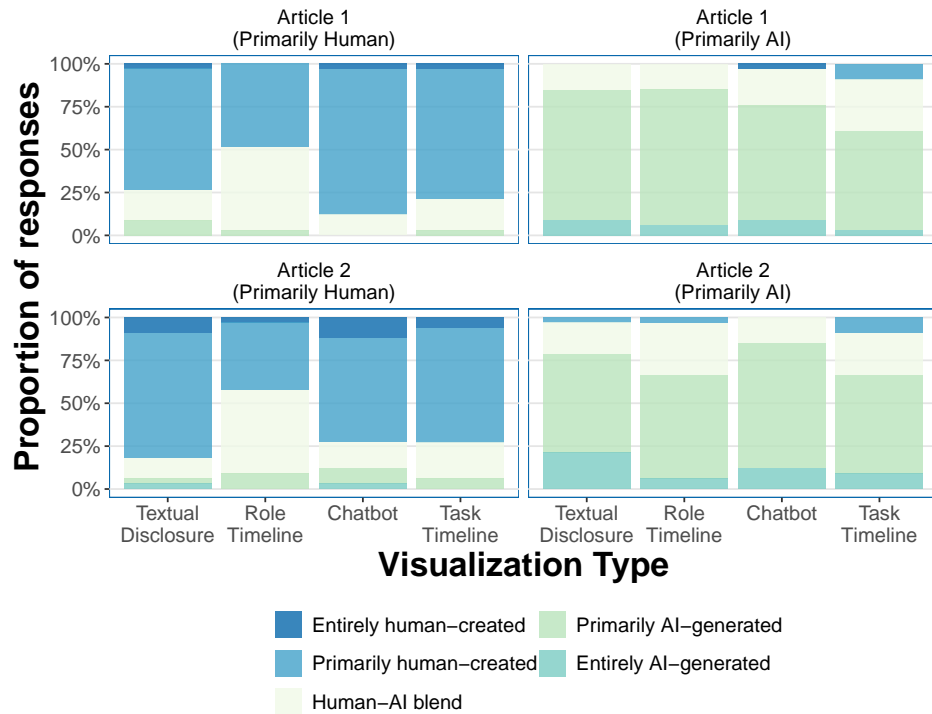


Figure 11: Proportion of responses of perceived human-AI collaboration per article and disclosure visualization

were perceived most clear. Role-based Timeline was rated significantly clearer than Textual Disclosure ($\beta=-1.55$, $SE=0.50$, $z=-3.08$, $p=.006$, $OR(95\%CI)=0.21$ (0.06, 0.80)) and Chatbot ($\beta=1.96$, $SE=0.52$, $z=3.79$, $p<.001$, $OR(95\%CI)=7.11$ (1.81, 27.89)). Task-based Timeline was also rated clearer than Chatbot ($\beta=-1.43$, $SE=0.50$, $z=-2.86$, $p=.008$, $OR(95\%CI)=0.24$ (0.06, 0.89)).

4.4.5 Perceived informativeness. Disclosure visualization significantly affected perceived informativeness, $p = .018$. Post hoc comparisons showed that the Chatbot was rated as most informative compared to all other visualizations (TD: $\beta=-1.64$, $SE=0.49$, $z=-3.35$, $p=.004$, $OR(95\%CI)=0.19$ (0.05, 0.71), RT: $\beta=-1.57$, $SE=0.49$, $z=-3.20$, $p=.004$, $OR(95\%CI)=0.21$ (0.06, 0.76), TT: $\beta=1.30$, $SE=0.49$, $z=2.63$, $p=.017$, $OR(95\%CI)=3.66$ (0.99, 13.49)).

4.4.6 Perceived easiness to understand. A significant main effect of disclosure visualization was found for perceived easiness to understand, $p < .001$. Post hoc tests showed that the Chatbot was significantly less easy to understand than the other visualizations (Textual Disclosure: $\beta=1.54$, $SE=0.49$, $z=3.13$, $p=.004$, $OR(95\%CI)=4.67$ (1.27, 17.18), Role-based Timeline: $\beta=2.89$, $SE=0.58$, $z=4.99$, $p<.001$, $OR(95\%CI)=17.97$ (3.90, 82.79), and Task-based Timeline: $\beta=-2.27$, $SE=0.54$, $z=-4.19$, $p<.001$, $OR(95\%CI)=0.10$ (0.02, 0.43)). Additionally, the Role-based Timeline was significantly easier to understand than Textual Disclosure ($\beta=-1.35$, $SE=0.52$, $z=-2.58$, $p=.015$, $OR(95\%CI)=0.26$ (0.07, 1.03)). Differences between Role-based Timeline and Task-based Timeline were not significant.

4.4.7 Providing an overview. CLMM revealed a significant main effect of disclosure visualization, $p < 0.001$. Post hoc comparisons showed that the other three visualizations were significantly better at providing an overview than the Chatbot: Role-based Timeline ($\beta=2.82$, $SE=0.56$, $z=5.01$, $p<.001$, $OR(95\%CI)=16.70$ (3.79, 73.53)), Task-based Timeline ($\beta=-3.14$, $SE=0.58$, $z=-5.42$, $p<.001$, $OR(95\%CI)=0.04$ (0.01, 0.20)), and Textual Disclosure ($\beta=1.59$, $SE=0.51$, $z=3.14$, $p=.003$, $OR(95\%CI)=4.90$ (1.29, 18.59)). Furthermore, the Role-based and Task-based Timeline provided a significantly better overview than Textual Disclosure (RT: $\beta=-1.23$, $SE=0.49$, $z=-2.50$, $p=.015$, $OR(95\%CI)=0.29$ (0.08, 1.07), TT: $\beta=-1.55$, $SE=0.51$, $z=-3.05$, $p=.003$, $OR(95\%CI)=0.21$ (0.06, 0.81)).

4.4.8 Providing in-depth information. There is a significant main effect of disclosure visualization $p < 0.001$. Post hoc comparisons showed that the Chatbot visualization was significantly better at providing in-depth information compared to the other three visualizations: Textual Disclosure ($\beta=-3.58$, $SE=0.60$, $z=-5.99$, $p<.001$, $OR(95\%CI)=0.03$ (0.01, 0.13)), Role-based Timeline ($\beta=-3.49$, $SE=0.59$, $z=-5.88$, $p<.001$, $OR(95\%CI)=0.03$ (0.01, 0.15)), and Task-based Timeline ($\beta=2.95$, $SE=0.58$, $z=5.10$, $p<.001$, $OR(95\%CI)=19.12$ (4.15, 88.03)).

4.4.9 Providing an understanding of specific steps in the process. CLMM results for providing an understanding in specific steps in the process also revealed a significant main effect of visualization type $p < .001$. Post hoc comparisons showed that the Textual Disclosure performed significantly worse than all other visualizations at showing the specific steps: Role-based Timeline ($\beta=-1.65$, $SE=0.48$,




Figure 12: Perceived AI involvement per role by visualization and collaboration ratio. Significance bar in black refer to significant main effect of visualization. Dashed line show average value per article version.

$z=-3.42$, $p=.002$, $OR(95\%CI)=0.19$ (0.05, 0.69)), Chatbot ($\beta=-1.05$, $SE=0.45$, $z=-2.33$, $p=.030$, $OR(95\%CI)=0.35$ (0.11, 1.15)), and Task-based Timeline ($\beta=-2.54$, $SE=0.51$, $z=-5.01$, $p<.001$, $OR(95\%CI)=0.08$ (0.02, 0.30)). Additionally, the Task-based Timeline was significantly better at showing the specific steps than the Chatbot ($\beta=-1.50$, $SE=0.48$, $z=-3.09$, $p=.004$, $OR(95\%CI)=0.22$ (0.06, 0.80)).

