The HIPUD Model: A Framework for Characterizing Generative AI Chatbots as Information Environments

Anna Sophie Kümpel & Michael V. Reiss

Department of Media and Communication, LMU Munich

Author Note

Anna Sophie Kümpel https://orcid.org/0000-0001-7184-4057

Michael V. Reiss https://orcid.org/0000-0002-6094-9985

Correspondence concerning this article should be addressed to Anna Sophie Kümpel, Department of Media and Communication, LMU Munich, Oettingenstr. 67, 80537 Munich, Germany. Email: anna.kuempel@ifkw.lmu.de

Abstract

Gen-AI chatbots such as ChatGPT or Claude are becoming central tools for everyday information use. We argue that this shift is epistemic, not merely technological. Rather than simply mediating access to existing content, chatbots synthesize it into novel informational artifacts, changing how people encounter, process, and use information. To capture this shift, we propose the HIPUD model, which characterizes Gen-AI chatbots as information environments that are **H**yper-personalized, **I**nteractive, **P**robabilistic, **U**niversal, and **D**e-contextualized. The model organizes these characteristics in three layers: features of the system (probabilistic, universal), default features of the output (hyper-personalized, de-contextualized), and interactivity as the affordance through which users can partly adjust these defaults. We discuss broader transformations resulting from Gen-AI chatbot use and illustrate the model's analytical leverage with examples from political, health, and consumer communication research. Overall, the model equips future research with a shared vocabulary and analytical lens for understanding information use in Gen-AI environments.

Keywords: generative AI, chatbots, information use, information environments, defaults, epistemic authority, HIPUD model

The HIPUD Model: A Framework for Characterizing Generative AI Chatbots as Information Environments

Until recently, seeking unfamiliar information almost invariably meant turning to a search engine. This is no longer a given. Today, hundreds of millions of users consult generative AI (Gen-AI) chatbots to address their information needs (Chatterji et al., 2025). Within a few years of their release in late 2022, Gen-AI chatbots such as ChatGPT, Gemini, or Claude have become embedded in everyday information use—consulted for medical, legal, and political questions as well as for cooking, travel, and countless other tasks. This role shows every sign of deepening further (Simon et al., 2026), with far-reaching consequences for how information is encountered, processed, and used.

This shift builds on, yet departs from, a longer history in which new media have reconfigured the relationship between users and available information (Aspray & Hayes, 2011; Briggs et al., 2020). Libraries and print organized information through physical availability and human intermediaries such as librarians, who helped people translate vague needs into productive inquiries. Radio and television extended reach and immediacy but left users with little control over sequence, pace, or depth. The internet and search engines reversed this balance, granting substantial user control over selection and navigation—a mode Bates (1989) influentially described as “berrypicking.” Social media platforms then blended active and passive use, as algorithmically curated feeds delivered content users had not explicitly sought (Kümpel, 2022; Nanz & Matthes, 2022). Across these shifts, however, one feature remained stable: Each environment mediated access to a pre-existing stock of shared, (mostly) public information—content authored by identifiable persons or institutions and, at least in principle, accessible to others.

We argue that Gen-AI chatbots mark a qualitative break with this logic. They, too, draw on existing content, whether encoded in their training data or retrieved at the moment of a query. Yet rather than mediating access to this content, they turn it into novel informational artifacts on demand: tailored to the individual user rather than addressed to a public, produced through dialogic exchange rather than delivered as finished content, statistically synthesized rather than retrieved, spanning virtually any topic and modality rather than confined to a specific domain, and disembedded from the source and social cues that have traditionally accompanied mediated communication. This rupture is not merely technological but epistemic, reshaping how people use information in everyday life.

This article proposes the HIPUD model to describe this shift and enable its systematic investigation. The model characterizes Gen-AI chatbots along five characteristics: They are **H**yper-personalized, **I**nteractive, **P**robabilistic, **U**niversal, and **D**e-contextualized. Rather than presenting these as a list of parallel properties, we organize them in three layers—features of the system, default features of its output, and interactivity as the affordance through which users can partly adjust these defaults. Because Gen-AI chatbots are conversational, some of their characteristics can be renegotiated through dialogue, though, as we argue, to a far lesser extent than it may seem.

The framework takes a user-centered, everyday-life perspective, keeping the focus on how people perceive and respond to AI-generated information. From this perspective, it makes three contributions. First, it offers an explicit organizing logic and conceptual precision: It moves the debate beyond vague assertions that "AI is different" and specifies which characteristics matter for which kinds of information use. Second, it is applicable across subfields of communication research: Rather than replacing established theories, it specifies the conditions

under which these theories apply to Gen-AI chatbots and where their assumptions require adjustment. Third, by showing which characteristics users can influence and which they cannot, it raises awareness of how Gen-AI chatbots function and thereby also informs AI literacy and governance.

In the remainder of the article, we first define our core terms. We then derive and organize the five characteristics, discuss each in turn, and examine how they jointly shape users' information use. Finally, we apply the model to three purposefully diverse scenarios from political, health, and consumer communication to illustrate its utility across communicative settings.

Gen-AI Chatbots, Information Use, and Information Environments

Before explicating the HIPUD model, three key concepts require clarification: Gen-AI chatbots, information (use), and information environments.

We define *Gen-AI chatbots* as conversational agents powered by large language models that generate novel, contextually responsive content in response to user input, drawing on information encoded during training and, in many systems, on content retrieved at the time of the query. Following Gil de Zúñiga et al. (2024), they exemplify artificial intelligence as a capability of non-human entities to perform tasks, communicate, and act in ways that resemble human action (p. 318). Prototypical examples include standalone conversational systems such as OpenAI's ChatGPT, Anthropic's Claude, or Google's Gemini. The boundaries of this category are fluid, however: Generative search features such as Google's AI Overviews blend retrieval with generative synthesis, AI assistants are increasingly embedded in other applications, and emerging agentic systems extend conversational interaction into autonomous task execution. Our

main focus is on standalone Gen-AI chatbots that users deliberately turn to, and we treat embedded variants as boundary cases.

Information, in the context of this framework, is understood broadly. Building on Hasebrink and Domeyer (2010; see also Kümpel et al., 2022), we adopt a holistic, social conception of information as content that is subjectively new and/or useful to the individual. This encompasses different information needs—for example, topic-related needs arising from personal interests, or problem-related needs focused on acquiring specific knowledge (e.g., finding a quick lunch recipe, fixing a coding error). In foregrounding such everyday information practices, we draw on Savolainen's (1995) concept of Everyday Life Information Seeking (ELIS), which attends to how people seek and use information to orient themselves in daily life—managing households, maintaining health, or navigating social relationships. Our framework, however, extends beyond information seeking in the narrow sense. While seeking refers to the active pursuit of information to address a perceived gap, *information use,* as understood here, encompasses not only finding information but also processing, evaluating, applying, and integrating it into one's existing knowledge and decision-making (Kari, 2010).

Finally, we conceptualize Gen-AI chatbots as *information environments*, that is, as the 'arena' in which information is encountered, processed, and used (also see Kümpel, 2022). In line with the "mix of attributes" approach (Eveland, 2003), we describe these environments in terms of general characteristics that are more or less pronounced across specific products and situations.

Deriving and Organizing the HIPUD Model

Selection Criteria

We derived the model's characteristics through a conceptual comparison with the information environments that preceded Gen-AI chatbots. Specifically, we compared them along five questions: how information comes into being, which needs it serves, to whom it is addressed, what context accompanies it, and how users obtain it. Candidate characteristics were retained if they met four criteria: (1) *distinctiveness*, meaning that the characteristic distinguishes Gen-AI chatbots from prior information environments in kind or in degree; (2) *relevance for information use*, meaning that it plausibly shapes how information is encountered, processed, and used; (3) *non-reducibility*, meaning that it cannot be subsumed under another characteristic; and (4) *generality*, meaning that it applies to standalone Gen-AI chatbots in general—if to varying degrees—rather than to specific products or versions.

