

Cross-Domain User Intent Modeling for Intelligent Conversational Search Systems

Manoj Dishra

Department of Computer Science, University of Houston, Houston, TX, USA.

manoj1985@uh.edu

Kanneth Rayes

Department of Computer Science, University of New Hampshire, Durham, NH, USA.

kennethhayes@unh.edu

Abstract

Intelligent conversational search systems must interpret user intent across multiple interaction turns and across heterogeneous application domains. This paper develops a system-level analysis of cross-domain user intent modeling for such systems, focusing on structural trade-offs rather than on a single algorithmic solution. The discussion begins with the conceptual foundations of user intent in conversational and contextual information retrieval, emphasizing that intent is dynamically negotiated through dialogue and is shaped by domain-specific constraints. It then examines architectural strategies for shared representation learning, domain adaptation, and conversational state management. A central concern is the tension between model sharing and domain isolation, because shared encoders improve cross-domain transfer while domain-specific adapters improve accountability and maintainability. The paper further discusses personalization and context-aware retrieval in local categorical query settings, drawing attention to the importance of context signals and dialogue history. Robustness, fairness, and governance are analyzed as first-order design constraints rather than post-deployment afterthoughts. The paper considers bias amplification, transparency, auditability, and the risks of large pretrained models. Deployment and operational sustainability are addressed through modular architectures, learning-to-rank components, and continuous evaluation. Finally, the policy implications of cross-domain intent modeling are discussed, including data governance, disparate impact, and the need for accountability frameworks. The paper argues that sustainable conversational search systems require coordinated attention to representation, infrastructure, fairness, deployment, and governance rather than isolated optimization of predictive accuracy.

Keywords

conversational search; cross-domain intent modeling; neural information retrieval; personalization; context-aware retrieval; fairness; system governance.

1. Introduction

The design of intelligent conversational search systems has become a central challenge at the intersection of information retrieval, natural language processing, and human-computer interaction. Unlike traditional query-response search engines, conversational systems must interpret mixed-initiative exchanges, maintain state across turns, and reconcile ambiguous expressions with evolving user goals [1]. The challenge is further complicated when user intent spans multiple domains, such as when a traveler asks about flights, hotel availability, restaurant reservations, and local weather within a single dialogue. Cross-domain user intent

modeling addresses this complexity by representing, transferring, and refining intent signals beyond isolated domain boundaries. Neural approaches to conversational AI have substantially improved the capacity of such systems to learn rich representations of queries and responses [2]. At the same time, transformer-based attention mechanisms and pretrained language models have supplied general-purpose semantic encoders that can be adapted to domain-specific retrieval tasks [3], [4].

The shift from single-turn retrieval to multi-turn conversational interaction changes the nature of intent modeling. In conversational search, intent is not a fixed query property but a dynamically negotiated construct that evolves through clarification, reformulation, and feedback [1]. Early work on user intent prediction in information-seeking conversations demonstrated that dialogue context, prior utterances, and response structure can be used to anticipate the type of information need under negotiation [5]. Spoken conversational search further complicates intent inference because voice interfaces introduce disfluency, turn-taking uncertainty, and prosodic signals that are absent from typed interaction [6]. Relevance, a foundational concept in information retrieval, has also been reconceptualized as a multidimensional and situational phenomenon rather than a simple topical match [7]. The contextual turn in information science emphasizes that information seeking and retrieval must be integrated with the user's work task, domain knowledge, and situational constraints [8].

Cross-domain intent modeling therefore requires architectural and governance decisions that go far beyond the selection of a particular neural model. System designers must decide how to encode domain-specific knowledge, how to share parameters across domains, how to manage privacy and data provenance, and how to ensure that intent predictions do not amplify harmful stereotypes or structural inequalities. The present paper develops a system-level analysis of cross-domain user intent modeling for conversational search. It does not propose a single algorithm, but instead examines the structural trade-offs that arise in representation, personalization, retrieval, deployment, and governance. The discussion draws on neural information retrieval, interactive information retrieval, fairness literature, and operational perspectives to identify enduring tensions in building such systems. By situating the technical design space within broader socio-technical infrastructures, the paper offers a forward-looking perspective on sustainable and responsible conversational search.

2. Conceptual Foundations of Cross-Domain User Intent

User intent in search has historically been modeled through query classification, click behavior, and session context. In conventional web search, intent categories such as informational, navigational, and transactional provide a coarse but useful abstraction for ranking and evaluation. However, conversational search systems require finer-grained intent representations because an utterance may perform multiple functions simultaneously. A user statement can assert a constraint, request a recommendation, clarify a previous turn, or change the topic. Cross-domain modeling introduces additional heterogeneity because the same surface utterance can signal different intents depending on whether the current domain is retail, travel, health, or local services. Domain-specific user interest models, such as deep interest networks for click-through prediction, illustrate the value of learning rich behavioral representations from high-volume interaction logs [9].