4.4.10 Gaze duration. A linear mixed-effects model showed a significant main effect of disclosure visualization, $F(3,57) = 31.4$, $p<0.001$, partial $\eta^2=.62$. Fig. 14 shows the distributions of all gaze features, with significance bars from LMMs. Post hoc comparisons showed that the Chatbot had significantly longer gaze durations than all other visualizations (all $p<0.0001$), with very large effect sizes (vs. RT: $d=2.58$, vs. TD: $d=2.41$, vs. TT: $d=2.52$). Which is in line with the time spent on the surveys (Sec. 4.4.1), and with the expectations as the Chatbot required more interactions. Additionally, individuals with higher AI literacy score had longer gaze durations ($\beta = 23.57$, $p = .019$), while users with more ChatGPT experience

had shorter gaze durations ($\beta = -18.95$, $p = .013$), as more experience requires less time to process information.

4.4.11 Fixation duration. A significant main effect of disclosure visualization was also found for fixation duration, $F(3,57) = 45.96$, $p < 0.001$, partial $\eta^2=.71$. Participants had significantly longer fixations for Role-based Timeline than all other types (all $p < 0.001$, vs. TD: $d=3.27$, vs. C: $d=3.15$, vs. TT: $d=1.99$). The fixations on Task-based Timeline were also significantly longer compared to Textual Disclosure ($p < 0.001$, $d = 1.28$), and Chatbot ($p < 0.001$, $d = 1.16$). The Chatbot did not differ significantly from Textual Disclosure. Participants with more experience in using ChatGPT had a significantly shorter fixation duration ($\beta = -19.45$, $p = .002$).

4.4.12 Fixation count. There was a significant main effect of disclosure visualization for fixation count, $F(3, 57) = 44.85$, $p < 0.001$, partial $\eta^2=.70$. The Chatbot had significantly more fixations compared to all other visualizations (all $p < 0.001$, vs. TD: $d=2.69$, vs. RT: $d=3.22$, vs. TT: $d=2.98$). The fixation count was also influenced




Figure 13: Plots of response variables per disclosure visualization.

by the AI literacy of the participants, as a higher literacy had an increase in fixation count ($\beta = 69.26, p = .021$).

4.4.13 Saccade count. For saccade count there was also a significant main effect found, $F(3, 57) = 34.44, p < 0.001$, partial $\eta^2 = .64$. Similar to fixation count, saccade counts were significantly higher for Chatbot than all others (all $p < 0.001$, vs. TD: $d = 2.33$, vs. RT: $d = 2.85$, vs. TT: $d = 2.59$). For this feature, both ChatGPT experience and AI literacy were significant covariates. AI literacy increased the saccade count ($\beta = 66.05, p = .046$), while ChatGPT experience decreased this count ($\beta = -49.89, p = .045$).

4.4.14 Saccade length. There was a significant main effect of disclosure visualization, $F(3, 57) = 11.11, p < 0.001$, partial $\eta^2 = .37$. Post hoc test showed that Role-based Timeline had significantly longer saccades than all other visualizations (vs. TD: $p < 0.001, d = 1.33$, vs. TT: $p = 0.003, d = 1.06$, vs. C: $p < 0.001, d = 1.75$).

4.5 Qualitative results

We analyzed the data with deductive thematic analysis [12]. Two authors separately analyzed the qualitative data before meeting to discuss our analysis. First, we coded it according to the interview questions. Then within each topic, we analyzed emerging themes. Recurring ideas across participants were noted down, but due to the semi-structured nature of the interview, it is important to mention that if a participant did not mention something, it does not mean they disagree. Participants are labeled P1-P32. Four themes emerged: initial interpretations of the visualizations, preferences for different designs, reflections on interactivity and information structure, and disclosures in real-world news contexts.

4.5.1 Interpretation of the visualizations. The initial reaction of participants on the disclosures was overall positive. They were described as “useful and helpful for the average consumer reading the news” [P19], and “illustrative, functional” [P30]. Three participants mentioned they had never seen these kinds of visualizations before “I hadn’t really seen those before, so I was kind of surprised by them.




Figure 14: Plots of gaze features per disclosure visualization.

And I was just curious to see what was done by journalists and what was done by an AI tool” [P14]. Four participants appreciated the transparency “I think it’s nice to be transparent about AI use” [P4].

Participants also reflected on how the visualizations shaped their understanding of how the article was created. Seven said they could not initially tell AI was used “I didn’t really have an idea on what was done by AI and what was done by the journalist” [P14]. For six others, the visualizations matched their expectations “this is written by AI and then the visualization confirmed it for me” [P17]. Additionally, two participants mentioned the visualizations created more nuance “It still made me change from the opinion that this was an AI generated article to actually a blend of AI and human generated article” [P16]. However, two other participants also expressed some skepticism “if it’s stated like a journalist wrote an article, I’m not sure whether to believe that or not” [P1], “I feel like some parts were not true. Maybe it’s just the wording in it” [P9].

4.5.2 *User preferences per disclosure visualization.* Overall, preferences were divided, but clear patterns emerged. Task-based Timeline was most frequently ranked first, described as “the best combination of informative, in-depth and visually clear as well. So it immediately provides you an overview” [P25]. Yet, nine participants overlooked the hover interaction, with several stating they “didn’t know [it] was interactive” [P5]. Some disliked the interaction entirely, “I didn’t like having to hover. If you’re going to do a visualization, then just show it directly” [P7]. Suggestions included changing the element to “a small text, above or below the circles, that you could click if you wanted to dive deeper” [P25]. Role-based Timeline was also frequently favored, as it was clear and easy to understand: “the level of AI against human input is clearly stated here” [P32], and “the pictures make it more easy to understand” [P27]. Participants also liked that “it gives a good overview and it’s to the point, you don’t have to hover” [P20]. Still, others found it less detailed than the task-based version, “it did not provide as much detail as the task-based timeline” [P16], and potentially misleading, “it gives you

the idea that the roles are done half half, then you need to read it [P12].

Chatbot was described as most informative, valued for depth and transparency: *“I liked the ability to interact with it...it could explain what AI can do. It has so many possibilities to really go in depth, and especially with journalism, I’d like it to be very transparent and adherent to integrity”* [P11]. It was also appreciated for giving control: *“it was nice to be given the option to consume the information that you are interested in”* [P31]. However, it was often critiqued as overwhelming, *“a bit too much information. I’m not sure if people are really going to read this”* [P8], and some struggled to navigate between questions. It was seen as clear and novel, but lacking an overview. By contrast, Textual Disclosure was often ranked lowest, described as *“too textual, looks like a legal disclaimer rather than an informative thing”* [P10]. Others noted it *“didn’t grab my attention”* [P11] and was *“not very detailed”* [P16]. Still, seven participants valued its simplicity, *“it has a short overview of all the information needed”* [P14], appreciating that it was easy to understand.

4.5.3 Interaction style and information structure. Preferences reflected broader design considerations: interaction style (static vs. interactive) and information structure (role-based vs. task-based). Nine participants favored interactive disclosures for exploration and engagement, e.g., *“I understand the process better when it’s interactive because you actively click on something because you want to know what happened”* [P3], and control of information was appreciated *“you can look for more information if you want it”* [P15]. In contrast, six participants favored the static disclosures for the overview and completeness. Some doubted they would interact with it on a news page *“I don’t think it’s really used if it would be in an article”* [P5] and liked having *“all the information already printed out in a schematic way. So it’s good that you don’t need to even move a finger to get other information”* [P10].

Participants were also divided in their preferences for representing information as role or task-based. Fourteen participants preferred role-based structures, *“I was more concerned about the roles since I already knew the steps”* [P10]. On the other hand, eleven participants preferred the task-based structures *“I felt the process depicts a timeline better”* [P31]. Remarkably, familiarity with the journalistic workflow shaped preferences. A participant with writing experience preferred task-based information: *“Because I write articles. So then I want to know what they did and every step”* [P15]. This suggests that disclosure visualizations should account for varying levels of domain familiarity among news readers.