Each of the five characteristics of the HIPUD model captures how Gen-AI chatbots answer one of the questions introduced above differently: Information is generated rather than retrieved (probabilistic), serves virtually any need through a single interface (universal), is addressed to one user rather than a public (hyper-personalized), arrives without source and social cues (de-contextualized), and is obtained through dialogue rather than as finished content (interactive). The number of characteristics thus follows from the comparison and the criteria rather than from a target count. Further candidates did not qualify. For example, the ability to edit and revise outputs is a form of interactivity (non-reducibility), and displays of a model's reasoning are offered only by some products (generality). Human-likeness and sycophancy, finally, are better understood as a user perception and a system behavior that primarily arise from

individual characteristics—the interactive and the hyper-personalized one, respectively—than as characteristics in their own right (non-reducibility; see below).

Three Layers

To organize the characteristics of the HIPUD model, we draw on Evans et al.'s (2017) distinction between features, affordances, and outcomes. Here, *features* are properties of a technology (which we extend to its output), *affordances* are relational possibilities for action that emerge between users and a technology and vary in how they are used, and *outcomes* are the consequences or effects of that use. Applied to Gen-AI chatbots, this yields three layers (see Table 1).

Probabilistic and *universal* are features of the system. They describe how Gen-AI chatbots produce content and how far they reach, and users cannot change them through exchanges with the chatbot. *Hyper-personalized* and *de-contextualized* are default features of the output. They describe the information users receive unless they act otherwise. Importantly, they are not outcomes in Evans et al.'s (2017) sense: They describe the informational artifact (i.e., the output) as it is encountered, not its consequences for the user or society at large. *Interactive*, finally, is the only characteristic that constitutes an affordance in the narrow sense: It is relational, it varies in how users take it up, and it is the means through which the defaults of the output can be partly renegotiated. The outcomes of using Gen-AI chatbots (e.g., implications for a shared epistemic basis or privacy) are discussed separately below.

The Five Characteristics

We discuss the characteristics layer by layer, beginning with the two features of the system, followed by the two default features of the output, and ending with interactivity. Because the model's logic lies in these layers rather than in its acronym, this order departs from the

sequence of the name HIPUD, which we retain as a mnemonic. For each characteristic, we outline what it describes, how it breaks with prior information environments, and what it implies for information use.

Table 1

The HIPUD Model: Five Characteristics of Gen-AI Chatbot Information Environments

Characteristic	**Layer**	**Break with prior information environments**
Probabilistic: information is statistically synthesized anew for each query rather than retrieved	*Feature* of the system	Retrieved, authored content → synthesized content
Universal: one interface answers across topics, languages, and modalities	*Feature* of the system	Gateway that points to domain-specific sources → one platform that answers across domains itself
Hyper-personalized: information itself, not only its selection, is tailored to the individual user	(Default) *Feature* of the output	Content addressed to a public → content tailored to one user
De-contextualized: information arrives without source and social cues	(Default) *Feature* of the output	Content embedded in source and social cues → content stripped of both
Interactive: information is obtained and shaped through real-time, reciprocal dialogue	*Affordance*	Finished content → content co-constructed in dialogue

Probabilistic

The first characteristic, *probabilistic*, captures how Gen-AI chatbots produce information: by generating "statistically likely continuations of word sequences" (Shanahan, 2024, p. 72) rather than by retrieving existing documents. Even systems that integrate live web content synthesize it into a generated response, so that the output is grounded in, but not identical to, existing sources. This marks the shift from retrieved, authored content to synthesized content. Each output is thus generated anew for the query at hand, as a combination of words that did not exist in this form before, regardless of who asks. One consequence is a form of epistemic fluidity: Identical queries can yield different answers, even under settings intended to make outputs deterministic (Atil et al., 2024).

Probabilistic generation is what allows Gen-AI chatbots to answer questions for which no single existing document provides an answer, by combining and reformulating information across sources. Yet it also means that not only the selection or presentation of information but its very *substance* may be unreliable. The coherence of Gen-AI outputs derives from statistical patterns in training data rather than from verified knowledge, and these patterns can be biased or incomplete. As Sundar and Liao (2023, p. 170) put it, "ChatGPT is more like the skilled imitator than the original painter." Gen-AI chatbots can thus produce fluent, confident, and seemingly authoritative text even when the underlying information is inaccurate or fabricated, as in the case of so-called *hallucinations* (Huang et al., 2025), which cannot be fully eliminated under current architectures (Z. Xu et al., 2025). Crucially, Gen-AI chatbots typically do not signal uncertainty: Outputs are delivered with uniform confidence, leaving users with few cues to tell accurate, well-supported answers from thinly sourced or fabricated ones. Moreover, because outputs gravitate toward the statistical center of the training data, less common perspectives may be systematically

underrepresented. Peterson (2025, p. 3264) warns that widespread reliance on such systems could lead to a "knowledge collapse," a reduction of available knowledge to the most dominant ideas. The risks of probabilistic generation thus concern not only what Gen-AI chatbots get wrong but also what they leave out.

The implications of probabilistic generation for information use depend largely on users' awareness of it. Users can calibrate their trust in outputs only if they recognize that these are probabilistically generated rather than retrieved. Such awareness cannot be taken for granted, as many people do not even recognize when they are using AI: Although nearly all U.S. adults use products with AI features, only about a third realize that they do (Maese, 2025). This problem may be compounded when probabilistic outputs align with users' preexisting beliefs: Information that confirms one's views is less likely to be questioned, regardless of its source. Indeed, Jacob et al. (2025) found evidence for what they term the "chat-chamber effect," that is, the tendency for Gen-AI chatbots to provide incorrect but proattitudinal information that users internalize without verification.

Universal

The second characteristic, *universal*, captures how Gen-AI chatbots serve as one interface for virtually all information needs. Search engines, too, are topically universal, but they act as gateways that point users to domain-specific sources. Gen-AI chatbots, by contrast, answer across domains themselves, diminishing (if not eliminating) the need to navigate elsewhere. This makes the technology a "convenient one-stop shop for information" (Reiss et al., 2026, p. 4): Instead of separate sources for legal guidance, health concerns, trip planning, or current events, Gen-AI chatbots operate across all these domains (Tafesse & Mamo, 2026).

This characteristic has three interrelated aspects. Thematically, Gen-AI chatbots are trained on vast and varied data and can augment responses with live web content, making them versatile across a wide range of topics (Y. Liu et al., 2025; Srivastava et al., 2023). Linguistically, most widely used systems operate across many natural and programming languages (Jiang et al., 2026; Y. Xu et al., 2025), allowing users to draw on globally available knowledge in their native language—although research indicates that performance is uneven across languages and reproduces dominant linguistic inequalities (Khanna & Li, 2025; Lai et al., 2023). Modally, Gen-AI chatbots process and produce text, audio, images, and video (Caffagni et al., 2024), which can deepen understanding but may also lend unwarranted credibility to false content when it is accompanied by visual evidence (Guo et al., 2025; Newman & Zhang, 2020). In practice, the same interface can help interpret a lab report, translate a letter from a foreign authority, or identify a broken appliance part from a photo, needs that previously required many different sources.

For information use, this has ambivalent implications. It lowers the cost of obtaining information across domains, but it may also undermine the very sources the technology draws on: As users bypass specialized platforms, domain-specific resources risk declining use. Within six months of ChatGPT's release, for instance, activity on the programming Q&A platform Stack Overflow dropped by about 25% relative to comparable forums (del Rio-Chanona et al., 2024). The reach of Gen-AI chatbots is further extended by their integration into other applications. Typically provided as centralized cloud services, they are now embedded in search engines, messengers, office software, or banking apps (van Der Vlist et al., 2024), so that users encounter them in contexts where they may not expect (or even notice) them. Such pervasive availability normalizes Gen-AI as an intermediary of everyday information while making it increasingly

invisible, which may foster less critical engagement, for instance, regarding data privacy (Susser, 2019).

