

Report on the First INformation Access in Uncertainty ScEnarios (INFUSE) Workshop at ECIR 2026

Alisa Rieger
GESIS
Cologne, Germany
alisa.rieger@gesis.org

Ran Yu
GESIS
Cologne, Germany
ran.yu@gesis.org

Amir Ebrahimi Fard
Ofcom
United Kingdom
amir.fard@ofcom.org.uk

Nicolas Mattis
University of Amsterdam
Amsterdam, The Netherlands
n.m.mattis@uva.nl

Johanne R. Trippas
RMIT University
Melbourne/Naarm, Australia
j.trippas@rmit.edu.au

Markus Bink, Stefan Dietze, David Elsweler, Liana Ermakova, Norbert Fuhr,
Luis Gomez, Vitalii Hirak, Taimoor Khan, Johannes Kiesel, Udo Kruschwitz,
Philipp Lorenz-Spreen, Oscar Ramirez Milian, Harrie Oosterhuis, and Marten
Risius*

Abstract

The First Workshop on INformation access in Uncertainty ScEnarios (INFUSE) [Rieger et al., 2026], in conjunction with the *Forty-eighth European Conference on Information Retrieval (ECIR) 2026*, focused on the complexities people face when seeking information on issues that lack a clear, definitive answer. These situations, defined as “uncertainty scenarios”, often involve scarce information, such as data voids, rapidly evolving knowledge, such as breaking news or rumours, or debated topics with multiple conflicting perspectives. A core problem is that modern information retrieval (IR) systems and large language models (LLMs) are primarily optimised to produce answers rather than communicate the limits of existing knowledge. Consequently, these systems may generate plausible-sounding responses even when reliable information is missing, thereby obscuring rather than exposing uncertainty. This dynamic introduces significant risks, including fostering a false sense of certainty, amplifying mis- and disinformation, and exposing users to viewpoint bias.

To address these challenges, this full-day, in-person event brought together a cross-disciplinary group of researchers and practitioners. The primary objectives of the workshop included:

- **Highlighting emerging risks:** Drawing attention to the potential harms caused by the rapid integration of LLM-enabled search services.

*Affiliation not shown for all authors due to space limitations (see Appendix A for details).

-
- **Building community:** Connecting professionals and researchers from various backgrounds who are focused on information access under uncertainty.
 - **Improving information systems:** Contributing to the creation of systems that promote critical thinking, awareness of complex nuances, and an informed citizenry.
 - **Fostering collaboration:** Creating a dedicated space for attendees to discover shared interests and form impactful research partnerships.

The INFUSE workshop was designed to be highly interactive, moving away from one-sided presentations. The program featured invited talks from interdisciplinary experts, participant presentations grouped by topic, “musical” roundtables for discussions, and a collaboration fair to help participants launch new research projects. The discussions and outcomes are compiled here as a collaborative report. Our aim is to share these insights with the broader IR community and help seed further dialogue on uncertainty in information access.

Date: 2 April 2026.

Website: <https://sites.google.com/view/infuse-workshop/>.

1 Introduction

The First Workshop on INformation access in Uncertainty ScEnarios (INFUSE),¹ held in conjunction with ECIR 2026,² provided a forum for the broader IR community to discuss emerging challenges, share diverse perspectives, and explore novel approaches to information access in contexts of uncertainty. To support collaborative discussion, participants were asked to submit short position statements, which were reviewed before the workshop, enabling attendees to engage deeply with one another’s perspectives.

The INFUSE Workshop brought together researchers and practitioners to examine the risks information access systems face in conditions of uncertainty, including scarce, evolving, or contested information. As LLM-enabled search services are rapidly adopted, many systems are optimised to generate answers rather than communicate what is unknown or uncertain. The workshop focused on closing this gap by supporting the development of information access systems that strengthen informed citizenship, encourage critical thinking, and better reflect the nuance of real-world information.

The workshop was motivated by the view that uncertainty is central to information seeking. Information matters because it can reduce uncertainty; if outcomes were fully known, information would have little value. However, contemporary IR research does not always represent uncertainty explicitly in its models and interfaces [Trippas and Culpepper, 2025]. This can contribute to systems that downplay uncertainty by producing fluent, plausible content even when reliable evidence is limited or unavailable.

The workshop was shaped by questions about specific uncertainty scenarios where information is scarce, evolving, or debated. Participants engaged with concerns such as: *How do we handle data voids where missing data can be easily exploited by low-quality content? How can systems provide reliable information during breaking news or when rumours spread as a collective sense-making mechanism? How do we prevent users from falling prey to confirmation bias or being*

¹<https://sites.google.com/view/infuse-workshop/home>

²<https://ecir2026.eu/>

misled by biased systems when exploring debated topics? How can we support users in their search process of search while dealing with epistemic uncertainty?

While the workshop took the stakes of public knowledge and democratic processes seriously, it maintained a constructive, cross-disciplinary tone. Through interactive “musical” tables, participants proposed approaches for aggregating and presenting evidence, managing information scarcity, and designing systems that support exploration and suspended judgement. The workshop also aimed to connect related disciplines, including communication science, to advance more robust, uncertainty-aware information access systems.

1.1 Invited Talks

Social Tipping Points: The Spiral of Attention and Platform Power

Philipp Lorenz-Spreen

This talk moves from everyday attention demands, notifications, news feeds, and social media, to the large platforms that profit from them. Over the past two decades, digital technologies have transformed public discourse: increasing interconnectedness has enabled more self-organized debate, while platforms, their algorithms, and now generative AI have gained unprecedented influence over what we see and share. I argue that power over discourse lies less in direct content control and more in the structural design of platforms. Drawing on complexity and behavioral science, the talk explores why platforms can support democratic outcomes in some contexts while fueling polarization, mistrust, and autocratic resilience in others. Finally, I discuss how we can measure these effects and what alternative platform designs might look like.