4.5.4 AI disclosures in real-world news. Participants reflected on how these disclosures might influence their engagement with news. Five noted the impact would depend on the topic, especially political or societal issues: *“then I think it’s really important some things are properly fact checked. And then I would want to know who did what”* [P11]. Thirteen emphasized disclosures as more important in high-impact contexts, *“especially who has done the research”* [P8] and *“the specific steps in which AI is involved become way more important”* [P25]. Nine participants wanted more detail about article creation, for example: *“if it’s more high-impact... then I would want to know if AI made that assumption or a person”* [P30]. Two stated they would not read AI-generated news at all: *“I wouldn’t want to read an article that would be partly created by AI if it would be a critical article.*

Because I think you should think about the impact that an article can make. So then I think it is more wise to only have a journalist write it” [P5]. In contrast, two participants did not care about AI usage: *“I don’t really want to know if the AI did it or the human did it. I don’t care actually”* [P2]. Two others said they would scan more quickly if they knew an article was AI-written. At the same time, eight said disclosures would make them read more critically, *“if you see how much AI has done... you might become a bit more critical or pay more attention”* [P21]. Two others noted the visualizations gave them a better sense of AI’s capabilities: *“it made me realize that AI could do a lot more than I thought”* [P18].

Fifteen participants would like to see such visualizations on news platforms, as *“transparency is really important”* [P17]. However, five expressed skepticism, they mentioned *“It’s like a romantic idea of journalism and investigative journalism still being very much primarily human based. If I were to see an AI label, I’d probably become skeptical”* [P11]. Two said they might even stop reading outlets that disclosed AI usage: *“I would steer more towards other platforms or news sites if I knew that that platform uses AI”* [P5]. Placement preferences were divided. Fifteen wanted disclosures upfront for ethical transparency, e.g., *“I think it’s more ethical to make people aware as soon as possible”* [P11]. Eleven preferred them at the bottom to avoid distraction: *“I think it takes away from the article if you immediately show this was generated with AI”* [P30]. A few suggested optional or hybrid approaches, such as *“a small icon above the article and then the full visualization under it”* [P25], or side placement *“just because sometimes maybe you don’t finish the whole article”* [P4].

5 Discussion

5.1 How can we visually represent the ratios of human-AI collaboration output in news production and beyond?

A primary goal of this research (RQ1) was to explore how human-AI collaboration in news production can be effectively communicated to readers through visual disclosures. To address this, we drew on principles from HCI and Information Visualization, and involved designers and HCI experts in co-design sessions. These sessions generated 69 design ideas, spanning both static and interactive formats, and explored a wide design space of how to represent human-AI collaboration ratios. Despite this broad scope, we found two overarching strategies emerged: ratio-focused and process-focused. Ratio-focused designs emphasized proportions through familiar visual methods such as pie charts, bar charts, percentages, or stamps, echoing common ratio visualizations used in media [7, 93]. Process-focused designs instead emphasized the steps taken by journalists or AI, aligning with Wittenberg et al. [115]’s framework for process-based disclosures. While these approaches reveal complementary ways of framing human-AI collaboration for readers, a key challenge across designs was how to communicate “who did what” without overwhelming readers with detail, a point reminiscent of how consent banners (e.g., cookies for accepting General Data Protection Act) can exhibit dark patterns [42, 77], with users getting habituated to such consent interfaces [45]. As such, designers in our study explored varying levels of information depth, from simple overview visualizations to multi-layered designs,

drawing on principles of progressive disclosure [97]. Furthermore, this draws on Burrus et al. [15]’s approach, who argue that readers should be informed without overwhelming them with technical details. While we find that there is no one-size-fits-all solution to visualizing human–AI collaboration in news articles, we find that such visualizations must balance: (1) proportional representation of human and AI contributions, (2) process-based explanations of how those contributions were made, and (3) levels of information detail that can support readers with different engagement goals and AI literacy levels.

Taken together, we contribute an initial design space where we consider nuanced approaches in AI system disclosure design, with the objective of clearly communicating human–AI collaboration. This is to ensure readers to make more informed judgments about credibility and accountability. While our scope was limited to these designs, our work raises questions about optimal disclosure placement (cf., [69]), layering, and device as well platform-based disclosure adaptation, since for example disclosures that work on a desktop site may not easily translate to mobile or social media feeds. Furthermore, beyond textual content, these early visualizations provide a starting point toward adaptation to other modality-dependent approaches, for example provenance indicators for AI-edited images [106], timeline visualizations for video production [56], or even metadata encryption in audio files [105]. We believe our work serves as a key starting point for ensuring more nuanced disclosure design in this era of blended work [22], within journalism and beyond. This includes key domains such as science communication or within educational activities, where transparency about human–AI collaboration is equally important [44, 114, 118]. For example, similarly to how journalists would carefully disclose their usage in a given format, we believe the same process can be applied towards AI usage disclosure for academic manuscripts, medical practitioners, and possibly any domain where AI usage is (a) composed of tasks with differing influences of AI contributions (b) (eventually) governed by the need for disclosure (e.g., through policy or best practices). In a future where combinations of humans and AI (cf., [108]) will increasingly become a de facto standard, our work takes a leap toward envisioning a media ecosystem where audiences (whether news consumers or otherwise) are not only told that AI was involved, but can meaningfully understand and evaluate its role, in a usable manner.

5.2 How do different disclosure visualizations communicate human–AI collaboration to readers?

We evaluated four human–AI collaboration disclosure visualizations on user perception (RQ2) in a lab-based study combining eye tracking, questionnaires, and interviews. We found that participants distinguished between high and low AI contributions across all designs (Fig. 11), validating disclosures as effective tools to communicate collaboration dynamics. This echoes Zier & Diakopoulos [120], who state that showing both human and AI involvement is necessary to demystify assumptions about AI. Overall, all visualizations were rated positively, but each revealed distinct strengths suited for certain purposes or contexts. Surprisingly, the visualization format itself influenced how readers attributed human and

AI roles. For primarily AI-written articles, the Task-based Timeline made readers see greater human involvement for writing the entire article, while Role-based Timelines shifted primarily human-written articles towards the perception of more AI contribution. At specific role level, only fact-checking was perceived with more AI involvement due to the Role-based Timeline, in both context. These amplifications show that disclosure designs (likely in leading users to focus on different aspects of the information and visual design) do not just communicate information but can actively alter the balance of human and AI contribution understanding, depending on how roles and tasks are represented.

Additionally, AI literacy and ChatGPT experience influenced perception of human-AI collaboration and visual engagement. Individuals with higher AI literacy demonstrated longer gaze durations, and more fixations and saccades, but perceived less AI involvement to headline creation and topic selection. This indicates to a more critical evaluation. However, more experienced ChatGPT users spent shorter time, and had fewer and shorter fixations and saccades, but attributed more AI involvement to topic selection. These findings reveal that the communication of human-AI collaboration by a visualization also depends on the readers’ experience and literacy. Below we list findings related to each of our visualizations:

Textual Disclosure is a more collaborative focused resemblance of current platform labels, but was often described as too general or overlooked. Eye tracking confirmed this with shorter fixations suggesting scanning rather than deep processing. Participants described it as “too textual” or “like a legal disclaimer” (Section 4.5.2), which helps explain why it was least effective at conveying workflow steps or providing an overview. Additionally, it highlights the importance of Epstein et al.’s [31] observation that the language used in labels has an immense influence on user perceptions. Although a common format nowadays, our findings suggest that it is not the most effective in communicating human-AI collaboration.

Role-based Timeline provided a significantly better understanding of the steps, also without requiring any user interaction. Participants valued its static format for quick scanning, and eye tracking data indicated longer fixations and broader eye movements. Both timelines were perceived significantly clearer than the Chatbot and Textual Disclosure, and the Role-based Timeline was also seen as easier to understand than Textual Disclosure, which was also supported by qualitative findings (Section 4.5.2). However, icons sometimes suggested a more balanced collaboration than intended, showing the need to refine the design to avoid misleading readers. Furthermore, this visualization may not easily scale toward representing multiple journalists and other editorial roles, as roles executed in parallel would disrupt the linear sequence and can then result in visual clutter. As Altay & Gilardi [2] note, distinctions between human and AI contributions must be very clear so users do not overestimate or underestimate AI’s role.

Chatbot enabled in-depth exploration through progressive disclosure [97] which gave users more control. It led to the longest dwell times and highest fixation and saccade counts, indicating high engagement, but not necessarily more information processing per element. Quantitative findings showed it was the most informative, but also significantly harder to understand than all other visualizations, confirming interview comments that it could feel overwhelming (Section 4.5.2). While valued in high-stake contexts,

some found it offered too much to read, where its lack of immediate overview could additionally discourage casual news readers. This aligns with Burrus et al.'s [15] perspective of cautioning against overloading users with technical detail.