Hyper-Personalized

The third characteristic, *hyper-personalized*, captures how Gen-AI chatbots tailor not only the selection but the content of information to individual users, adapting to their linguistic style, their prompts, their inferred needs, and, where memory functions are enabled, their interaction history (Ai et al., 2025, p. 118; see also Ramaul et al., 2024). Algorithmic personalization on social media and in recommender systems operates through the curation of pre-existing content (Just & Latzer, 2017; Kümpel, 2022). Gen-AI chatbots, by contrast, do not merely curate but create, adjusting explanations to match "the user's reading level, interests, and conversational preferences" (Trippas et al., 2025, p. 83). This marks the shift from content addressed to a public to content tailored to a single user. Importantly, hyper-personalization goes beyond the fact that every output is generated anew (see Probabilistic): It refers to the adaptation of this output to the individual who asks. Notably, this tailoring operates from the very first turn: Chatbot responses largely reflect the framing and assumptions already embedded in a prompt, a user-input bias (N. Sharma et al., 2024) that requires no iterative exchange. Tailoring thus does not presuppose an ongoing dialogue.

As a default feature of the output, this tailoring is set by providers through training that optimizes for user approval, through system prompts, and through memory functions. In Bozdag's (2013) terms, the default corresponds to *implicit* personalization performed by the system, whereas users' requests to change it, for instance by asking for neutral, multi-perspective, or less tailored answers, constitute *explicit* personalization. Such requests, however, change the response only partly (see Interactive). Moreover, optimization for user approval pulls

outputs back toward agreement, as preference-based training can reward responses that match users' views over accurate ones (M. Sharma et al., 2024). This tendency, known as *sycophancy* (see also Cheng et al., 2025), is best understood as the behavior that emerges where tailoring to the user meets such optimization. It can compromise accuracy, which is particularly consequential in domains such as health (Chen et al., 2025).

Hyper-personalization may enhance comprehension and perceived relevance, for instance, when a user asks a Gen-AI chatbot to explain medical test results in plain language or to break down a legal contract. It thus also carries inclusive potential, for example, for people with ADHD, dyslexia, or autism, for whom Gen-AI chatbots may lower barriers to information access (Pierrès et al., 2025; X. Zhao et al., 2025). At the same time, hyper-personalized tailoring creates opportunities for manipulation: Information adapted to individual users can be leveraged to promote products under the guise of neutral advice or to present political actors in a favorable light, leaving "entry points for targeted propaganda efforts" (Mattis & de Vreese, 2025, p. 3616). Moreover, because Gen-AI chatbots are query-driven, users typically receive what they ask for rather than what might challenge them, which reduces incidental exposure to diverse viewpoints that social media can provide (Nanz & Matthes, 2022). Beyond the individual, such tailoring may also contribute to a fragmented information landscape, an implication we return to below.

De-Contextualized

The fourth characteristic, *de-contextualized*, captures how Gen-AI chatbots present information without the cues that typically accompany mediated communication. Two kinds of cues are affected. Source cues—authorship, institutional origin, and markers of credibility conveyed through design and tone (Fogg et al., 2001; Metzger et al., 2010; Robins & Holmes, 2008)—are muted or absent, and often even the origin of information remains opaque (Burrell,

2016; Hassija et al., 2024; Liao & Wortman Vaughan, 2024). Social cues—who shared a piece of information, how others reacted to it, and how widely it circulated (Haim et al., 2018; Lin et al., 2016)—are absent altogether. Social media, too, detach content from its original context, but they embed it in a new social layer (Kümpel, 2022). Gen-AI chatbots remove both layers and present information in the system's own voice. This marks the shift from content embedded in source and social cues to content stripped of both. In addition, Gen-AI chatbots lack the feedback loops of earlier environments. Unlike on social media, where individual actions shape what others encounter (Eg et al., 2023), or on search engines, which rely on aggregated user behavior (Baeza-Yates et al., 2004; Joachims, 2002), individual use remains disconnected from collective experience, making it a dyadic (Seering et al., 2019) and largely asocial experience.

As a default feature of the output, this characteristic is set through interface and response design. Users can request sources, and some systems display them by default. Yet this recovers only a fraction of the original context. Whether the citations reflect what actually informed an answer remains uncertain (see below), and much of what accompanies the original source is lost: its design and layout, its own language and style, its emphases, and any embedded or related content. Information is still rendered in the chatbot's voice, and social cues cannot be recovered at all. The core of the characteristic thus persists even when users ask for sources.

Again, for information use, the implications are ambivalent. On the one hand, the absence of cues may reduce distraction by keeping the focus on the content itself and attenuate biases associated with socially embedded information, such as the bandwagon effects of popularity cues (Haim et al., 2018). On the other hand, it deprives users of the cues they (used to) rely on to assess credibility and quality, and content of a distinctive character, such as satire or irony, may

be flattened. Moreover, outputs are ephemeral and visible only to the individual user. Gen-AI outputs are thus publicly invisible and difficult to scrutinize or contest.

Interactive

The fifth characteristic, *interactive*, captures how information use with Gen-AI chatbots unfolds as a real-time, reciprocal exchange. Rather than submitting queries and receiving static lists of results, users can refine their questions, request clarification, revise outputs, or explore tangential topics as the conversation progresses (Amer & Elboghdadly, 2024). Interactive information environments are, of course, not new as such: Search engines allow users to refine queries, and social media enable exchange among users. What is new is that the interlocutor is also the producer of the content. This marks the shift from finished content to content co-constructed in dialogue—an exchange that, from an audience perspective, can feel like a genuinely interpersonal one (Folk et al., 2024; Kleinert et al., 2026).

As an affordance in the narrow sense, interactivity is the means through which users can adjust the default features of the output described above: They can ask for sources, for a neutral framing, or for a less tailored answer. Such renegotiation, however, is limited in two respects.

First, renegotiation remains partial. The two features of the system (probabilistic, universal) lie beyond its reach altogether: No request can yield information that is not probabilistically generated, and users can narrow their own use but not the system's scope. Even the default features of the output can be adjusted only at the margins: Whatever users request, the response remains tailored to them and stripped of most of the context of its sources. A request for neutrality is itself a prompt to which the system adapts, and requested sources are generated as well: In an audit of generative search engines, only about half of the generated sentences were fully supported by their citations, and about a quarter of the citations did not support the sentence

they accompanied (N. Liu et al., 2023). Likewise, a model's account of how it reached an answer may not reflect the process that actually produced it (Turpin et al., 2023).

Second, defaults tend to persist. Across domains, preset options exert a robust influence on behavior, not least because people read them as an implicit recommendation by whoever set them (Jachimowicz et al., 2019). Whether users renegotiate a default depends on three conditions (Davis & Chouinard, 2016, pp. 245–246): *perception* (do they recognize it as a default?), *dexterity* (can they formulate an effective request to change it?), and *cultural and institutional legitimacy* (does it seem appropriate to question the system?). Neither of the first two conditions can be taken for granted in Gen-AI chatbot environments. Perception is weakest in embedded interfaces such as AI Overviews, where users may not perceive themselves to be in an interaction at all and may continue to act as they would with a conventional search engine. Dexterity, in turn, is often limited: Non-experts frequently struggle to translate their goals into effective prompts (Zamfirescu-Pereira et al., 2023), and when dissatisfied with a response, users frequently make no attempt to correct it (Kim et al., 2024). Accordingly, many exchanges leave little room for renegotiation in the first place: In a large corpus of logged ChatGPT interactions, most conversations consisted of a single user turn (W. Zhao et al., 2024). At the population level, what most users receive is therefore the default.

The interactivity of Gen-AI chatbots lowers the barrier to obtaining relevant information. Users need not translate a vague need into precise search terms from the outset; instead, the chatbot takes over part of what Taylor (1968) described as the "question negotiation" once performed by librarians, clarifying and narrowing a request over successive turns. Information use thereby becomes a collaborative process in which users co-construct the information they

receive (Ju & Stewart, 2024), so that what they end up with depends on how they steer the exchange.

Yet this means of renegotiation is not neutral. Users can employ it to override defaults, but also to push the system toward the answers they prefer: AI assistants frequently abandon correct answers when users merely question them (M. Sharma et al., 2024). Moreover, reliability tends to decline over the course of an exchange, as models commit to early assumptions and fail to recover from them (Laban et al., 2025). Each turn also tells the system more about the user, deepening the hyper-personalized tailoring described above and encouraging gradual self-disclosure (Ischen et al., 2020). Finally, natural dialogue is a central reason why Gen-AI chatbots are perceived as social and human-like (Kang et al., 2026; Schuetzler et al., 2020). Consistent with our selection criteria, we treat such human-likeness not as a characteristic of the information environment but as a user perception arising from the interactive characteristic. This perception is double-edged (Peter et al., 2025; Reinecke et al., 2025): It may increase engagement and willingness to disclose information, but it may also invite what Peter et al. (2025, p. 2) term "anthropomorphic seduction," an uncritical acceptance of AI-generated content driven by the system's communicative abilities rather than by the accuracy of its outputs.