Uncertain, Incomplete and Contradictory Knowledge in IR

Norbert Fuhr

Information retrieval is about uncertainty, vagueness, and context in information access. Traditionally, IR has addressed the uncertainty inherent in document representations. Today, the knowledge described in a document itself is regarded as being potentially uncertain and incomplete. Moreover, documents vary in their credibility, and the information from different documents may be contradictory. As a logical framework for dealing with these phenomena, we propose a probabilistic four-valued logic that extends two-valued logic with the truth values unknown and inconsistent, such that the probabilities of the four truth values sum to one. Also, the framework uses an open-world assumption, treating missing facts as unknown. As an application example for aggregating incomplete and inconsistent knowledge, we present the summarization of hotel reviews. Credibility plays an important role in medical retrieval, where consideration of the level of evidence of medical studies improves retrieval quality.

2 Vision Statements

To foster an interactive environment and prioritise meaningful discussion, participants were asked to submit short abstracts and position statements rather than full papers [Azzopardi et al., 2024; Clarke et al., 2025; Azzopardi et al., 2026]. These submissions address the complex challenges of

information access under uncertainty, focusing on human engagement and the design of information systems. The collected visions encompassed a dynamic cross-section of the field, including ongoing projects, previously presented relevant work, and novel perspectives. The statements below summarise their viewpoints in their own words and are presented here in alphabetical order by first author.

Beyond Correctness: Behavioural Epistemic Learning Indicators for Search in Uncertain Information Environments

Markus Bink, Marten Risius, Udo Kruschwitz, and David Elswailer

People frequently search for information in situations where evidence is incomplete, evolving, or contested, for example, in health contexts [White and Horvitz, 2009], emerging topics [Rieger et al., 2025], or socio-political debates [Aslett et al., 2024]. In such scenarios, traditional user-oriented evaluation of search systems focuses on endpoints such as correctness, satisfaction, or interaction markers (e.g., query and click counts or depth) [Bateman et al., 2012; Elswailer et al., 2025; Bink et al., 2026a,b]. These measures say little about how users reason with uncertain information and can obscure risky behaviours, such as prematurely terminating a search upon encountering a plausible answer or uncritically accepting AI-generated summaries.

Accordingly, we introduce B-ELS (Behavioural Epistemic Learning in Search), a conceptual model that links epistemic orientations to observable interaction patterns during search. B-ELS operationalises epistemic engagement through indicators such as premature closure regulation, evidence-oriented query reformulation, stance diversity, and evidence quality weighting. The model draws on theories from information behaviour [Belkin et al., 1980; Kuhlthau, 2005; Bates, 1989], epistemic cognition [Hofer and Pintrich, 1997; Greene and Yu, 2016; Chinn et al., 2014], and credibility research [Rieh, 2002; Savolainen, 2007; Metzger et al., 2010], and is motivated by the need for evaluation approaches that capture how users engage with uncertainty rather than only what outcomes they reach.

We illustrate the potential of B-ELS through a reanalysis of a preregistered user study comparing a baseline search system with companion-system conditions that provide explicit interventions to explain the benefits of epistemically productive actions, such as clarifying information needs, formulating more effective queries to obtain higher-quality evidence, and exploring multiple results [Bink et al., 2026b]. Using B-ELS indicators, we identify distinct behavioural profiles that are not visible with traditional performance or engagement metrics. Baseline users tend to cluster into profiles characterised by low stance diversity and low-quality evidence use, patterns consistent with epistemically risky behaviour. In contrast, users in the intervention conditions exhibit broader profiles associated with greater evidence engagement and more regulated search closure. A principal component analysis of the indicators suggests separable dimensions of evidence-landscape engagement and epistemic regulation. Notably, correct and incorrect task outcomes are mixed within these profiles, highlighting the limitations of binary correctness-based evaluation under conditions of uncertainty.

We conclude by discussing how B-ELS can inform the design and evaluation of information access systems in uncertainty scenarios, including generative AI overviews, pedagogical scaffolds that treat AI errors as learning opportunities, and evaluation metrics that capture verification and evidence-seeking behaviour. Our contribution is both a theoretical framework and a set

of practical indicators for analysing search behaviour under uncertainty, aligned with workshop themes on user behaviour, ambiguity, misinformation robustness, and evaluation in uncertain information environments.

Reliable Multimodal Information Access through Training-Free Uncertainty Estimation in Cross-Modal Retrieval

Lluís Gomez, *Computer Vision Center, Universitat Autònoma de Barcelona*

Modern information access systems increasingly rely on large vision–language models (VLMs), such as CLIP [Radford et al., 2021], to retrieve relevant content across modalities. While these models achieve strong retrieval performance, they typically provide little indication of when their predictions may be unreliable. In real-world settings—such as content moderation, medical search, or decision-support tools—acting on incorrect retrieval results can have significant consequences, making *uncertainty awareness* a critical requirement for trustworthy information access.

This abstract presents a training-free framework for estimating *predictive uncertainty* in cross-modal retrieval systems built on top of pre-trained VLMs [Gomez, 2025]. Rather than modifying the underlying model or learning probabilistic embeddings [Chun et al., 2021; Upadhyay et al., 2023], the proposed approach derives practical uncertainty signals directly from retrieval behavior, enabling deployment without additional training cost or dataset-specific adaptation.

We begin with a simple yet effective baseline that interprets the cosine similarity between a query and its top-ranked result as a proxy for confidence. Building on this intuition, we introduce methods that quantify uncertainty through retrieval stability:

1. **Top-1 consistency under posterior sampling.** We measure how often the top-1 retrieved item remains the same across stochastic forward passes. Using Monte Carlo Dropout (MCD) [Gal and Ghahramani, 2016]—which approximates Bayesian inference by enabling dropout at test time to sample from the *posterior predictive distribution*—or Deep Ensembles [Lakshminarayanan et al., 2017], we obtain multiple retrieval rankings and track whether the top result is stable. Instability signals high *epistemic uncertainty*, i.e., uncertainty arising from the model’s lack of knowledge rather than from inherent data noise. In practice, we use 50 MCD samples with a dropout rate of 0.2, or an ensemble of 12 independently trained ViT-L/14 models. The computational overhead scales linearly with the number of samples but is trivially parallelizable, resulting in modest wall-clock impact.
2. **Adversarial perturbation-based robustness.** We estimate the minimal perturbation in embedding space required to flip the top-1 retrieval result, solved via Projected Gradient Descent. Confident rankings require large perturbations to change, while uncertain ones are easily disrupted. This provides a margin-based uncertainty measure analogous to decision-boundary distance in classification.