Task-based Timeline provided a good overview, insight into specific steps, and was most frequently preferred. Questionnaire and gaze data indicated effective information processing. However, the hover interaction was not always discovered, suggesting further refinement by replacing hover with inline text (Section 4.5.2), again aligning with Burrus et al. [15] to inform users without too much technical detail.

To summarize, we find that effectiveness depends on the intended use and context. Textual disclosure is familiar, and while increasingly pervasive across work practices [22], is the least effective in communicating collaboration. Timelines provide strong overviews and insights into the specific steps of the process, where we believe this is generally suitable for human-AI collaboration tasks with (currently) clear workflows (e.g., producing an academic research manuscript [29], or practicing design work [66]). For high-stake articles, and more generally safety-critical domains (e.g., medical AI systems [73]) that may require detailed information, the Chatbot would be most effective. The amplification effects raise the question: do disclosures simply inform readers, or do they also re-frame how AI's role is understood? Our findings show that design choices can tilt perceptions toward human or AI involvement, making disclosures active framings of contributions rather than neutral sources of information. This is in line with recent research that positions framing as a core dimension of design practice, actively shaping narratives, not just showing data [92]. To that end, if we are to maximize both transparency and usability, disclosure formats should be adapted to a news article's topic, desired information depth, and the intended level of reader engagement. However, we caution about blanket personalization of disclosure visualizations and instead recommend consistency, specifically due to any risks of end-user misinterpretation especially should disclosures be used within sensitive settings (e.g., patient-centered health care [80] and medical AI system design [73]).

5.3 Design considerations

Placement as a balance between visibility and bias. Participants were divided on whether disclosures should appear at the beginning of articles to encourage critical reading or at the end to avoid introducing negative bias (Sec. 4.5.4). Prior studies show that AI-labeled headlines are perceived as less accurate or less shareable [2, 39, 103], reflecting the transparency dilemma where disclosures intended to increase trust instead end up eroding it [89]. Whether our human-AI collaboration visualizations, which are more detailed, influence users similarly remains an open question. However, it does raise the consideration of how to design disclosures so they are visible enough to support critical engagement without introducing unintended bias (cf., progressive disclosures [97]), and whether other means such as toggles may be more suitable.

Disclosure formats and stake dependence. Although our user study stimuli consisted of only low-stakes articles, our qualitative findings support Gamage et al. [38]'s findings, where we found that participants considered stake contexts in news articles, which often

demanded different formats. For low-stakes articles, a simple textual label may suffice, even if less effective overall. While high-stakes or sensitive topics may require more detailed visualizations, such as timelines or a chatbot format (Section 4.5.4) to provide more nuance and accountability. News organizations could consider developing internal guidelines to match disclosure formats to article type and stake dependence, similar to editorial policies.

Supporting AI literacy through disclosures. Several participants reported becoming more critical readers or more aware of AI's capabilities after seeing the disclosures (Section 4.5.4). This suggests that disclosures may function as implicit AI literacy (cf., [110]) interventions. Repeated exposure could increase literacy dimensions such as understanding, detection, and ethical awareness, in line with the commonly used Meta AI Literacy Scale (MAILS) [17]. While not necessarily designed as educational tools, meaningful human-AI collaboration disclosure designs could have the added benefit of increasing AI literacy among audiences with limited GenAI knowledge.

Disclosures for AI-generated disclosures. Despite our detailed visualizations, some participants still questioned the truthfulness of the disclosures themselves 4.5.1. Given that LLMs can hallucinate or fabricate believable but incorrect information [57], this raises concerns for any AI-generated disclosures, most prominently Chatbot-based disclosures. In this regard, AI-generated disclosures raise a new problem: what if the disclosure itself is either subject to AI hallucination or deliberately populated with incorrect information? Does this necessitate an additional disclosure on the AI-generated disclosure? Future work should explore how to prevent disclosures from becoming sources of misinformation by indicating how, by whom, and with what information they were created, and by investigating how trust indicators can aid users in evaluating their accuracy and trustworthiness.

Encoding disclosures into provenance signals AI usage detection across modalities such as watermarking [63], structural marking [10], metadata standards [106], or cryptographic signatures [64] are key steps toward detecting AI-generated content, but typically do not include human-AI collaboration ratios. For example, prior work on metadata standards like C2PA have shown that provenance labels can improve trust [106], whereas others such as TransparentMeta [105] aim to encrypt transparency and attribution metadata directly into audio files. Our findings open the possibility of extending provenance frameworks to include metadata on how humans and AI collaborated across tasks and roles, and to encode these disclosures as machine-readable provenance signals.

5.4 Limitations and future work

First, it can be argued that since we are concerned with the news production cycle, we should have adopted a more participatory AI approach (cf., [9, 46] where we involved journalists directly, as this may limit insights into the feasibility of implementing such disclosures within editorial workflows. However, this was intentionally scoped as our study was primarily focused on establishing whether and how readers make sense of human-AI collaboration disclosure visualizations, before turning to questions of journalist adoption and/or their preferred methods of disclosure creation. Nevertheless,

we believe this provides a rich avenue for future work: from establishing how journalists would like to disclose AI usage, to testing disclosure visualizations on live news platforms, and across articles spanning low and high stakes, and with keeping the target demographic in mind (e.g., race or gender [19] or political orientation [4]). Another cautionary aspect: while it is known that LLM users do engage in passive non-disclosure or active concealment of the usage of AI [119], this may not necessarily generalize to journalism, where trust and accountability are crucial. Therefore, we contend that for true adoption, collaborations with journalists and news organizations will be essential to evaluate workflow integration and feasibility with respect to this emergent type of AI disclosures.

Third, our study was limited to evaluating only four prototype designs. While these were selected to create diverse visualizations and keeping the amount feasible for evaluation, other design ideas were generated which can be further explored. Additionally, all prototypes communicated the same core information about human and AI contributions, but differed in the level of detail provided. In particular, the Chatbot enabled users to retrieve more details through the layered predefined questions, incorporating progressive disclosure (cf., [97]). This difference in information granularity was intentional to explore how readers respond to varying disclosure formats that naturally present different information depths. This choice ensured diversity in disclosure design, despite introducing a potential confound. Moreover, our disclosure designs were restricted primarily to textual news content, where we do not examine multimodal disclosures (e.g., in audio, images, or video). For example, image-based disclosures such as C2PA visual provenance approaches were not explored, where embedding C2PA provenance labels in images has shown to improve trust in AI-generated content [106]. Fourth, there are clear limits to our ecological validity; despite that our study was conducted with a controlled web application, and aimed to mimic web news consumption, this does not fully reflect the context of real-world news platforms, which also include branding that may influence trust [68]. Relatedly, our study only explored short-term effects of AI disclosures, where it remains future work to better understand the long-term impact of repeated exposures to such disclosures, where participants may get desensitized and stop noticing or reading them after the initial short-term engagement (cf., [39]).

Finally, given that users have different disclosure visualization preferences (cf., [88, 111]), it is worth asking to what extent media organizations should personalize these disclosures to end-users visiting their websites. While newsrooms can assess user preferences and tailor such interfaces, with a manageable workflow should such metadata tagging become standardized, the real risk lies in differing interpretations communicated, where AI roles and contributions can be more/less amplified depending on the disclosure visualization. This may pose risks to news organizations, especially given that prior work has shown that AI-generated news is perceived as less credible [70], and differences in trust with some disclosure designs more than others [38, 88], even if limited to high stakes scenarios. Nevertheless, we speculate that while in the future users may have a minimal (standardized) disclosure visualization available for all, there could be optional expansion layers of tailored visualizations for readers that desire this. This would provide a

combination of media transparency, user control, and importantly, editorial consistency.