HIPUD in Concert: Outcomes of Gen-AI Chatbot Use

The five characteristics do not operate in isolation. Taken together, they point to broader transformations that, in Evans et al.'s (2017) terms, constitute the *outcomes* of using Gen-AI chatbots: their emergence as epistemic authorities, the potential erosion of a shared epistemic basis, and new forms of informational vulnerability.

Epistemic Authority

Epistemic authorities are the persons, institutions, or technologies people rely on to know better than they do themselves. Following Bartsch et al. (2025), who build on Jäger (2024), a source fulfills this functional role when it is subjectively recognized as having an advanced epistemic position, socially claimed as such, and objectively justified, that is, actually competent and reliable in helping reach one's epistemic goals. Bartsch et al. (2025) explicitly include AI among potential epistemic authorities. The HIPUD model specifies why Gen-AI chatbots are likely to be recognized as such, and why this recognition may outrun its justification.

Recognition is fostered by several characteristics at once. Gen-AI chatbots speak fluently and confidently on virtually any topic, whether or not users know anything about it themselves, combining the breadth of universality with the uniform confidence of probabilistic generation. Interactivity allows them to respond to this user's specific question and follow-ups, and hyper-personalization pitches explanations at the user's level of understanding. De-contextualization, moreover, obscures authorship and origin, making it difficult to trace where claims come from or whose perspectives they reflect. As a result, information may be attributed to the chatbot itself rather than to any underlying human or institutional source, a dynamic consistent with the machine heuristic, whereby users rely on the perceived competence of a system rather than evaluating its outputs (Sundar & Liao, 2023).

Justification, by contrast, is limited by probabilistic generation: Outputs are often accurate but not reliably so, and fluency is easily mistaken for truth. The social claim, finally, is ambivalent. Providers market their systems as highly capable, and the outputs themselves adopt the assured, authoritative style that Bartsch et al. (2025, p. 38) identify as a non-epistemic means of claiming authority. At the same time, outputs come with disclaimers that the system 'can

make mistakes.' Gen-AI chatbots thus straddle the two forms of false authority distinguished by Jäger (2024): Like fake authorities, they are presented as competent; like unintended authorities, they formally disclaim the very status they are granted, much as influencers mark their advice as being just for entertainment (see also Bartsch et al., 2025, p. 43). This ambivalence makes the risk of false authority structural rather than incidental: Providers benefit from users' recognition while shifting responsibility for verification onto users, who, as argued above, are often poorly equipped to do so.

The result is a form of epistemic authority that is both diffuse and potent. Unlike scientists or journalists, whose credibility is grounded in transparent, institutionalized processes of knowledge production that are subject to public scrutiny, Gen-AI chatbots derive their authority primarily from perceptual cues. This may discourage verification, reduce engagement with primary sources, and foster dependency on single systems. It also makes their authority derivative: It rests on the very sources the chatbots synthesize, while their use may weaken the visibility, use, and legitimacy of precisely these sources (see Universal). This dynamic may be reinforced where established authorities contradict one another or are distrusted (Bartsch et al., 2025): Confronted with competing signals of truth, users may welcome a source that delivers an orderly, confident synthesis.

Erosion of a Shared Epistemic Basis

In many traditional information environments, people may disagree about interpretations while still orienting themselves toward common institutions or sources. Gen-AI chatbot use, by contrast, is characterized by outputs generated for a specific user in a specific interaction. Hyper-personalization creates individualized micro-realities of meaning, de-contextualization removes shared reference points, public visibility, and the feedback loops that connect individual use to

collective experience, and universality extends this logic across nearly all life domains. As a result, users may increasingly rely on singularized information environments that are difficult for others to inspect and contest. This does not necessarily imply total fragmentation. Because probabilistic generation privileges central and widely represented patterns, a few dominant models may even standardize prevailing interpretations. Emerging monetization models may further reinforce this centralizing tendency: If providers integrate commercially driven response priorities or content placement, the information users receive would also be shaped by the economic interests of a small number of corporations. Such standardization would coexist with individualized presentation, producing a paradoxical combination of centralization and singularization. The broader consequence may be not only a weakening of common reference points for public discourse but also a more subjective mode of knowledge formation, in which what counts as relevant becomes increasingly personalized.

Privacy and Informational Vulnerability

The capacity of Gen-AI chatbots to hyper-personalize outputs depends, at least in part, on access to detailed information about the users. Interactivity encourages users to disclose such information gradually and with little friction, especially where they perceive the exchange as human-like (see above), and memory functions retain these disclosures across conversations. Universality expands the breadth of disclosure, because the same system accompanies users across work, health, finances, relationships, and other sensitive domains, and multimodal input deepens it further when users upload medical records or personal photographs. The result is a broader and more intimate data relation than in earlier information environments. Privacy risks thus arise not only from isolated acts of data collection but from the cumulative integration of personal, cross-domain information into centralized infrastructures. This information meets a

technology increasingly capable of inferring personal attributes from it (Staab et al., 2024) and may feed into model training data (King et al., 2025), raising risks of data leakage or unauthorized repurposing (Carlini et al., 2021; Nasr et al., 2025). Gen-AI chatbots may thus encourage users to share more than they realize, across more domains than they intend, with less control over downstream use than they assume.

These outcomes are interconnected. Recognizing Gen-AI chatbots as epistemic authorities may facilitate the erosion of a shared epistemic basis, as users who trust a system uncritically are less likely to seek out common reference points, and both dynamics may lower attentiveness to privacy risks. How pronounced these outcomes become, however, is not fixed: It depends on the specific constellation of characteristics users encounter in Gen-AI chatbot information environments—some fixed by the system, others partly open to renegotiation—and on how they engage with it.

Applying the HIPUD Model in Communication Research

A key claim of the HIPUD model is that it offers an analytical lens applicable across a range of subfields in communication research—not by providing a one-size-fits-all template, but by making visible how specific constellations of characteristics shape information use in different contexts. Rather than replacing established theories, the model specifies the conditions under which they apply to Gen-AI chatbots. To illustrate this, we present three application scenarios spanning political communication, health communication, and consumer communication. These are intended as starting points; each scenario opens onto a broader range of questions and approaches than we can develop here.

Political Communication

Consider a citizen preparing for an upcoming election. Rather than consulting party websites or voting advice applications, she opens a Gen-AI chatbot and asks: "Which party best matches my views on climate and migration policy?" The chatbot responds with a tailored comparison that is not retrieved from a static database, but generated in real time based on her stated preferences (expressed directly or inferred from prior interactions) and the model's probabilistic synthesis.

Four characteristics are particularly salient here. Probabilistic generation means that the comparison reflects statistical patterns in the training data rather than verified policy positions, creating the risk of subtle inaccuracies that may go unnoticed because the output reads as authoritative. Such concerns are not hypothetical: In an audit of a Gen-AI chatbot before the 2023 Swiss and German state elections, about a third of answers contained factual errors (AlgorithmWatch & AI Forensics, 2023). Outputs may also be systematically skewed in ways users cannot trace: Which arguments and sources a response draws on depends on what enters training data and retrieval pipelines—and thus potentially on factors unrelated to epistemic quality, such as licensing agreements between providers and publishers (OpenAI, 2023). Whether a comparison favors one party's position because it is better supported or because of such factors cannot be discerned from the output itself. Hyper-personalization means that the political information she receives is not a standardized overview but an individualized response shaped by her prompting, her choice of policy domains, and potentially her interaction history. De-contextualization strips the information of the source cues that typically accompany political communication: Unlike a newspaper article or a social media post, the chatbot's response is not attributable to a journalist, party, or fellow citizen and thus escapes the kind of collective

scrutiny that characterizes public discourse. Finally, interactivity allows her to ask follow-up questions, clarify political terms, or request the parties' own statements, although such requests recover the original context only in part, and few users are likely to make them.