Together, these signals provide complementary views of epistemic uncertainty in ranking decisions. *Calibration*—the degree to which predicted confidence scores align with actual retrieval accuracy—is assessed via Spearman rank correlation and R^2 scores between binned uncertainty levels and Recall@1 performance.

Empirical evaluation on MSCOCO [Lin et al., 2014] and Flickr30K [Young et al., 2014] shows that the proposed methods achieve strong calibration between predicted confidence and actual

retrieval performance ($-SR^2$ scores up to 0.99), while remaining computationally efficient and outperforming learned probabilistic approaches such as ProbVLM [Upadhyay et al., 2023]. Importantly, uncertainty estimates enable effective *rejection strategies*: filtering highly uncertain queries leads to substantial improvements in retrieval reliability, as demonstrated by rejection curve analysis. This illustrates how uncertainty can support safer system behavior rather than merely post-hoc analysis.

Beyond performance gains, this work highlights a broader design perspective for information access in uncertainty scenarios. Training-free uncertainty signals can act as operational tools that allow retrieval systems to abstain, request clarification, trigger fallback mechanisms, or route decisions to humans when confidence is low. Such capabilities are increasingly important as retrieval models become core components of multimodal assistants and retrieval-augmented generation (RAG) pipelines. By connecting predictive uncertainty with actionable system strategies, this work contributes toward more transparent, robust, and dependable multimodal information environments.

Temporal Dynamics of Information and Attention on Emerging Topics

Vitalii Hirak, Alisa Rieger, Ran Yu, and Stefan Dietze

Emerging topics — novel issues gaining traction in public discourse — pose challenges for information retrieval systems due to the uncertain and evolving nature of the information surrounding them. Understanding how such topics develop over time, in terms of information availability, user attention, and semantic content, is essential for building retrieval systems that can reliably support users in acquiring information on emerging topics. This work-in-progress aims to address this challenge, contributing to the INFUSE workshop’s focus on information access in uncertainty scenarios, such as when information about a topic is scarce, incomplete, or rapidly evolving.

In this work, we investigate the temporal dynamics of information development and user attention on emerging topics using Wikipedia as a data source. Leveraging Wikipedia’s open structure, we collected time-series data on page views, page edits, and page size for approximately 300,000 articles created between January 2021 and June 2022, serving as proxies for user attention and information saturation. To support analysis across topic types, we assign each article a top-level schema.org category (Person, Place, Event, CreativeWork, Product, Organization, or Intangible), using golden annotations where available and the Qwen 2.5 large language model for classifying the remaining articles based on title and introductory sentence. Additionally, for each article, we extracted a range of descriptive statistics from the time series, including minimum, maximum, mean, median, and standard deviation of daily views, edits, and page size, as well as total views and edits and the dates on which peak values were reached.

Through unsupervised cluster analysis of these and additional statistical and temporal features, we aim to identify divergences across emerging and non-emerging topics, characterize distinct patterns of emergent topic categories, and delineate different stages of a topic’s lifecycle.

Beyond user attention and information saturation, we are interested in the semantic dimension of topic emergence. Specifically, we plan to investigate vocabulary and semantic shifts in Wikipedia articles through diachronic word embeddings, comparing the linguistic characteristics of articles at early emergence versus stabilization. This would complement the extrinsic temporal patterns with a deeper understanding of how the content and framing of emerging topics evolve over time.

Together, these lines of work contribute to a richer, multi-dimensional understanding of how emerging topics develop in online information ecosystems. We believe our findings will have direct implications for designing retrieval systems that can better navigate the uncertainty inherent in emerging and evolving information.

A Unified Methodology for Learning Uncertainty-Aware Ranking Policies from User Interactions

Oscar Ramirez Milian and Harrie Oosterhuis

Modern search and recommendation systems learn from user interactions with their rankings to optimize the predictive models underlying their ranking policies. User interactions bring the notorious challenges of bias and noise, since there are many factors that affect interactions that have to be disentangled before one can learn from them. Despite these inherent difficulties, existing methods for this task are purely frequentist and give no indication of confidence or uncertainty in their predictions. Motivated by this substantial shortcoming, we present our ongoing research project which is divided into two main stages: (i) uncertainty modelling and quantification to support deep-learning from users' interactions with ranked lists, and (ii) risk-aware reinforcement-learning of ranking policies through portfolio theory.

The first stage of our research focuses on the development of models that naturally incorporate uncertainty as an intrinsic dimension of the observed phenomenon, which in our case consists of users' clicks on ranked lists. We adopt a Bayesian viewpoint from the perspective of the parameters of click models. Rather than treating relevance and bias as classical one-single-value latent variables, we model them as distributional latent variables. Consequently, during the development and implementation of learning methods, these variables are represented as two-dimensional quantities that capture both location and dispersion. Using this framework, deep learning can be used to provide contextual predictions of relevance together with their associated uncertainty. This approach of quantifying uncertainty is commonly known in the literature as epistemic deep learning. Although this approach presents both strengths and limitations and remains an active area of research, we are integrating these developments into learning-to-rank frameworks. One may expect a degradation in predictive performance from adding uncertainty quantification to the problem. However, to our surprise, we found no such trade-off in our experiments, instead, we found an increased robustness against over-fitting, faster convergence and more accurate predictions than baselines without uncertainty quantification. These observations indicate that incorporating uncertainty quantification in predictive models can even benefit settings where uncertainty is not a desiderata in itself.