6 Conclusion

This work aimed to examine how human-AI collaboration in news production can be communicated and perceived through visual disclosures. Based on co-design sessions (N=10), we derive four disclosure visualizations and evaluate them in a controlled within-subjects lab study (N=32). We examined how these disclosure visualizations (Textual, Role-based Timeline, Task-based Timeline, Chatbot) and collaboration ratios (Primarily Human-written vs. Primarily AI) influenced visualization perceptions, gaze patterns, and post-experience responses. We found that (a) all prototypes effectively communicated human-AI collaboration output ratios, but for varied purposes; (b) textual disclosures were least effective in communicating human-AI collaboration; (c) Role-based and Task-based Timelines provided clearer overviews of editorial steps, while the Chatbot offered the most in-depth information; and (d) role-based timelines amplified AI contribution in primarily human-written articles, whereas task-based timeline shifted perceptions toward human involvement in primarily AI articles. Our work provides empirically validated human-AI collaboration disclosure prototypes⁸ and design considerations that can support researchers, practitioners, and news organizations in creating more nuanced disclosures. Through such disclosures, we envision a future where human and AI contributions can be better captured and communicated, especially for critical domains such as journalistic editorial processes.

AI Usage Disclosure Statement

To align with the disclosure practices we investigated, we conclude with a disclosure statement of our own collaboration with AI (see Fig 15). This illustrates how human-AI collaboration disclosure visualizations could also be adapted to research, with five research phases commonly used in HCI [29]. We emphasize that the intellectual framing, initial content created, critical perspectives, and final interpretations and crafting of the discussion, remain our own.

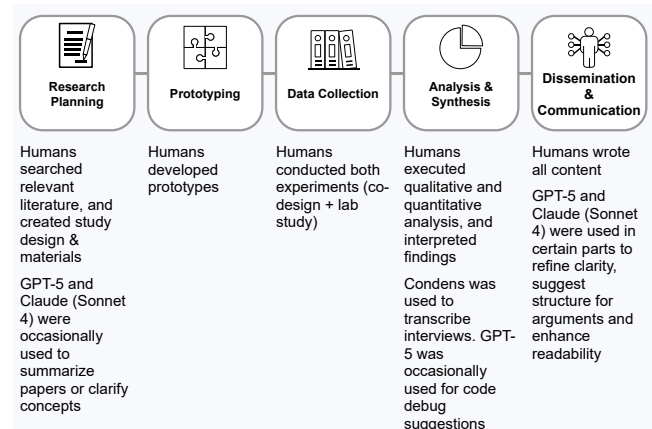


Figure 15: Task-based Timeline to disclose our AI usage.

⁸GitHub: <https://github.com/cwi-dis/CHI2026-AIDisclosureVisualizations>

References

- [1] Bill Albert and Tom Tullis. 2022. *Measuring the user experience: Collecting, analyzing, and presenting UX metrics*. Morgan Kaufmann.
- [2] Sacha Altay and Fabrizio Gilardi. 2024. People are skeptical of headlines labeled as AI-generated, even if true or human-made, because they assume full AI automation. *PNAS nexus* 3, 10 (2024), pgae403.
- [3] Agnieszka Andrychowicz-Trojanowska. 2018. Basic terminology of eye-tracking research. *Applied Linguistics Papers* 25/2 (2018), 123–132.
- [4] Theo Araujo, Anna Brosius, Andreas C Goldberg, Judith Möller, and Claes de Vreese. 2023. Humans vs. AI: the role of trust, political attitudes, and individual characteristics on perceptions about automated decision making across Europe. *International Journal of Communication* 17 (2023), 28.
- [5] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48. doi:10.18637/jss.v067.i01
- [6] Elisa Bertino, Finale Doshi-Velez, Maria Gini, Daniel Lopresti, and David Parkes. 2020. Artificial Intelligence & Cooperation. *arXiv e-prints*, Article arXiv:2012.06034 (Dec. 2020), arXiv:2012.06034 pages. arXiv:2012.06034 [cs.CY] doi:10.48550/arXiv.2012.06034
- [7] Francesco Bianconi. 2024. *Proportions*. Springer Nature Switzerland, Cham, 25–52. doi:10.1007/978-3-031-57051-3_3
- [8] Monika Bickert. 2024. Our Approach to Labeling AI-Generated Content and Manipulated Media. <https://about.fb.com/news/2024/04/metasp-approach-to-labeling-ai-generated-content-and-manipulated-media/> Accessed: 2025-01-27.
- [9] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Jason Gabriel, and Shakir Mohamed. 2022. Power to the People? Opportunities and Challenges for Participatory AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization* (Arlington, VA, USA) (EAAMO '22). Association for Computing Machinery, New York, NY, USA, Article 6, 8 pages. doi:10.1145/3551624.3555290
- [10] Adam Block, Ayush Sekhari, and Alexander Rakhlin. 2025. GaussMark: A Practical Approach for Structural Watermarking of Language Models. arXiv:2501.13941 [cs.CR] <https://arxiv.org/abs/2501.13941>
- [11] Michał Boni. 2021. The ethical dimension of human-artificial intelligence collaboration. *European View* 20, 2 (2021), 182–190.
- [12] Virginia Braun and Victoria Clarke. 2012. Thematic analysis. In *APA Handbook of Research Methods in Psychology, Vol. 2: Research Designs: Quantitative, Qualitative, Neuropsychological, and Biological*, Harris Cooper, Paul M. Camic, Deborah L. Long, A. T. Panter, David Rindskopf, and Kenneth J. Sher (Eds.). American Psychological Association, 57–71. doi:10.1037/13620-004
- [13] Karin Breckner, Thomas Neumayr, Martina Mara, Marc Streit, and Mirjam Augstein. 2025. The Changing Nature of Human-AI Relations: A Scoping Review on Terminology and Evolution in the Scientific Literature. *International Journal of Human-Computer Interaction* 0, 0 (2025), 1–58. arXiv:https://doi.org/10.1080/10447318.2025.2482742 doi:10.1080/10447318.2025.2482742
- [14] Peter Bro, Kenneth Reinecke Hansen, and Ralf Andersson. 2016. Improving productivity in the newsroom? Deskillling, reskilling and multiskilling in the news media. *Journalism Practice* 10, 8 (2016), 1005–1018.
- [15] Olivia Burrus, Amanda Curtis, and Laura Herman. 2024. Unmasking AI: Informing Authenticity Decisions by Labeling AI-Generated Content. *Interactions* 31, 4 (2024), 38–42.
- [16] C2PA. 2024. Introducing Official Content Credentials Icon - C2PA — c2pa.org. <https://c2pa.org/post/contentcredentials/>. [Accessed 17-01-2024].
- [17] Astrid Carolus, Martin J Koch, Samantha Straka, Marc Erich Latoschik, and Carolin Wiernich. 2023. MAIIS-Meta AI literacy scale: Development and testing of an AI literacy questionnaire based on well-founded competency models and psychological change-and meta-competencies. *Computers in Human Behavior: Artificial Humans* 1, 2 (2023), 100014.