Here, HIPUD shifts the analytical focus from *whether* Gen-AI chatbots differ from traditional media to *how* their specific constellation of characteristics shapes political communication. Research on media credibility and perceived bias, for instance, has typically examined messages from identifiable sources that reach a broad audience. Chatbot outputs, by contrast, are commonly neither attributable to such sources nor shared with others. This raises questions such as: How does this configuration affect the perceived neutrality of political information? Does cross-domain use (i.e., universality) moderate users' critical engagement with political content? Methodologically, the distinction between features of the system and default features of the output supports targeted experimental designs. As defaults set by providers, hyper-personalization and de-contextualization can be varied systematically, for instance, by manipulating whether chatbot responses include source attributions or by comparing tailored with standardized outputs, while measuring effects on political knowledge or voting intentions. Experimental designs can also capture whether participants renegotiate these defaults when given the opportunity, for example, by requesting sources or alternative perspectives.

Health Communication

A second application concerns the health domain. Imagine a user who notices a persistent skin irritation, takes a photo with their smartphone, and uploads it with the prompt: "I've had this skin issue for a few days, what could it be?" Unlike a search engine query returning a list of links, the chatbot responds conversationally and may ask follow-up questions about their age, skin care, or known allergies. What unfolds is not a one-shot retrieval of information but a

dialogic process that progressively shapes both the content provided and the user's understanding of their condition.

Here, universality (in the form of multimodality), interactivity, and hyper-personalization are tightly intertwined. The photo already allows the system to tailor its first response, and each follow-up turn adds further health-related information that deepens this tailoring. Users may even disclose more than they would to a human, as the exchange is perceived as more anonymous and non-judgmental (Croes et al., 2024). The result is a quasi-clinical interaction that is neither subject to professional oversight nor constrained by the ethical obligations governing medical advice. Probabilistic generation adds particular risk in this domain, as inaccurate assessments are delivered with the same confidence as accurate ones. Hyper-personalization, in turn, may lead the chatbot to confirm a user's self-diagnosis rather than challenge it, given the sycophantic tendencies discussed above. De-contextualization compounds this risk: Unlike medical information websites, the chatbot's output lacks institutional markers that signal provenance and accountability. And while providers increasingly embed disclaimers in medical contexts, such disclaimers generally appear to have little effect on users' verification behavior (Knor et al., 2026).

For health communication research, HIPUD offers a framework that goes beyond measuring whether Gen-AI chatbots provide accurate medical information (though this remains important). A model's accuracy alone, after all, says little about what users obtain: In a preregistered randomized study, large language models alone identified relevant conditions in about 95% of medical scenarios, whereas participants using them did so in fewer than 35% of cases and performed no better than a control group (Bean et al., 2026). Rather than focusing on isolated capabilities, HIPUD directs attention to the communicative processes through which

health information is co-constructed and acted upon. Research on online health information seeking and self-disclosure has largely examined disclosure to human professionals or the retrieval of existing content; in chatbot interactions, by contrast, users' disclosures become the very input from which the health information they receive is generated. This raises questions such as: How do multimodal capabilities and interactivity alter disclosure behavior compared to more static web searches? Under what conditions does hyper-personalization foster uncritical acceptance of health misinformation? Methodologically, focusing now on qualitative research (e.g., interviews with think-aloud protocols or diary studies) as an example, HIPUD provides a coding vocabulary that captures which characteristics users perceive in health-related interactions and how these perceptions shape their information evaluation and subsequent medical decisions.

Consumer Communication

A third exemplary application extends the model to consumer decision-making. Consider a user planning to purchase an e-bike who turns to a Gen-AI chatbot: "I need an e-bike for city commuting, budget around 2,000 euros. What do you recommend?" The chatbot synthesizes a comparison of models for him, weighing features against the user's stated needs. While this resembles advice found on product review websites, the underlying information dynamics differ in ways HIPUD makes analytically visible.

De-contextualization fundamentally alters the availability of cues. Traditional consumer information environments are rich in social and institutional context: Review platforms aggregate user ratings, display verified-purchase labels, and host product discussions; even commercial websites, though optimized for brand presentation, still make their provenance and intent legible. The chatbot's recommendation, by contrast, arrives as an institutionally and

socially disembedded artifact that obscures its sources as it synthesizes them. Universality and hyper-personalization combine to position the chatbot as a comprehensive, individually attuned advisor, which may reduce the motivation to consult additional sources. This raises questions about commercial neutrality, particularly as Gen-AI providers begin to explore monetization models such as sponsored placements or affiliate partnerships.

HIPUD enables consumer communication research to pose more precise questions. Research on persuasion knowledge, for instance, assumes that consumers can recognize a persuasive attempt and its source, which allows them to activate coping strategies (Friestad & Wright, 1994)—an assumption developed largely in the context of traditional advertising (Isaac & Calder, 2025). A de-contextualized, individually tailored recommendation offers few such signals. This raises questions such as: How does the absence of visual brand identity or social context cues affect trust in product recommendations? How does the chatbot's role as a universal, personalized advisor shape the breadth of consumers' information search? Methodologically, computational analyses of naturally occurring chatbot interaction logs offer a promising avenue. For this, HIPUD provides the analytical categories to systematically code such interaction traces at scale: for instance, identifying how frequently users iterate on product requirements (interactivity), what personal information they disclose (which feeds hyper-personalization), or whether they explicitly request sources (renegotiation of de-contextualization). This allows researchers to move beyond surface-level prompt analyses toward an understanding of consumer behavior grounded in the characteristics of Gen-AI information environments.

Conclusion

Gen-AI chatbots reconfigure everyday information use. Where earlier media delivered content that already existed elsewhere—authored by identifiable persons or institutions and, in principle, available to others—Gen-AI chatbots produce their outputs in the moment of the query: statistically synthesized, available for virtually any need, tailored to the individual user, stripped of the contextual anchors people used to rely on, and shaped through dialogue. The HIPUD model gives this transformation a name and a structure.

Its value lies less in predicting specific outcomes than in equipping researchers to investigate them with precision. By resolving broad claims about Gen-AI's novelty into five characteristics organized in three layers—features of the system, default features of the output, and an affordance—HIPUD makes it possible to ask which characteristic matters in which setting, how characteristics interact, and which of them users can, at least partly, renegotiate. As the application scenarios illustrate, the model does not replace established communication theories but specifies where their assumptions require adjustment.

Beyond research, HIPUD points to concrete goals for AI literacy and governance. From a user-centered perspective, AI literacy would need to enable people to recognize the default features of the output as defaults, to renegotiate them where possible, and to understand the inherent limits of such renegotiation, which are set by the features of the system. Training users to prompt for diversity is one such strategy, but its reach remains bounded by the very characteristics it seeks to counter. Moreover, because defaults tend to persist, literacy alone will not suffice: What most users receive is what providers have set. This makes defaults, and other implications of Gen-AI chatbots, a matter of governance and HIPUD in turn an analytical lens for regulators to understand and study their relevance. For instance, current regulation primarily

addresses whether people know that they are interacting with AI (e.g., the EU AI Act requires that they be informed when this is the case). From the perspective of HIPUD, regulators could also address the default features of the output, such as whether sources are displayed or how far outputs are tailored to users and their stored interaction history.

Several limitations mark this starting point. HIPUD has been developed primarily with standalone Gen-AI chatbots in mind. Its boundary conditions remain to be established for embedded applications such as AI Overviews, as well as for AI companions and agentic systems. Moreover, HIPUD does not itself supply the mechanisms through which its characteristics shape trust, disclosure, or persuasion; these remain the domain of established theories such as those of source credibility or self-disclosure. What HIPUD offers is the *lens* that shows where and why these theories need to be specified for Gen-AI chatbots. Nor does it, on its own, arbitrate between optimistic and critical accounts of Gen-AI's role in public life; rather, it provides the analytical precision required to make such debates empirically tractable. Finally, the technology will continue to evolve. Still, there are good reasons to expect the five characteristics to remain relevant, as each is anchored in something more stable than particular products. Better retrieval may reduce errors, for instance, but answers will still be synthesized rather than retrieved; and as long as providers compete for users, they have every incentive to keep their systems broad and their outputs tailored. What is likely to change is the relative weight of the characteristics, not their relevance.