How should a model of uncertainty be exploited at decision time to yield robust outcomes in practice? Diverse strategies at the time of decision-making enhance robustness when confronted with unforeseen scenarios. This principle reflects the idea that relying on a single action in uncertain environments is inherently fragile, whereas diversification constitutes a rational strategy. This principle has been extensively studied in the field of portfolio management which provides tools that, when properly adapted, naturally yield diverse strategies. In the second stage of our research, we conceptualize information access as a decision-making process. Within this framework, we employ reinforcement learning to design ranking policies that explicitly incorporate risk and acceptability functionals defined over previously learned distributions. This stage constitutes

a natural continuation of the first; by doing so, it enables the integration of uncertainty-aware learning with policy design, thereby closing a virtuous cycle of learning and decision-making.

In summary, our ongoing research project aims to develop an uncertainty management framework for learning from user interactions with ranking applications in information retrieval and recommendation system settings, for robust and reliable ranking optimization

Epistemic Labour in Contemporary Information Access

Johanne R. Trippas and Alisa Rieger

Online information and search environments are rapidly expanding, and the cues people use to evaluate information credibility are often unreliable. Users rely on *epistemic signals* such as source attribution, provenance markers, and verification indicators to judge whether content is trustworthy. These signals are frequently missing, overlooked, or easy to fake [Lajewska et al., 2024]. As a result, users must do more than determine whether the text, video, or audio they encounter is authentic. For example, they must also assess how it is framed or from whose perspective it is presented to draw conclusions. Research on misinformation and disinformation examines false and misleading content, but it does not fully address how users navigate information environments where credibility cues are weakened. This statement introduces *epistemic labour* as a central concept for explaining how users manage ongoing uncertainty and situates it within cognitive load theory and interactive information retrieval (IIR).

From Information Scarcity to Epistemic Instability. The challenge facing contemporary users is no longer *how much* information is available, but which information can be relied upon, for what purposes, and under which constraints it should be interpreted. Information access was long framed as a scarcity problem [Fischer and Stevens, 1991]. Access was limited by availability, distribution barriers, and high time and effort costs to retrieve information. As digital systems have scaled, this framing has shifted toward information abundance, reframing challenges around overload, attention, and decision-making [Bawden and Robinson, 2009; Kashlot et al., 2026]. In offline and online contexts, people encounter information that may not reflect reality. Cosmetic, technical, and algorithmic enhancements obscure the cues that users rely on to assess credibility [Ross Arguedas et al., 2024; Metzger and Flanagin, 2013]. In online environments, particularly in generated, sponsored, or heavily edited content, signals of source, provenance, and intent are weakened. For IIR, this shift creates a fundamental challenge. Users must continually judge information, even as the usual cues for relevance, credibility and trust become unreliable. This creates constant uncertainty. Users must repeatedly recheck sources, evidence, and signals of authority, which turns information access into an epistemic trust problem and increases the cognitive demands of accessing information.

Epistemic Trust and Labour. Epistemic trust refers to the acceptance of information as accurate, credible, and relevant [Knaben et al., 2024]. Individuals use learned signals, such as stylistic conventions, institutional affiliation, or professional norms, to judge trustworthiness. These signals have never guaranteed truth, but have historically been sufficiently reliable. Information access systems implicitly depend on epistemic trust. Rankings treat relevance signals as meaningful, summaries treat synthesis as acceptable, and fluent interfaces treat coherence as a proxy

for reliability. When epistemic trust is stable, these assumptions lower cognitive effort and support sensemaking. When epistemic trust is unstable, these assumptions no longer hold. Users must question sources, verify claims and re-interpret credibility cues. This ongoing checking and recalibrating trust is epistemic labour, which increases the cognitive load of accessing information.

Beyond Misinformation and Disinformation. Research on misinformation and disinformation has identified how false and misleading claims spread online [Pathak and Spezzano, 2024]. However, this research conceptualises misinformation as inaccuracies within an information environment presumed to be largely reliable. Truth is often assumed to be stable, and disruption occurs when false content is added. The erosion of epistemic trust challenges this assumption. Users increasingly encounter environments in which *all information is potentially suspect*, regardless of accuracy. The task is no longer identifying false claims but deciding to what degree any source is dependable. Rather than asking “Is this claim true?”, users should ask “What can I justifiably take away from the information, given its framing, context, and other constraints?”. This distinction is critical. Misinformation research focuses on whether content is correct. Epistemic trust concerns whether information can be justifiably relied upon, given all its constraints. When trust deteriorates, the cognitive burden on users increases.

Cognitive Load Theory and Uncertainty Costs. Cognitive load theory [Sweller, 1994] offers a framework for analysing the effects of eroding epistemic trust. Traditionally, cognitive load is divided into intrinsic, extraneous, and germane components. Epistemic instability reshapes all three. *Intrinsic load*, the mental effort required to process the material, may increase when evaluating authenticity, plausibility, and provenance becomes part of the task rather than a background assumption. *Extraneous load*, the effort of information presentation, increases when system design promotes fluent outputs while concealing uncertainty, thereby requiring additional effort to interpret reliability. *Germane load*, the effort invested in building and refining mental models, may decline if users hesitate to integrate information they do not fully trust, which can impede learning. The resulting overload does not stem from increased informational complexity, but from the constant epistemic vigilance needed. Users allocate cognitive resources to determining whether understanding is permissible, not to understanding.

Epistemic Labour and Uncertainty in IIR. Within IIR research, uncertainty has traditionally been modelled as task-bounded and progressively reducible [Belkin et al., 1980]. Query uncertainty, topical unfamiliarity, and relevance ambiguity are assumed to diminish through interaction. Epistemic uncertainty often remains after users find information and does not disappear when they add more. Epistemic labour can be seen as managing uncertainty rather than reducing it. Users stay effective not by removing doubt, but by continually working with it. This shifts the goal of information access from correctness and efficiency toward maintaining sensemaking.