- [18] Chaomei Chen. 2010. Information visualization. *Wiley Interdisciplinary Reviews: Computational Statistics* 2, 4 (2010), 387–403.
- [19] Inyoung Cheong, Alicia Guo, Mina Lee, Zhehui Liao, Kowe Kadoma, Dongyoung Go, Joseph Chee Chang, Peter Henderson, Mor Naaman, and Amy X. Zhang. 2025. Penalizing Transparency? How AI Disclosure and Author Demographics Shape Human and AI Judgments About Writing. arXiv:2507.01418 [cs.CY] <https://arxiv.org/abs/2507.01418>
- [20] Rune Haubo B Christensen. 2018. Cumulative link models for ordinal regression with the R package ordinal. *Submitted in J. Stat. Software* 35 (2018), 1–46.
- [21] Jacob Cohen. 2013. *Statistical power analysis for the behavioral sciences*. routledge.
- [22] Marios Constantinides, Himanshu Verma, Shadan Sadeghian, and Abdallah El Ali. 2025. The Future of Work is Blended, Not Hybrid. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25)*. Association for Computing Machinery, New York, NY, USA, Article 28, 13 pages. doi:10.1145/3729176.3729202
- [23] Hannes Cools and Nicholas Diakopoulos. 2024. Uses of Generative AI in the Newsroom: Mapping Journalists' Perceptions of Perils and Possibilities. *Journalism Practice* 0, 0 (2024), 1–19. arXiv:https://doi.org/10.1080/17512786.2024.2394558 doi:10.1080/17512786.2024.2394558
- [24] Patrick Corrigan. 2024. LinkedIn Adopts C2PA Standard. <https://www.linkedin.com/pulse/linkedin-adopts-c2pa-standard-patrick-corrigan-kwldf/?trackingId=4gnjmapwRsmugUNqwj0fRw%3D%3D> Accessed: 2025-01-27.
- [25] Weiwei Cui, Xiaoyu Zhang, Yun Wang, He Huang, Bei Chen, Lei Fang, Haidong Zhang, Jian-Guan Lou, and Dongmei Zhang. 2019. Text-to-viz: Automatic generation of infographics from proportion-related natural language statements. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 906–916.
- [26] Google DeepMind. 2024. SynthID. <https://deepmind.google/technologies/synthid/>. Accessed: 2024-1-19.
- [27] Nicholas Diakopoulos, Hannes Cools, Charlotte Li, Natali Helberger, Ernest Kung, Aimee Rinehart, and L Gibbs. 2024. Generative AI in journalism: the evolution of newswork and ethics in a generative information ecosystem.
- [28] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. 2024. Transparent AI Disclosure Obligations: Who, What, When, Where, Why, How. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 342, 11 pages. doi:10.1145/3613905.3650750
- [29] Passant Elagroudy, Jie Li, Kaisa Väänänen, Paul Lukowicz, Hiroshi Ishii, Wendy E. Mackay, Elizabeth F Churchill, Anicia Peters, Antti Oulasvirta, Rui Prada, Alexandra Diening, Giulia Barbareschi, Agnes Gruenerbl, Midori Kawaguchi, Abdallah El Ali, Fiona Draxler, Robin Welsch, and Albrecht Schmidt. 2024. Transforming HCI Research Cycles using Generative AI and “Large Whatever Models” (LWMs). In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 584, 5 pages. doi:10.1145/3613905.3643977
- [30] Louis Engelbrecht, Adele Botha, and Ronell Alberts. 2015. Designing the visualization of information. *International Journal of Image and Graphics* 15, 02 (2015), 1540005.
- [31] Ziv Epstein, Mengying C Fang, Antonio A Arechar, and David G Rand. 2023. What label should be applied to content produced by generative AI? doi:10.31234/osf.io/v4mfz
- [32] Ziv Epstein, Aaron Hertzmann, Investigators of Human Creativity, Memo Akten, Hany Farid, Jessica Fjeld, Morgan R Frank, Matthew Groh, Laura Herman, Neil Leach, et al. 2023. Art and the science of generative AI. *Science* 380, 6650 (2023), 1110–1111.
- [33] European Parliament and Council. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence and amending certain Union legislative acts. <https://artificialintelligenceact.eu/article/50/> Official Journal of the European Union.
- [34] Richard Fletcher and R Nielsen. 2024. What does the public in six countries think of generative AI in news? URL https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2024-05/Fletcher_and_Nielsen_Generative_AI_and_News_Audiences.pdf (2024), 42 pages.
- [35] Tania Forja-Pena, Berta García-Orosa, and Xosé López-García. 2024. The Ethical Revolution: Challenges and Reflections in the Face of the Integration of Artificial Intelligence in Digital Journalism. *Communication & Society* 37, 3 (Jun. 2024), 237–254. doi:10.15581/003.37.3.237-254
- [36] Paul Formosa, Sarah Bankins, Rita Matulionyte, and Omid Ghasemi. 2025. Can ChatGPT be an author?: Generative AI creative writing assistance and perceptions of authorship, creatorship, responsibility, and disclosure. *AI and Society* 40, 5 (June 2025), 3405–3417. doi:10.1007/s00146-024-02081-0
- [37] George Fragiadakis, Christos Diou, George Kousiouris, and Mara Nikolaidou. 2025. Evaluating Human-AI Collaboration: A Review and Methodological Framework. arXiv:2407.19098 [cs.HC] <https://arxiv.org/abs/2407.19098>
- [38] Dilrukshi Gamage, Dilki Sewwandi, Min Zhang, and Arosha K Bandara. 2025. Labeling Synthetic Content: User Perceptions of Label Designs for AI-Generated Content on Social Media. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 814, 29 pages. doi:10.1145/3706598.3713171
- [39] Fabrizio Gilardi, Sabrina Di Lorenzo, Juri Ezzaini, Beryl Santa, Benjamin Streiff, Eric Zurluh, and Emma Hoes. 2024. Willingness to Read AI-Generated News Is Not Driven by Their Perceived Quality. *arXiv e-prints*, Article arXiv:2409.03500 (Sept. 2024), arXiv:2409.03500 pages. arXiv:2409.03500 [cs.CY] doi:10.48550/arXiv.2409.03500
- [40] Gregory Gondwe. 2023. Exploring the multifaceted nature of generative AI in journalism studies: A typology of scholarly definitions. *Available at SSRN 4465446* (2023).
- [41] Google. [n.d.]. Disclosing use of altered or synthetic content. <https://support.google.com/youtube/answer/1432849?hl=en&co=GENIE.Platform%3DAndroid> Accessed: 2025-01-27.