Ultimately, the questions HIPUD opens extend beyond Gen-AI chatbots. They concern how communication theory must accommodate environments in which information is less something accessed than something produced in the moment of inquiry, and in which the institutional and social scaffolds of mediated communication can no longer be taken for granted.

Engaging with these questions is, in our view, among the field's central theoretical tasks in the coming years.

References

Ai, Q., Dou, Z., & Zhang, M. (2025). Improving Generative Information Retrieval Systems Based on User Feedback. In R. W. White & C. Shah (Eds.), *Information Access in the Era of Generative AI* (pp. 111–133). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-73147-1_5

AlgorithmWatch & AI Forensics. (2023). *Prompting Elections: The Reliability of Generative AI in the 2023 Swiss and German Elections. An Analysis of Microsoft's Bing Chat*. https://algorithmwatch.org/de/wp-content/uploads/2023/12/AlgorithmWatch_AIForensics_Bing_Chat_Report.pdf

Amer, E., & Elboghdadly, T. (2024). The End of the Search Engine Era and the Rise of Generative AI: A Paradigm Shift in Information Retrieval. *2024 International Mobile, Intelligent, and Ubiquitous Computing Conference (MIUCC)*, 374–379. https://doi.org/10.1109/MIUCC62295.2024.10783559

Aspray, W., & Hayes, B. M. (Eds.). (2011). *Everyday Information: The Evolution of Information Seeking in America*. MIT Press.

Atil, B., Aykent, S., Chittams, A., Fu, L., Passonneau, R. J., Radcliffe, E., Rajagopal, G. R., Sloan, A., Tudrej, T., Ture, F., Wu, Z., Xu, L., & Baldwin, B. (2024). *Non-Determinism of "Deterministic" LLM Settings* (arXiv:2408.04667). arXiv. https://doi.org/10.48550/arXiv.2408.04667

Baeza-Yates, R., Hurtado, C., & Mendoza, M. (2004). Query Recommendation Using Query Logs in Search Engines. In W. Lindner, M. Mesiti, C. Türker, Y. Tzitzikas, & A. I.

Vakali (Eds.), *Current Trends in Database Technology—EDBT 2004 Workshops* (Vol. 3268, pp. 588–596). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-30192-9_58

Bartsch, A., Neuberger, C., Stark, B., Karnowski, V., Maurer, M., Pentzold, C., Quandt, T., Quiring, O., & Schemer, C. (2025). Epistemic Authority in the Digital Public Sphere. An Integrative Conceptual Framework and Research Agenda. *Communication Theory*, *35*(1), 37–50. https://doi.org/10.1093/ct/qtae020

Bates, M. J. (1989). The Design of Browsing and Berrypicking Techniques for the Online Search Interface. *Online Review*, *13*(5), 407–424. https://doi.org/10.1108/eb024320

Bean, A. M., Payne, R. E., Parsons, G., Kirk, H. R., Ciro, J., Mosquera-Gómez, R., Hincapié M, S., Ekanayaka, A. S., Tarassenko, L., Rocher, L., & Mahdi, A. (2026). Reliability of LLMs as Medical Assistants for the General Public: A Randomized Preregistered Study. *Nature Medicine*, *32*(2), 609–615. https://doi.org/10.1038/s41591-025-04074-y

Bozdag, E. (2013). Bias in Algorithmic Filtering and Personalization. *Ethics and Information Technology*, *15*(3), 209–227. https://doi.org/10.1007/s10676-013-9321-6

Briggs, A., Burke, P., & Ytreberg, E. (2020). *A Social History of the Media*. Polity.

Burrell, J. (2016). How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms. *Big Data & Society*, *3*(1), 2053951715622512. https://doi.org/10.1177/2053951715622512

Caffagni, D., Cocchi, F., Barsellotti, L., Moratelli, N., Sarto, S., Baraldi, L., Baraldi, L., Cornia, M., & Cucchiara, R. (2024). The Revolution of Multimodal Large Language Models: A Survey. *Findings of the Association for Computational Linguistics: ACL 2024*, 13590–13618. https://doi.org/10.18653/v1/2024.findings-acl.807

Carlini, N., Tramèr, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., & Raffel, C. (2021). *Extracting Training Data from Large Language Models*. 2633–2650. https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting

Chatterji, A., Cunningham, T., Deming, D., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). *How People Use ChatGPT* (No. 34255). National Bureau of Economic Research. https://doi.org/10.3386/w34255

Chen, S., Gao, M., Sasse, K., Hartvigsen, T., Anthony, B., Fan, L., Aerts, H., Gallifant, J., & Bitterman, D. S. (2025). When Helpfulness Backfires: LLMs and the Risk of False Medical Information Due to Sycophantic Behavior. *npj Digital Medicine*, *8*(1), 605. https://doi.org/10.1038/s41746-025-02008-z

Cheng, M., Yu, S., Lee, C., Khadpe, P., Ibrahim, L., & Jurafsky, D. (2025). *ELEPHANT: Measuring and Understanding Social Sycophancy in LLMs* (arXiv:2505.13995). arXiv. https://doi.org/10.48550/arXiv.2505.13995

Croes, E. A. J., Antheunis, M. L., van der Lee, C., & de Wit, J. M. S. (2024). Digital Confessions: The Willingness to Disclose Intimate Information to a Chatbot and its Impact on Emotional Well-Being. *Interacting with Computers, 36*(5), 279–292. https://doi.org/10.1093/iwc/iwae016

Davis, J. L., & Chouinard, J. B. (2016). Theorizing Affordances: From Request to Refuse. *Bulletin of Science, Technology & Society*, *36*(4), 241–248. https://doi.org/10.1177/0270467617714944

del Rio-Chanona, R. M., Laurentsyeva, N., & Wachs, J. (2024). Large Language Models Reduce Public Knowledge Sharing on Online Q&A Platforms. *PNAS Nexus*, *3*(9), pgae400. https://doi.org/10.1093/pnasnexus/pgae400

Eg, R., Demirkol Tønnesen, Ö., & Tennfjord, M. K. (2023). A Scoping Review of Personalized User Experiences on Social Media: The Interplay Between Algorithms and Human Factors. *Computers in Human Behavior Reports*, *9*, 100253. https://doi.org/10.1016/j.chbr.2022.100253

Evans, S. K., Pearce, K. E., Vitak, J., & Treem, J. W. (2017). Explicating Affordances: A Conceptual Framework for Understanding Affordances in Communication Research. *Journal of Computer-Mediated Communication*, *22*(1), 35–52. https://doi.org/10.1111/jcc4.12180

Eveland, W. P. (2003). A "Mix of Attributes" Approach to the Study of Media Effects and New Communication Technologies. *Journal of Communication*, *53*(3), 395–410. https://doi.org/10.1111/j.1460-2466.2003.tb02598.x

Fogg, B. J., Marshall, J., Laraki, O., Osipovich, A., Varma, C., Fang, N., Paul, J., Rangnekar, A., Shon, J., Swani, P., & Treinen, M. (2001). What Makes Web Sites Credible? A Report on a Large Quantitative Study. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 61–68. https://doi.org/10.1145/365024.365037

Folk, D., Yu, S., & Dunn, E. (2024). Can Chatbots Ever Provide More Social Connection Than Humans? *Collabra: Psychology*, *10*(1), 117083. https://doi.org/10.1525/collabra.117083

Friestad, M., & Wright, P. (1994). The Persuasion Knowledge Model: How People Cope with Persuasion Attempts. *Journal of Consumer Research*, *21*(1), 1–31. https://doi.org/10.1086/209380

Gil de Zúñiga, H., Goyanes, M., & Durotoye, T. (2024). A Scholarly Definition of Artificial Intelligence (AI): Advancing AI as a Conceptual Framework in Communication Research. *Political Communication*, *41*(2), 317–334. https://doi.org/10.1080/10584609.2023.2290497