3 Musical Chairs and Breakout Group Summary

3.1 Labels and Annotations

During the workshop, we discussed the need for and shortcomings of labels (e.g., labelling each search result) in Information Retrieval and Human Information Interaction research. In the musical chair discussion, we clarified, referring to different types of labels:

- Labels required to improve IR system performance (e.g., relevance)
- Labels needed for user study data analysis (e.g., credibility, viewpoints)
- Labels that are displayed to users to support their information interactions, for instance, navigating uncertainty (e.g., information nutrition labels [Fuhr et al., 2018])

Assumptions and Biases. Search result attributes for which we often try to assign labels, such as relevance, credibility, or viewpoints, can not be captured in a fully objective manner. Instead, they are interpretations shaped by the assumptions and perceptions of the people who created the labelling schema and who labelled the items of information/ search results [Eickhoff, 2018]. In a similar vein, the same label may also be interpreted differently by users depending on their past experiences and personal characteristics [Mattis et al., 2025].

Complex Realities. Reality is often complex, with many nuances and interdependencies. However, to evaluate and optimise information access systems, we need labels to categorise information according to simplified attributes. The complexities and nuances of reality are rarely captured by such labels. For instance, labelling the stance of a search result as *supporting* a given statement does not provide any information on what arguments it contains, how they are positioned with respect to each other, and what logic of evaluation was applied [Draws et al., 2022]. This also applies to user-facing labels, which need to strike a balance between conciseness and specificity.

LLMs as Judges. We discussed how using LLMs as a judge could collapse disagreement among annotators, which might be beneficial for tasks with a clear ground truth, where disagreement stems from individual annotator biases. However, this can be detrimental when valuable information lies in the disagreement, for example, revealing the nuances and ambiguities of the information. This is the case in many labelling scenarios, prompting cautious interpretations or the need to refine the labels and annotation process.

Best Practices. When using labels for our research, we should be aware of the assumptions, norms, and power structures that shape labels and the labelling process, and that have shaped established taxonomies. If possible, we should choose labels that capture nuances and try to employ strategies to mitigate biases during the labelling process. In our work, we should acknowledge the inherent subjectivity and simplifications of labels and explicitly describe the assumptions and the context (e.g., cultural, political) in which labels were created and assigned.

Key Challenges.

- Investigating how labelling schemas and practices should be designed to reflect the nuances of real-world information, while still enabling meaningful categorisations
- Identifying strategies to assess and minimise the influence of annotator or labelling schema biases and assumptions.
- Understanding when labelling disagreement should be leveraged to reveal ambiguity, multiplicity of perspectives, or uncertainty, rather than dismissed as noise.

3.2 Data Access

Due to commercial interests and privacy concerns for user data, accessing and sharing data has been a widely recognized challenge in the research community. In this session, we discussed the challenges, EU regulations, and possible solutions for acquiring data from commercial platforms for research purposes. Specifically, the participants shared views on several key topics.

EU Digital Services Act (DSA). The participants discussed experiences related to the DSA. They highlighted different experiences in acquiring data under DSA 40/12 and 40/4. Some reported successful cases of obtaining data from Meta under DSA 40/12, while none were aware of any successful cases under DSA 40/4. We also discussed what constitutes public data and the differences between platforms, such as social media, search engines, and chat applications, in terms of the difficulty of data acquisition.

Data Sharing & Donation. Participants shared opinions, experiences, and relevant resources related to data donation, including tools and existing projects, especially for generative AI-based conversational systems such as ChatGPT. In addition, we covered data-sharing-related topics, including measures that should be taken, such as anonymization, before sharing data. Participants also discussed concerns about whether data acquired under the DSA can be shared for research purposes and how to share such data while complying with GDPR requirements, particularly regarding users' rights to trace or delete their data.

Community Efforts. Another focus was on community efforts. We discussed possibilities to avoid different researchers or research groups redundantly spending effort and resources on the same data acquisition procedures, for example, writing similar DSA data applications from scratch, by sharing documents and resources. We also discussed potential events to coordinate the sharing of resources and experiences for data acquisition.

Legal Aspects. Finally, we discussed the legal risks in different steps of data access and sharing, especially concerning specific commercial platforms. Participants noted that researchers may need to consider the resources required for legal support and funding, as well as the potential benefits of establishing collaborations with legal experts.

Key Challenges.

- Varying levels of compliance with the DSA among commercial organisations.
- Navigating legal and privacy-related risks when collecting and sharing the data.
- Reaching particular groups of participants for data donations, such as members of marginalised groups or people with low levels of institutional trust.

3.3 Users and *Good* Search Behaviour

A recurring theme in the workshop was how uncertainty shapes interaction with information access systems. We discussed assumptions these systems make about users, the behaviours they reward, and the support they offer when users are unsure how to proceed. Uncertainty was framed as a fundamental issue for information seeking. This highlighted a central design challenge, that is, many information systems are built around an idealised user, while users often approach them with incomplete knowledge, moving goals, and limited confidence in how/what to search.

A central theme in the workshop was that the idea of “good” search behaviour is misleading when it assumes that users start with clarity, precision, and confidence. In many information systems, success is tied to behaviours such as formulating accurate queries, choosing appropriate terminology, efficiently narrowing scope, and quickly judging relevance. Participants challenged these assumptions, arguing that they privilege experienced users and those familiar with a domain while creating barriers for others. They stressed that uncertainty is not a sign of poor reasoning or ineffective behaviour. Instead, it is often a normal part of exploration, learning, and problem formulation. The discussion questioned fixed standards of “good” searching and called for systems that better reflect how people actually seek information.