- [42] Colin M. Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. 2021. Dark Patterns and the Legal Requirements of Consent Banners: An Interaction Criticism Perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 172, 18 pages. doi:10.1145/3411764.3445779
- [43] Matthew Groh, Aruna Sankaranarayanan, Nikhil Singh, Dong Young Kim, Andrew Lippman, and Rosalind Picard. 2023. Human Detection of Political Speech Deepfakes across Transcripts, Audio, and Video. arXiv:2202.12883 [cs.HC]
- [44] Lars Guenther, Jessica Kunert, and Bernhard Goodwin. 2025. My New Colleague, ChatGPT? How German Science Journalists Perceive and Use (Generative) Artificial Intelligence. *Journalism Practice* 0, 0 (2025), 1–18. arXiv:https://doi.org/10.1080/17512786.2025.2502794 doi:10.1080/17512786.2025.2502794
- [45] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. “Okay, whatever”: An Evaluation of Cookie Consent Interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 621, 27 pages. doi:10.1145/3491102.3501985
- [46] Nicolai Brodersen Hansen, Christian Dindler, Kim Halskov, Ole Sejer Iversen, Claus Bossen, Ditte Amund Basballe, and Ben Schouten. 2020. How Participatory Design Works: Mechanisms and Effects. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction* (Fremantle, WA, Australia) (OzCHI '19). Association for Computing Machinery, New York, NY, USA, 30–41. doi:10.1145/3369457.3369460
- [47] Winston Haynes. 2013. Benjamini–hochberg method. In *Encyclopedia of systems biology*. Springer, 78–78.
- [48] Jessica He, Stephanie Houde, and Justin D. Weisz. 2025. Which Contributions Deserve Credit? Perceptions of Attribution in Human-AI Co-Creation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 540, 18 pages. doi:10.1145/3706598.3713522
- [49] Roy S Hessels, Chantal Kemner, Carljin van den Boomen, and Ignace TC Hooge. 2016. The area-of-interest problem in eyetracking research: A noise-robust solution for face and sparse stimuli. *Behavior research methods* 48, 4 (2016), 1694–1712.
- [50] Kenneth Holmqvist, Marcus Nyström, and Fiona Mulvey. 2012. Eye tracker data quality: what it is and how to measure it. In *Proceedings of the Symposium on Eye Tracking Research and Applications* (Santa Barbara, California) (ETRA '12). Association for Computing Machinery, New York, NY, USA, 45–52. doi:10.1145/2168556.2168563
- [51] Steffen Holter and Mennatallah El-Assady. 2024. Deconstructing Human-AI Collaboration: Agency, Interaction, and Adaptation. *Computer Graphics Forum* 43, 3 (2024), e15107. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.15107 doi:10.1111/cgf.15107
- [52] Ignace TC Hooge, Gijs A Holleman, Nina C Haukes, and Roy S Hessels. 2019. Gaze tracking accuracy in humans: One eye is sometimes better than two. *Behavior Research Methods* 51, 6 (2019), 2712–2721.
- [53] Runsheng Huang, Liam Dugan, Yue Yang, and Chris Callison-Burch. 2024. MiRAGeNews: Multimodal Realistic AI-Generated News Detection. arXiv:2410.09045 [cs.CV] https://arxiv.org/abs/2410.09045
- [54] Christophe Hurter, Alexandru Telea, and Bernice Rogowitz. 2024. *Information Visualization*. CRC Press, 223–262. doi:10.1201/9781003495147-7 Publisher Copyright: © 2025 selection and editorial matter, Constantine Stephanidis and Gavriel Salvendy..
- [55] Angel Hsing-Chi Hwang, Q. Vera Liao, Su Lin Blodgett, Alexandra Olteanu, and Adam Trischler. 2025. “It was 80% me, 20% AI”: Seeking Authenticity in Co-Writing with Large Language Models. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW122 (May 2025), 41 pages. doi:10.1145/3711020
- [56] Ayae Ide, Tory Park, Jaron Mink, and Tanusree Sharma. 2025. Signals of Provenance: Practices & Challenges of Navigating Indicators in AI-Generated Media for Sighted and Blind Individuals. arXiv:2505.16057 [cs.HC] https://arxiv.org/abs/2505.16057
- [57] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM computing surveys* 55, 12 (2023), 1–38.
- [58] Haiyan Jia, Alyssa Appelman, Mu Wu, and Steve Bien-Aimé. 2024. News bylines and perceived AI authorship: Effects on source and message credibility. *Computers in Human Behavior: Artificial Humans* 2, 2 (2024), 100093. doi:10.1016/j.chbah.2024.100093
- [59] Tingting Jiang, Zhumo Sun, Shiting Fu, and Yan Lv. 2024. Human-AI interaction research agenda: A user-centered perspective. *Data and Information Management* 8, 4 (2024), 100078. doi:10.1016/j.dim.2024.100078
- [60] Cameron R. Jones and Benjamin K. Bergen. 2025. Large Language Models Pass the Turing Test. arXiv:2503.23674 [cs.CL] https://arxiv.org/abs/2503.23674
- [61] Mitt Nowshade Kabir. 2024. Unleashing Human Potential: A Framework for Augmenting Co-Creation with Generative AI. In *Proceedings of the International Conference on AI Research*. Academic Conferences and publishing limited.
- [62] Maurits Kaptein. 2016. *Using Generalized Linear (Mixed) Models in HCI*. Springer International Publishing, Cham, 251–274. doi:10.1007/978-3-319-26633-6_11
- [63] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. A watermark for large language models. In *International Conference on Machine Learning*. PMLR, 17061–17084.
- [64] Pallav Laskar. 2025. Cryptographic Provenance and the Future of Media Authenticity: Technical Standards and Ethical Frameworks for Generative Content. *Journal of Computer Science and Technology Studies* 7, 6 (Jun. 2025), 967–972. doi:10.32996/jcsts.2025.7.114
- [65] Fan Li, Ya Yang, et al. 2024. Impact of Artificial Intelligence–Generated Content Labels on Perceived Accuracy, Message Credibility, and Sharing Intentions for Misinformation: Web-Based, Randomized, Controlled Experiment. *JMIR Formative Research* 8, 1 (2024), e60024.
- [66] Jie Li, Hancheng Cao, Laura Lin, Youyang Hou, Ruihao Zhu, and Abdallah El Ali. 2024. User Experience Design Professionals’ Perceptions of Generative Artificial Intelligence. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 381, 18 pages. doi:10.1145/3613904.3642114
- [67] Jian Li, Jinsong Huang, Jiayang Liu, and Tianqi Zheng. 2022. Human-AI co-operation: modes and their effects on attitudes. *Telematics and Informatics* 73 (2022), 101862.
- [68] Bingjie Liu and Lewen Wei. 2019. Machine authorship in situ: Effect of news organization and news genre on news credibility. *Digital journalism* 7, 5 (2019), 635–657.
- [69] Jacob A Long, Tabitha Oyewole, Maryam Goli, Jacqueline M Keisler, Saud Alyaqout, Michael D Rodgers, and Arielle N'Diaye. 2019. The Disclosure Dilemma: How AI Attribution Affects Reactions to Public Health Messages. (In.d.).
- [70] Chiara Longoni, Andrey Fradkin, Luca Cian, and Gordon Pennycook. 2022. News from Generative Artificial Intelligence Is Believed Less. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 97–106. doi:10.1145/3531146.3533077
- [71] Tambiama Madiega. 2023. Generative AI and watermarking. *European Parliament* (2023). https://www.europarl.europa.eu/RegData/etudes/BRIE/2023/757583/EPRS_BRI(2023)757583_EN.pdf
- [72] Andrew Vande Moere and Helen Purchase. 2011. On the role of design in information visualization. *Information Visualization* 10, 4 (2011), 356–371.
- [73] Pedro A. Moreno-Sánchez, Javier Del Ser, Mark van Gils, and Jussi Hernesniemi. 2026. A design framework for operationalizing trustworthy artificial intelligence in healthcare: Requirements, tradeoffs and challenges for its clinical adoption. *Information Fusion* 127 (2026), 103812. doi:10.1016/j.inffus.2025.103812
- [74] Sophie Morosoli, Laurens Naudts, Hannes Cools, Karthikeya Venkatraj, Natali Helberger, and Claes de Vreese. 2025. Public accountability and regulatory expectations for AI in journalism: qualitative evidence from focus groups with Dutch citizens. *AI & SOCIETY* (2025), 1–13.
- [75] Jakob Nielsen. 2005. *Ten usability heuristics*. http://www.nngroup.com/articles/ten-usability-heuristics/. 01163.
- [76] Sachita Nishal and Nicholas Diakopoulos. 2024. Envisioning the Applications and Implications of Generative AI for News Media. arXiv:2402.18835 [cs.CY] https://arxiv.org/abs/2402.18835
- [77] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
- [78] Johannes Oberpriller, Melina de Souza Leite, and Maximilian Pichler. 2022. Fixed or random? On the reliability of mixed-effects models for a small number of levels in grouping variables. *Ecology and Evolution* 12, 7 (2022), e9062. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/ece3.9062 doi:10.1002/ece3.9062 e9062 ECE-2022-01-00007.R1.
- [79] Anneli Olsen. 2012. The Tobii I-VT fixation filter. *Tobii Technology* 21, 4-19 (2012), 5.
- [80] Roy Perlis. 2025. AI Disclosure and Patient Consent in Health Care. *JAMA* 334, 11 (09 2025), 961–961. arXiv:https://jamanetwork.com/journals/jama/articlepdf/2837952/-jama_perlis_2025_en_250008_1757361763.57499.pdf doi:10.1001/jama.2025.14026
- [81] Stanislaw Piasecki, Sophie Morosoli, Natali Helberger, and Laurens Naudts. 2024. AI-generated journalism: Do the transparency provisions in the AI Act give news readers what they hope for? *Internet Policy Review* 13, 4 (2024), 1–28. doi:10.14763/2024.4.1810
- [82] Colin Porlezza and Aljosha Karim Schapals. 2024. AI ethics in journalism (studies): An evolving field between research and practice. *Emerging Media* 2, 3 (2024), 356–370.