Guo, S., Zhong, Y., & Hu, X. (2025). People Are More Susceptible to Misinformation with Realistic AI-Synthesized Images That Provide Strong Evidence to Headlines. *Harvard Kennedy School Misinformation Review*. https://doi.org/10.37016/mr-2020-189

Haim, M., Kümpel, A. S., & Brosius, H.-B. (2018). Popularity Cues in Online Media: A Review of Conceptualizations, Operationalizations, and General Effects. *Studies in Communication | Media*, *7*(2), 186–207. https://doi.org/10.5771/2192-4007-2018-2-58

Hasebrink, U., & Domeyer, H. (2010). Zum Wandel von Informationsrepertoires in konvergierenden Medienumgebungen [On the Transformation of Information Repertoires in Converging Media Environments]. In M. Hartmann & A. Hepp (Eds.), *Die Mediatisierung der Alltagswelt* (pp. 49–64). VS Verlag. https://doi.org/10.1007/978-3-531-92014-6_4

Hassija, V., Chamola, V., Mahapatra, A., Singal, A., Goel, D., Huang, K., Scardapane, S., Spinelli, I., Mahmud, M., & Hussain, A. (2024). Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence. *Cognitive Computation*, *16*(1), 45–74. https://doi.org/10.1007/s12559-023-10179-8

Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., & Liu, T. (2025). A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Transactions on Information Systems, 43*(2), Article 42. https://doi.org/10.1145/3703155

Isaac, M. S., & Calder, B. J. (2025). Thirty Years of Persuasion Knowledge Research: From Demonstrating Effects to Building Theory to Increasing Applicability. *Consumer Psychology Review*, *8*(1), 3–14. https://doi.org/10.1002/arcp.1107

Ischen, C., Araujo, T., Voorveld, H., Van Noort, G., & Smit, E. (2020). Privacy Concerns in Chatbot Interactions. In A. Følstad, T. Araujo, S. Papadopoulos, E. L.-C. Law, O.-C. Granmo, E. Luger, & P. B. Brandtzaeg (Eds.), *Chatbot Research and Design* (Vol. 11970, pp. 34–48). Springer International Publishing. https://doi.org/10.1007/978-3-030-39540-7_3

Jachimowicz, J. M., Duncan, S., Weber, E. U., & Johnson, E. J. (2019). When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects. *Behavioural Public Policy*, *3*(2), 159–186. https://doi.org/10.1017/bpp.2018.43

Jacob, C., Kerrigan, P., & Bastos, M. (2025). The Chat-Chamber Effect: Trusting the AI Hallucination. *Big Data & Society*, *12*(1), 20539517241306345. https://doi.org/10.1177/20539517241306345

Jäger, C. (2024). False Authorities. *Acta Analytica*, *39*(4), 643–661. https://doi.org/10.1007/s12136-024-00594-3

Jiang, J., Wang, F., Shen, J., Kim, S., & Kim, S. (2026). A Survey on Large Language Models for Code Generation. *ACM Transactions on Software Engineering and Methodology*, *35*(2), 1–72. https://doi.org/10.1145/3747588

Joachims, T. (2002). Optimizing Search Engines Using Clickthrough Data. *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 133–142. https://doi.org/10.1145/775047.775067

Ju, B., & Stewart, J. B. (2024). Empowering Users with ChatGPT and Similar Large Language Models (LLMs): Everyday Information Needs, Uses, and Gratification. *Proceedings of the Association for Information Science and Technology*, *61*(1), 172–182. https://doi.org/10.1002/pra2.1018

Just, N., & Latzer, M. (2017). Governance by Algorithms: Reality Construction by Algorithmic Selection on the Internet. *Media, Culture & Society*, *39*(2), 238–258. https://doi.org/10.1177/0163443716643157

Kang, B., Kim, J., Yun, T., Bae, H., & Kim, C.-E. (2026). Identifying Features That Shape Perceived Consciousness in LLM-Based AI: A Quantitative Study of Human Responses. *Computers in Human Behavior Reports*, *21*, 100901. https://doi.org/10.1016/j.chbr.2025.100901

Kari, J. (2010). Diversity in the Conceptions of Information Use. *Information Research, 15*(3). https://informationr.net/ir/15-3/colis7/colis709.html

Khanna, S., & Li, X. (2025). *Invisible Languages of the LLM Universe* (arXiv:2510.11557). arXiv. https://doi.org/10.48550/arXiv.2510.11557

Kim, Y., Lee, J., Kim, S., Park, J., & Kim, J. (2024). Understanding Users' Dissatisfaction with ChatGPT Responses: Types, Resolving Tactics, and the Effect of Knowledge Level. *Proceedings of the 29th International Conference on Intelligent User Interfaces, IUI '24*, 385–404. https://doi.org/10.1145/3640543.3645148

King, J., Klyman, K., Capstick, E., Saade, T., & Hsieh, V. (2025). User Privacy and Large Language Models: An Analysis of Frontier Developers' Privacy Policies. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, *8*(2), 1465–1477. https://doi.org/10.1609/aies.v8i2.36646

Kleinert, T., Waldschütz, M., Blau, J., Heinrichs, M., & Schiller, B. (2026). AI Outperforms Humans in Establishing Interpersonal Closeness in Emotionally Engaging Interactions, but Only When Labelled as Human. *Communications Psychology*, *4*(1), 23. https://doi.org/10.1038/s44271-025-00391-7

Knor, E. L., Reiss, M. V., Möller, J., & Merten, L. (2026). Determinants of Verification Behavior in Generative Search: Evidence from a Conjoint Experiment. *Computers in Human Behavior Reports*, *22*, 101056. https://doi.org/10.1016/j.chbr.2026.101056

Kümpel, A. S. (2022). Social Media Information Environments and Their Implications for the Uses and Effects of News: The PINGS Framework. *Communication Theory*, *32*(2), 223–242. https://doi.org/10.1093/ct/qtab012

Kümpel, A. S., Anter, L., & Unkel, J. (2022). What Does "Being Informed" Mean? Assessing Social Media Users' Self-Concepts of Informedness. *Media and Communication*, *10*(3), 93–103. https://doi.org/10.17645/mac.v10i3.5310

Laban, P., Hayashi, H., Zhou, Y., & Neville, J. (2025). *LLMs Get Lost In Multi-Turn Conversation* (arXiv:2505.06120). arXiv. https://doi.org/10.48550/arXiv.2505.06120

Lai, V., Ngo, N., Pouran Ben Veyseh, A., Man, H., Dernoncourt, F., Bui, T., & Nguyen, T. (2023). ChatGPT Beyond English: Towards a Comprehensive Evaluation of Large Language Models in Multilingual Learning. *Findings of the Association for Computational Linguistics: EMNLP 2023*, 13171–13189. https://doi.org/10.18653/v1/2023.findings-emnlp.878

Liao, Q. V., & Wortman Vaughan, J. (2024). AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *Harvard Data Science Review*, (Special Issue 5). https://doi.org/10.1162/99608f92.8036d03b

Lin, X., Spence, P. R., & Lachlan, K. A. (2016). Social Media and Credibility Indicators: The Effect of Influence Cues. *Computers in Human Behavior*, *63*, 264–271. https://doi.org/10.1016/j.chb.2016.05.002

Liu, N., Zhang, T., & Liang, P. (2023). Evaluating Verifiability in Generative Search Engines. *Findings of the Association for Computational Linguistics: EMNLP 2023*, 7001–7025. https://doi.org/10.18653/v1/2023.findings-emnlp.467

Liu, Y., Cao, J., Liu, C., Ding, K., & Jin, L. (2025). Datasets for Large Language Models: A Comprehensive Survey. *Artificial Intelligence Review, 58*(12), Article 403. https://doi.org/10.1007/s10462-025-11403-7

Maese, E. (2025, January 15). *Americans Use AI in Everyday Products Without Realizing It*. Gallup. https://news.gallup.com/poll/654905/americans-everyday-products-without-realizing.aspx

Mattis, N., & de Vreese, C. (2025). Breaking the News? Generative AI's Impact on Journalism and Its Implications for Disinformation. *International Journal of Communication*, *19*(2025), 3602–3625.

Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010). Social and Heuristic Approaches to Credibility Evaluation Online. *Journal of Communication*, *60*(3), 413–439. https://doi.org/10.1111/j.1460-2466.2010.01488.x

Nanz, A., & Matthes, J. (2022). Democratic Consequences of Incidental Exposure to Political Information: A Meta-Analysis. *Journal of Communication*, *72*(3), 345–373. https://doi.org/10.1093/joc/jqac008

Nasr, M., Rando, J., Carlini, N., Hayase, J., Jagielski, M., Cooper, A. F., Ippolito, D., Choquette-Choo, C. A., Tramèr, F., & Lee, K. (2025). *Scalable Extraction of Training Data from*

Aligned, Production Language Models. The Thirteenth International Conference on Learning Representations. https://openreview.net/forum?id=vjel3nWP2a

Newman, E. J., & Zhang, L. (2020). Truthiness: How Non-Probative Photos Shape Belief. In R. Greifeneder, M. Jaffe, E. Newman, & N. Schwarz (Eds.), *The Psychology of Fake News: Accepting, Sharing, and Correcting Misinformation* (pp. 90–114). Routledge. https://doi.org/10.4324/9780429295379-8

OpenAI. (2023, December 13). *Partnership with Axel Springer to Deepen Beneficial Use of AI in Journalism*. https://openai.com/index/axel-springer-partnership/

Peter, S., Riemer, K., & West, J. D. (2025). The Benefits and Dangers of Anthropomorphic Conversational Agents. *Proceedings of the National Academy of Sciences, 122*(22), e2415898122. https://doi.org/10.1073/pnas.2415898122

Peterson, A. J. (2025). AI and the Problem of Knowledge Collapse. *AI & Society*, *40*(5), 3249–3269. https://doi.org/10.1007/s00146-024-02173-x

Pierrès, O., Darvishy, A., & Christen, M. (2025). Exploring the Role of Generative AI in Higher Education: Semi-Structured Interviews with Students with Disabilities. *Education and Information Technologies*, *30*(7), 8923–8952. https://doi.org/10.1007/s10639-024-13134-8

Ramaul, L., Ritala, P., & Ruokonen, M. (2024). Creational and Conversational AI Affordances: How the New Breed of Chatbots Is Revolutionizing Knowledge Industries. *Business Horizons*, *67*(5), 615–627. https://doi.org/10.1016/j.bushor.2024.05.006

Reinecke, M. G., Ting, F., Savulescu, J., & Singh, I. (2025). The Double-Edged Sword of Anthropomorphism in LLMs. *Proceedings, 114*(1), Article 4. https://doi.org/10.3390/proceedings2025114004

Reiss, M. V., Knor, E. L., Merten, L., & Möller, J. (2026). Will News Find Them with Generative AI? Exploring the Affordances and Potential of GenAI for News Consumption. *Weizenbaum Journal of the Digital Society*, *6*(1). https://doi.org/10.34669/wi.wjds/6.1.1

Robins, D., & Holmes, J. (2008). Aesthetics and Credibility in Web Site Design. *Information Processing & Management, 44*(1), 386–399. https://doi.org/10.1016/j.ipm.2007.02.003

Savolainen, R. (1995). Everyday Life Information Seeking: Approaching Information Seeking in the Context of "Way of Life." *Library & Information Science Research*, *17*(3), 259–294. https://doi.org/10.1016/0740-8188(95)90048-9

Schuetzler, R. M., Grimes, G. M., & Giboney, J. S. (2020). The Impact of Chatbot Conversational Skill on Engagement and Perceived Humanness. *Journal of Management Information Systems, 37*(3), 875–900. https://doi.org/10.1080/07421222.2020.1790204

Seering, J., Luria, M., Kaufman, G., & Hammer, J. (2019). Beyond Dyadic Interactions: Considering Chatbots as Community Members. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300680

Shanahan, M. (2024). Talking about Large Language Models. *Communications of the ACM*, *67*(2), 68–79. https://doi.org/10.1145/3624724

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2024). Towards Understanding Sycophancy in Language Models. *The Twelfth International Conference on Learning Representations*.

https://proceedings.iclr.cc/paper_files/paper/2024/hash/0105f7972202c1d4fb817da9f21a9663-Abstract-Conference.html

Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–17. https://doi.org/10.1145/3613904.3642459

Simon, F., Nielsen, R. K., & Fletcher, R. (2026). *Generative AI and News Report 2026: From Standalone Tools to Information Infrastructure*. Reuters Institute for the Study of Journalism. https://doi.org/10.60625/RISJ-9XX8-0K37

Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid, A., Fisch, A., Brown, A. R., Santoro, A., Gupta, A., Garriga-Alonso, A., Kluska, A., Lewkowycz, A., Agarwal, A., Power, A., Ray, A., Warstadt, A., Kocurek, A. W., Safaya, A., Tazarv, A., … Wu, Z. (2023). Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models. *Transactions on Machine Learning Research*. https://openreview.net/forum?id=uyTL5Bvosj

Staab, R., Vero, M., Balunovic, M., & Vechev, M. (2024). *Beyond Memorization: Violating Privacy via Inference with Large Language Models*. The Twelfth International Conference on Learning Representations. https://openreview.net/forum?id=kmn0BhQk7p

Sundar, S. S., & Liao, M. (2023). Calling BS on ChatGPT: Reflections on AI as a Communication Source. *Journalism & Communication Monographs*, *25*(2), 165–180. https://doi.org/10.1177/15226379231167135

Susser, D. (2019). Invisible Influence: Artificial Intelligence and the Ethics of Adaptive Choice Architectures. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 403–408. https://doi.org/10.1145/3306618.3314286

Tafesse, W., & Mamo, Y. (2026). A Comparison of Conversational Chatbots and the Internet for Consumer Information Search. *Behaviour & Information Technology*, *45*(2), 314–331. https://doi.org/10.1080/0144929X.2025.2517215

Taylor, R. S. (1968). Question-Negotiation and Information Seeking in Libraries. *College & Research Libraries, 29*(3), 178–194. https://doi.org/10.5860/crl_29_03_178

Trippas, J. R., Spina, D., & Scholer, F. (2025). Adapting Generative Information Retrieval Systems to Users, Tasks, and Scenarios. In R. W. White & C. Shah (Eds.), *Information Access in the Era of Generative AI* (pp. 73–109). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-73147-1_4

Turpin, M., Michael, J., Perez, E., & Bowman, S. R. (2023). *Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting* (arXiv:2305.04388). arXiv. https://doi.org/10.48550/arXiv.2305.04388

van Der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud Infrastructure Dependence and the Industrialisation of Artificial Intelligence. *Big Data & Society*, *11*(1), 20539517241232630. https://doi.org/10.1177/20539517241232630

Xu, Y., Hu, L., Zhao, J., Qiu, Z., Xu, K., Ye, Y., & Gu, H. (2025). A Survey on Multilingual Large Language Models: Corpora, Alignment, and Bias. *Frontiers of Computer Science*, *19*(11), 1911362. https://doi.org/10.1007/s11704-024-40579-4

Xu, Z., Jain, S., & Kankanhalli, M. (2025). *Hallucination Is Inevitable: An Innate Limitation of Large Language Models* (arXiv:2401.11817). arXiv. https://doi.org/10.48550/arXiv.2401.11817

Zamfirescu-Pereira, J. D., Wong, R. Y., Hartmann, B., & Yang, Q. (2023). Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. *Proceedings of the*

2023 CHI Conference on Human Factors in Computing Systems, CHI '23, 1–21. https://doi.org/10.1145/3544548.3581388

Zhao, W., Ren, X., Hessel, J., Cardie, C., Choi, Y., & Deng, Y. (2024). *WildChat: 1M ChatGPT Interaction Logs in the Wild* (arXiv:2405.01470). arXiv. https://doi.org/10.48550/arXiv.2405.01470

Zhao, X., Cox, A., & Chen, X. (2025). The Use of Generative AI by Students with Disabilities in Higher Education. *The Internet and Higher Education*, *66*, 101014. https://doi.org/10.1016/j.iheduc.2025.101014