Task Uncertainty and Articulation. A major theme in the workshop was that information seeking is often exploratory rather than a straightforward process of direct retrieval [White and Roth, 2009]. Users often search to understand a topic, compare options, discover what is available, or refine an emerging question, rather than to retrieve a single known item. This distinction is especially important when considering the difference between factoid (typically narrow) and non-factoid (broader) information needs. The workshop discussions focused largely on these non-factoid forms of seeking, where users may not begin with a fully formed question or a clearly defined endpoint. This uncertainty takes at least two forms. The first is **uncertainty about the task** itself. Users may not know exactly what they are trying to achieve, what information would be useful, or what would count as a successful outcome [Belkin et al., 1980]. This reflects uncertainty in intention, problem framing, and task definition. The second is **uncertainty in information need articulation**. Users may struggle to find the language to express their needs, especially in unfamiliar domains or when the issue is vague, emotionally charged, or conceptually complex [Taylor, 1962; Trippas et al., 2018]. Because information access systems rely heavily on language to convert intent into retrieval, these difficulties can further restrict access to relevant information. Participants identified both forms of uncertainty as common features of information seeking and noted that many systems fail to support them because they assume clear goals and precise language from the outset. The discussion highlighted the need for systems that support browsing, comparison, discovery, and query reformulation in addition to direct retrieval. It also emphasised the value of interfaces that help users clarify goals, identify possible directions, discover

appropriate terminology, and move from open-ended concerns to more actionable queries through prompts, examples, related concepts, and alternative framings.

We further discussed task uncertainty as arising from the interaction between two dimensions of *query uncertainty* and *model uncertainty*. Query uncertainty refers to the extent to which the user can articulate the task or information need, while model uncertainty concerns how reliably the system can infer and support that need from the available query and context. When both query and model uncertainty are low, the task is relatively well specified and the system is more likely to retrieve relevant information directly. However, when at least one dimension is high, the interaction may lead to less relevant or desirable outcomes. When query uncertainty is high, even a capable model may struggle and therefore needs to support articulation through clarification, examples, related concepts, and reformulation. When model uncertainty is high, even a clear query may produce poor results. Finally, when both query and model uncertainty are high, information retrieval becomes particularly challenging, as the user is still forming the task while the system is also uncertain about how to represent it, increasing the likelihood of unproductive search trajectories. In such cases, systems should support iterative refinement by clarification, exploration, and terminology development rather than assuming that one-shot retrieval is sufficient [Trippas et al., 2025].

Confidence, Trust, and Reassurance. The workshop also highlighted the importance of confidence, trust, and reassurance in shaping user interaction with information access systems. Information access is not only a cognitive process. It is also influenced by how secure users feel in their understanding, how far they trust the system, and whether they believe they are making progress [Söllner et al., 2016]. When uncertainty is high, users may hesitate, second-guess their choices, or abandon the search altogether. Trust can be weakened when a system seems opaque, rigid, or unresponsive to the user’s situation. Confidence is also shaped through interaction. Systems that provide clear cues, transparent options, and supportive feedback can help users persist through uncertainty. This suggests that effective design should combine retrieval functions with reassurance, explanation, and visible support for incremental progress.

System Contribution to User Uncertainty. Uncertainty cannot be understood only as a user deficit, and participants emphasised that information systems can create or intensify uncertainty through confusing interfaces, hidden assumptions, unexplained categories, and inflexible interaction pathways. Uncertainty often arises from the interaction between the user, task, and system. This shifts attention away from improving user behaviour alone and towards re-examining the design choices that shape how uncertainty is experienced [White, 2016]. Within this discussion, participants also identified a tension between guidance and over-direction. Systems may need to help users clarify their needs and orient themselves, especially in exploratory contexts, but too much intervention can become restrictive. Guidance is useful when it supports orientation and progress. It becomes problematic when it narrows possibilities too early, treats ambiguity as unacceptable, or pushes users into system-preferred pathways. This is particularly important in exploratory search, where users may benefit from keeping options open rather than reaching early conclusions. The discussion suggests that effective design should provide scaffolding without too much control. That is, systems should offer suggestions, structure, and points of orientation while preserving user agency. In practice, this means supporting multiple pathways, reversible decisions, and forms of interaction that do not punish uncertainty.

Uncertainty Beyond the Problem. A key insight from the workshop was that uncertainty can also be framed as a productive condition linked to curiosity, reflection, and learning. In many situations, users need time and space to explore possibilities rather than being pushed towards completing a task. Systems designed primarily for certainty may restrict this process by requiring users to appear more precise than they are [Zarikas, 2007]. By contrast, systems that treat uncertainty as a normal and potentially valuable part of information seeking may better support interpretation, sensemaking, and discovery. This perspective broadens the role of information systems. They deliver answers and also help users understand their questions.

Key Challenges.

- Understanding how information access systems can better support users whose tasks are not yet well defined and who struggle to articulate their information needs. This includes mechanisms for task clarification, query reformulation, terminology discovery, and exploratory interaction.
- Examining how interface design influences users’ confidence and trust under conditions of uncertainty. For example, studying how feedback, transparency, and explanatory support affect users’ willingness to persist in their search task.
- Investigating how systems create or intensify uncertainty through unclear interfaces, hidden assumptions, and rigid interaction pathways. This includes understanding how to balance helpful guidance with the need to preserve user agency in exploratory search.
- Exploring how uncertainty can be treated as a productive condition that supports curiosity, sensemaking, and learning. This would broaden current approaches to system evaluation beyond retrieval efficiency alone.

4 Final Note

The *First Workshop on INformation access in Uncertainty ScEnarios (INFUSE)*, held with ECIR 2026 [Rieger et al., 2026], brought together a cross-disciplinary community to examine how information access systems should respond when reliable knowledge is limited, evolving, or contested. Across invited talks, discussions, and breakout groups, the workshop highlighted challenges in communicating uncertainty, labelling and annotation, data access, and supporting users engaged in exploratory or ill-defined search tasks. Despite the risks identified, including the risk of false certainty or bias, the discussions were marked by optimism and a commitment to building information systems that better reflect complexity, encourage critical engagement, and support informed decision-making under uncertainty.