- [83] Annemarie Quispel, Alfons Maes, and Joost Schilperoord. 2018. Aesthetics and Clarity in Information Visualization: The Designer's Perspective. *Arts* 7, 4 (2018), 72. doi:10.3390/arts7040072
- [84] Michael J. Rees, Andrew White, and Bebo White. 2001. *Designing Web Interfaces* (1st ed.). Prentice Hall Professional Technical Reference.
- [85] Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. 2025. Can AI-Generated Text be Reliably Detected? arXiv:2303.11156 [cs.CL]. <https://arxiv.org/abs/2303.11156>
- [86] Gonesh Chandra Saha, Sanjay Kumar, Avinash Kumar, Hasi Saha, TK Lakshmi, and Niyati Bhat. 2023. Human-AI Collaboration: Exploring interfaces for interactive Machine Learning. *Tuijin Jishu/Journal of Propulsion Technology* 44, 2 (2023), 2023.
- [87] Elizabeth B-N Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *Co-design* 4, 1 (2008), 5–18.
- [88] Nicolas Scharowski, Michaela Benk, Swen J. Kühne, Léane Wettstein, and Florian Brühlmann. 2023. Certification Labels for Trustworthy AI: Insights From an Empirical Mixed-Method Study. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 248–260. doi:10.1145/3593013.3593994
- [89] Oliver Schilke and Martin Reimann. 2025. The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes* 188 (2025), 104405. doi:10.1016/j.obhdp.2025.104405
- [90] Tommy Shane, Emily Saltz, and Claire Leibowicz. 2021. From deepfakes to TikTok filters: How do you label AI content? *Nieman Lab*, May 12 (2021).
- [91] Yi Shi and Lin Sun. 2024. How Generative AI Is Transforming Journalism: Development, Application and Ethics. *Journalism and Media* 5, 2 (2024), 582–594.
- [92] Prakash Shukla and Paul Parsons. 2025. Identifying Framing Practices in Visualization Design Through Practitioner Reflections. *arXiv e-prints*, Article arXiv:2508.20383 (Aug. 2025), arXiv:2508.20383 pages. arXiv:2508.20383 [cs.HC] doi:10.48550/arXiv.2508.20383
- [93] Harri Siirtola. 2014. Bars, Pies, Doughnuts & Tables: Visualization of Proportions. In *Proceedings of the 28th International BCS Human Computer Interaction Conference on HCI 2014 - Sand, Sea and Sky - Holiday HCI* (Southport, UK) (BCS-HCI '14). BCS Learning & Development Ltd, Swindon, GBR, 240–245. doi:10.14236/ewic/hci2014.30
- [94] Aswath Sivakumaran. 2023. Investigating Consumer Perception and Speculative AI Labels for Creative AI Usage In Media.
- [95] Alem Febri Sonni, Hasdiyanto Hafied, Irwanto Irwanto, and Rido Latuheru. 2024. Digital Newsroom Transformation: A Systematic Review of the Impact of Artificial Intelligence on Journalistic Practices, News Narratives, and Ethical Challenges. *Journalism and Media* 5, 4 (2024), 1554–1570.
- [96] Ian Spence and Stephan Lewandowsky. 1991. Displaying proportions and percentages. *Applied Cognitive Psychology* 5, 1 (1991), 61–77.
- [97] Aaron Springer and Steve Whittaker. 2020. Progressive disclosure: When, why, and how do users want algorithmic transparency information? *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10, 4 (2020), 1–32.
- [98] Marc Steen. 2013. Co-design as a process of joint inquiry and imagination. *Design issues* 29, 2 (2013), 16–28.
- [99] Mohsen Tavakol and Reg Dennick. 2011. Making sense of Cronbach's alpha. *International journal of medical education* 2 (2011), 53.
- [100] Sina Thäsler-Kordonouri. 2024. What Comes After the Algorithm? An Investigation of Journalists' Post-editing of Automated News Text. *Journalism Practice* 0, 0 (2024), 1–20. arXiv:<https://doi.org/10.1080/17512786.2024.2404692> doi:10.1080/17512786.2024.2404692
- [101] TikTok. [n. d.]. About AI-generated content. <https://support.tiktok.com/en/using-tiktok/creating-videos/ai-generated-content> Accessed: 2025-01-27.
- [102] Tobii Technology. 2011. *Accuracy and precision test method for remote eye trackers*. Technical Report. Tobii Technology.
- [103] Benjamin Toff and Felix M. Simon. 2024. “Or They Could Just Not Use It?": The Dilemma of AI Disclosure for Audience Trust in News. *The International Journal of Press/Politics* 0, 0 (2024), 19401612241308697. arXiv:<https://doi.org/10.1177/19401612241308697> doi:10.1177/19401612241308697
- [104] Richie Torres. 2023. H.R.3831 - AI Disclosure Act of 2023. <https://www.congress.gov/bill/118th-congress/house-bill/3831?s=1&r=1> Accessed: 2024-1-15.
- [105] Transparent Audio. 2025. TransparentMeta: A Standard for Compliance with AI Audio Transparency Legislation. <https://www.transparentaudio.ai/transparentmeta>. An open-source Python library that encrypts transparency & attribution metadata into WAV & MP3 files, cryptographically signed for AI audio provenance.
- [106] Christoph Trattner, Svenja Forstner, Alain Starke, and Erik Knudsen. 2025. C2PA Provenance Labels Increase Trust in News Platforms Across Western Countries. (05 2025). doi:10.31219/osf.io/pdhaz_v1
- [107] Jakob Trischler, Timo Dietrich, and Sharyn Rundle-Thiele. 2019. Co-design: from expert-to user-driven ideas in public service design. *Public management review* 21, 11 (2019), 1595–1619.
- [108] Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour* 8, 12 (2024), 2293–2303. doi:10.1038/s41562-024-02024-1
- [109] Raul Valverde. 2011. *Principles of Human Computer Interaction Design: HCI Design*. LAP Lambert Academic Publishing.
- [110] Johanna Velander, Nuno Otero, and Marcelo Milrad. 2024. *What is Critical (about) AI Literacy? Exploring Conceptualizations Present in AI Literacy Discourse*. Springer Nature Switzerland, Cham, 139–160. doi:10.1007/978-3-031-58622-4_8
- [111] Karthikeya Puttur Venkatraj, Sophie Morosoli, Hannes Cools, Laurens Naudts, Claes de Vreese, Natali Helberger, Pablo Cesar, and Abdallah El Ali. 2025. Understanding AI Disclosure Needs for News Production and Journalism. In *Proceedings of the 24th International Conference on Mobile and Ubiquitous Multimedia (MUM '25)*. Association for Computing Machinery, New York, NY, USA, 202–208. doi:10.1145/3771882.3771899
- [112] Jan Philip Wahle, Terry Ruas, Saif M. Mohammad, Norman Meuschke, and Bela Gipp. 2023. AI Usage Cards: Responsibly Reporting AI-Generated Content. In *2023 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. IEEE, 282–284. doi:10.1109/JCDL57899.2023.00060
- [113] Dakuo Wang, Elizabeth Churchill, Pattie Maes, Xiangmin Fan, Ben Shneiderman, Yuanchun Shi, and Qianying Wang. 2020. From Human-Human Collaboration to Human-AI Collaboration: Designing AI Systems That Can Work Together with People. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHIEA '20). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3334480.3381069
- [114] Peter Weingart and Lars Guenther. 2016. Science communication and the issue of trust. *Journal of Science communication* 15, 5 (2016), C01.
- [115] Chloe Wittenberg, Ziv Epstein, Adam J. Berinsky, and David G. Rand. 2024. Labeling AI-Generated Content: Promises, Perils, and Future Directions. *An MIT Exploration of Generative AI* (3 2024).
- [116] Chloe Wittenberg, Ziv Epstein, Gabrielle Péloquin-Skulski, Adam J Berinsky, and David G Rand. 2024. Labeling AI-Generated Media Online.
- [117] Yunfeng Zhang, Q. Vera Liao, and Rachel K. E. Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 295–305. doi:10.1145/3351095.3372852
- [118] Yue Zhang and Pascal Reusch. 2025. Trust in and Adoption of Generative AI in University Education: Opportunities, Challenges, and Implications. In *2025 IEEE Global Engineering Education Conference (EDUCON)*. 1–10. doi:10.1109/EDUCON62633.2025.11016490
- [119] Zhiping Zhang, Chenxinran Shen, Bingsheng Yao, Dakuo Wang, and Tianshi Li. 2025. Secret Use of Large Language Model (LLM). *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW163 (May 2025), 26 pages. doi:10.1145/3711061
- [120] Jessica Zier and Nicholas Diakopoulos. 2024. Labeling AI-Generated News Content: Matching Journalist Intentions with Audience Expectations. *informal Proceedings of the Computation + Journalism Symposium* (2024).