Acknowledgments

We thank all the speakers and participants for their contributions, insights, and engagement throughout the workshop. We are also grateful to the ECIR 2026 organising committee for supporting our event and helping create a positive and memorable conference experience. This work was partially funded by the DFG, German Research Foundation (“EmergentIR”, 548295069).

A Authors and Affiliations

First author tier:

- Alisa Rieger, GESIS, Germany
- Ran Yu, GESIS, Germany.
- Amir Ebrahimi Fard, Ofcom, United Kingdom.
- Nicolas Mattis, University of Amsterdam, The Netherlands.
- Johanne R. Trippas, RMIT University, Australia.

Second author tier:

- Markus Bink (University of Regensburg, Germany).
- Stefan Dietze (GESIS - Leibniz Institute for the Social Sciences, Heinrich-Heine-University Düsseldorf, Germany).
- David Elsweiler (University of Regensburg, Germany).
- Liana Ermakova (University of Brest, France).
- Norbert Fuhr (University of Duisburg-Essen, Germany).
- Lluís Gómez (Computer Vision Center – Universitat Autònoma de Barcelona, Spain).
- Vitalii Hirak (Heinrich-Heine-University Düsseldorf, Germany).
- Taimoor Khan (GESIS - Leibniz Institute for the Social Sciences, Germany)
- Johannes Kiesel (GESIS - Leibniz Institute for the Social Sciences, Germany)
- Udo Kruschwitz (University of Regensburg, Germany).
- Philipp Lorenz-Spreen (Dresden University of Technology, Germany).
- Oscar Ramirez Milian (Radboud University, The Netherlands).
- Harrie Oosterhuis (Radboud University, The Netherlands).
- Marten Risius (University of Applied Sciences Neu-Ulm, Germany).

References

- Kevin Aslett, Zeve Sanderson, William Godel, Nathaniel Persily, Jonathan Nagler, and Joshua A Tucker. Online searches to evaluate misinformation can increase its perceived veracity. *Nature*, 625(7995):548–556, 2024.
- Leif Azzopardi, Charles LA Clarke, Paul Kantor, Bhaskar Mitra, Johanne R Trippas, Zhaochun Ren, Mohammad Aliannejadi, Negar Arabzadeh, Raman Chandrasekar, Maarten de Rijke, et al. Report on The Search Futures Workshop at ECIR 2024. *ACM SIGIR Forum*, 58(1):1–41, 2024.
- Leif Azzopardi, Charles L. A. Clarke, Claudia Hauff, Yubin Kim, Adam Roegiest, Johanne R. Trippas, Zhaochun Ren, and Saber Zerhoubi. Report on The Search Futures Workshop at ECIR 2026. *ACM SIGIR Forum*, 60(1):1–28, 2026.
- Scott Bateman, Jaime Teevan, and Ryen W White. The search dashboard: how reflection and comparison impact search behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1785–1794, 2012.

-
- Marcia J Bates. The design of browsing and berrypicking techniques for the online search interface. *Online review*, 13(5):407–424, 1989.
- David Bawden and Lyn Robinson. The dark side of information: overload, anxiety and other paradoxes and pathologies. *Journal of information science*, 35(2):180–191, 2009.
- Nicholas J Belkin et al. Anomalous states of knowledge as a basis for information retrieval. *Canadian journal of information science*, 5(1):133–143, 1980.
- Markus Bink, Marten Risius, David Elsweiler, and Udo Kruschwitz. ”Can You Tell Me?”: Designing Copilots to Support Human Judgement in Online Information Seeking. CHIIR ’26, page 172–182, New York, NY, USA, 2026a. Association for Computing Machinery. ISBN 9798400724145. doi: 10.1145/3786304.3787866. URL <https://doi.org/10.1145/3786304.3787866>.
- Markus Bink, Marten Risius, Udo Kruschwitz, and David Elsweiler. Seek and You Shall Find: Design & Evaluation of a Context-Aware Interactive Search Companion. CHIIR ’26, page 361–365, New York, NY, USA, 2026b. Association for Computing Machinery. ISBN 9798400724145. doi: 10.1145/3786304.3787914. URL <https://doi.org/10.1145/3786304.3787914>.
- Clark A Chinn, Ronald W Rinehart, and Luke A Buckland. Epistemic cognition and evaluating information: Applying the air model of epistemic cognition. *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences*, pages 425–453, 2014.
- Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio de Rezende, Yannis Kalantidis, and Diane Larlus. Probabilistic embeddings for cross-modal retrieval. In *arXiv preprint arXiv:2101.05068*, 2021.
- Charles Clarke, Paul Kantor, Adam Roegiest, Ian Soboroff, Johanne Trippas, and Zhaochun Ren. The second search futures workshop at ecir’25. In *European Conference on Information Retrieval*, pages 313–318. Springer, 2025.
- Tim Draws, Oana Inel, Nava Tintarev, Christian Baden, and Benjamin Timmermans. Comprehensive Viewpoint Representations for a Deeper Understanding of User Interactions With Debated Topics. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*, CHIIR ’22, pages 135–145, New York, NY, USA, March 2022. Association for Computing Machinery. ISBN 978-1-4503-9186-3. doi: 10.1145/3498366.3505812.
- Carsten Eickhoff. Cognitive Biases in Crowdsourcing. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, WSDM ’18, pages 162–170, New York, NY, USA, February 2018. Association for Computing Machinery. ISBN 978-1-4503-5581-0. doi: 10.1145/3159652.3159654.
- David Elsweiler, Samy Ateia, Markus Bink, Gregor Donabauer, Marcos Fernández Pichel, Alexander Frummet, Udo Kruschwitz, David E. Losada, Bernd Ludwig, Selina Meyer, and Noel Pascual Presa. Query Smarter, Trust Better? Exploring Search Behaviours for Verifying News

-
- Accuracy. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '25, page 515–526, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400715921. doi: 10.1145/3726302.3730067. URL <https://doi.org/10.1145/3726302.3730067>.
- Gerhard Fischer and Curt Stevens. Information access in complex, poorly structured information spaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 63–70, 1991.
- Norbert Fuhr, Anastasia Giachanou, Gregory Grefenstette, Iryna Gurevych, Andreas Hanselowski, Kalervo Jarvelin, Rosie Jones, YiquN Liu, Josiane Mothe, Wolfgang Nejdl, et al. An information nutritional label for online documents. In *ACM SIGIR Forum*, volume 51, pages 46–66. ACM New York, NY, USA, 2018.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning (ICML)*, 2016.
- Lluís Gomez. Over the top-1: Uncertainty-aware cross-modal retrieval with clip. In *Proceedings of the 41st Conference on Uncertainty in Artificial Intelligence (UAI)*, 2025.
- Jeffrey A Greene and Seung B Yu. Educating critical thinkers: The role of epistemic cognition. *Policy Insights from the Behavioral and Brain Sciences*, 3(1):45–53, 2016.
- Barbara K Hofer and Paul R Pintrich. The development of epistemological theories: Beliefs about knowledge and knowing and their relation to learning. *Review of educational research*, 67(1): 88–140, 1997.
- Salah Mohammed Kashlot, Ian Ruthven, and Yashar Moshfeghi. Strategies for managing information overload: a systematic review. *Journal of Documentation*, 82(2):249–265, 2026.
- Saskia Knapen, Wilma E Swildens, Wendy Mensink, Adriaan Hoogendoorn, Joost Hutsebaut, and Aartjan TF Beekman. The development and psychometric evaluation of the questionnaire epistemic trust (qet): A self-report assessment of epistemic trust. *Clinical Psychology & Psychotherapy*, 31(1):e2930, 2024.
- Carol Collier Kuhlthau. Information search process. *Hong Kong, China*, 7(2005):226, 2005.
- Weronika Łajewska, Krisztian Balog, Damiano Spina, and Johanne Trippas. Can users detect biases or factual errors in generated responses in conversational information-seeking? In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 92–102, 2024.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014.

-
- Nicolas Mattis, Tim Groot Kormelink, Philipp K Masur, Judith Moeller, and Wouter van Attevelde. Nudging news readers: A mixed-methods approach to understanding when and how interface nudges affect news selection. *Digital Journalism*, 13(9):1647–1667, 2025.
- Miriam J Metzger and Andrew J Flanagin. Credibility and trust of information in online environments: The use of cognitive heuristics. *Journal of pragmatics*, 59:210–220, 2013.
- Miriam J Metzger, Andrew J Flanagin, and Ryan B Medders. Social and heuristic approaches to credibility evaluation online. *Journal of communication*, 60(3):413–439, 2010.
- Royal Pathak and Francesca Spezzano. An empirical analysis of intervention strategies’ effectiveness for countering misinformation amplification by recommendation algorithms. In *European Conference on Information Retrieval*, pages 285–301. Springer, 2024.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763. PMLR, 2021.
- Alisa Rieger, Stefan Dietze, and Ran Yu. A research vision for web search on emerging topics. *arXiv preprint arXiv:2509.10212*, 2025.
- Alisa Rieger, Ran Yu, Amir Ebrahimi Fard, Nicolas Mattis, and Johanne R. Trippas. INFUSE Workshop: The First Workshop on INformation access in Uncertainty ScEnarios. In *Proceedings of the European Conference on Information Retrieval (ECIR’26)*, ECIR ’26, pages 1–9, 2026.
- Soo Young Rieh. Judgment of information quality and cognitive authority in the web. *Journal of the American society for information science and technology*, 53(2):145–161, 2002.
- Amy A Ross Arguedas, Sumitra Badrinathan, Camila Mont’Alverne, Benjamin Toff, Richard Fletcher, and Rasmus Kleis Nielsen. Shortcuts to trust: Relying on cues to judge online news from unfamiliar sources on digital platforms. *Journalism*, 25(6):1207–1229, 2024.
- Reijo Savolainen. Media credibility and cognitive authority. the case of seeking orienting information. *Information Research: An International Electronic Journal*, 12(3):n3, 2007.
- Matthias Söllner, Axel Hoffmann, and Jan Marco Leimeister. Why different trust relationships matter for information systems users. *European Journal of Information Systems*, 25(3):274–287, 2016.
- John Sweller. Cognitive load theory, learning difficulty, and instructional design. *Learning and instruction*, 4(4):295–312, 1994.
- Robert S Taylor. The process of asking questions. *American documentation*, 13(4):391–396, 1962.
- Johanne R. Trippas and J. Shane Culpepper. Report from the Fourth Strategic Workshop on Information Retrieval in Lorne (SWIRL 2025). *SIGIR Forum*, 59(1), 2025.

-
- Johanne R Trippas, Damiano Spina, Lawrence Cavedon, Hideo Joho, and Mark Sanderson. Informing the design of spoken conversational search: Perspective paper. In *Proceedings of the 2018 conference on human information interaction & retrieval*, pages 32–41, 2018.
- Johanne R Trippas, Damiano Spina, and Falk Scholer. Adapting generative information retrieval systems to users, tasks, and scenarios. In *Information access in the era of generative AI*, pages 73–109. Springer, 2025.
- Uddeshya Upadhyay, Shyamgopal Karthik, Massimiliano Mancini, and Zeynep Akata. Probylm: Probabilistic adapter for frozen vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1899–1910, 2023.
- Ryen W White. *Interactions with search systems*. Cambridge University Press, 2016.
- Ryen W White and Eric Horvitz. Cyberchondria: studies of the escalation of medical concerns in web search. *ACM Transactions on Information Systems (TOIS)*, 27(4):1–37, 2009.
- Ryen W White and Resa A Roth. *Exploratory search: Beyond the query-response paradigm*. Number 3. Morgan & Claypool Publishers, 2009.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014.
- Vasilios Zarikas. Modeling decisions under uncertainty in adaptive user interfaces. *Universal Access in the Information Society*, 6(1):87–101, 2007.