The conceptual foundation of cross-domain intent modeling lies in separating domain-invariant intent structure from domain-specific expressions of that structure. Domain-invariant elements include information need type, urgency, specificity, and the user's willingness to accept recommendations. Domain-specific elements include entities, attributes,

constraints, and service actions. A robust cross-domain system can learn shared representations for intent dynamics while allowing domain adapters to capture vocabulary and task constraints. This separation is not merely an engineering convenience; it reflects the observation that information seeking is a general human activity that is realized differently across domains [8]. Search engines have long combined general retrieval algorithms with domain-specific indexing and ranking features [10]. Neural ranking models extend this logic by learning semantic matches between queries and documents across heterogeneous collections [11], [12].

Cross-domain intent modeling also requires a clear account of transferability. Not all intent signals are equally transferable. A user's preference for immediate answers may generalize across domains, while a user's preference for luxury hotels may not transfer to restaurant selection without careful abstraction. Pretrained language models provide a basis for transfer because they encode broad linguistic regularities that can be fine-tuned with domain data [4]. Continued pretraining frameworks further illustrate how general-purpose representations can be adapted to multiple domains without catastrophic forgetting [13]. Sentence encoding methods for natural language inference demonstrate that task-specific modifications can preserve general semantic capabilities while improving performance on targeted objectives [14]. These technical findings support a layered architecture in which general intent representation, domain adaptation, and task-specific decoding are treated as distinct but jointly optimized components.

3. Architectural Considerations for Conversational Search Systems

Architectural choices shape how cross-domain intent models are built, evaluated, and maintained. A conversational search system typically includes automatic speech recognition or text input processing, dialogue state tracking, intent recognition, query rewriting, retrieval, ranking, response generation, and personalization modules. Cross-domain capability can be introduced at different layers. Shared intent encoders may be used across domains, domain-specific projection layers may map general intent representations into task spaces, and retrieval indexes may remain isolated to respect data governance boundaries. The structural trade-off is between shared capacity and domain isolation. Shared capacity improves sample efficiency and enables zero-shot transfer, but it also increases the risk of interference when domains have conflicting feature distributions. Domain isolation improves accountability and maintainability, but it may reduce the system's ability to generalize rare intents.

The conversational nature of the system introduces stateful processing requirements that are less prominent in single-turn search. Dialogue state tracking must maintain a belief over the user's goal, constraints, and preferences across turns. Intent recognition is therefore not a one-shot classification problem but a recurrent inference process over an evolving state [1], [5]. Transformer architectures allow the system to condition on long dialogue histories and attend to relevant previous turns [3]. However, long context modeling is computationally expensive, and deployed systems often impose context windows that can truncate important user constraints. The architecture must balance representational richness against latency, memory, and cost. This is particularly important in spoken conversational search, where users expect low-latency natural interaction [6].

A further architectural concern is the distinction between generative and discriminative components. Large language models can generate natural language responses that blend retrieved information with inferred intent, but they are not always reliable as retrieval engines. Retrieval-augmented architectures separate retrieval from generation, allowing the model to

condition on explicit evidence rather than relying on parametric memory alone. This separation is useful for cross-domain systems because retrieval corpora can be updated independently of inference models. Ranking components can also be modularized so that different domains maintain different ranking models while sharing a common semantic encoder [11], [12]. Such modularity supports continuous deployment, targeted evaluation, and rollback without retraining the entire system. It also creates a natural boundary where domain-specific business rules and policy constraints can be enforced.

4. Intent Representation and Cross-Domain Transfer

Representing user intent in a computational form requires a set of modeling choices about granularity, temporality, and domain specificity. Early click models used click-through history to infer relevance and user preferences, treating observed interaction logs as a supervisory signal [15]. Such models are effective in high-volume domains where interaction data are abundant, but they are less suitable for sparse conversational settings. Cross-domain transfer can mitigate data sparsity by pooling similar intent types across domains. However, pooling must be guided by a clear notion of similarity. For example, a purchase intent in e-commerce and a booking intent in travel share a transactional structure but differ in attributes, constraints, and risk. A well-designed cross-domain intent model learns to represent the transactional structure as a shared latent variable while preserving domain-specific attribute spaces.

The temporal dimension of intent is particularly important in conversational search. User intent can shift rapidly within a session, and the system must distinguish between a genuine shift and an anaphoric continuation. Dialogue history encoding methods based on bidirectional transformers [4] or enhanced sequential encoders [14] can capture local dependencies, but they require careful handling of long-range dependencies. The model must also account for the fact that users may express partial preferences before committing to a final decision. Exploratory search models have long recognized that information seeking is often iterative, uncertain, and influenced by what the user encounters during the search process [16]. Cross-domain intent models should therefore not treat intent as a fixed label to be predicted once; they should model intent as a probability distribution over latent goals that is refined as new evidence arrives.

Another representational issue is the degree of interpretability. Black-box neural intent models may achieve high predictive accuracy while offering little insight into why a particular clarification question was selected or why a result was retrieved. Interpretable representations are valuable for debugging, fairness auditing, and user trust. Hybrid architectures can pair neural encoders with structured domain ontologies or attribute-value representations to provide a symbolic anchor for intent decisions. This approach is particularly relevant in cross-domain settings because domain experts can inspect whether the shared intent representation has captured the correct abstraction. It also supports policy-level governance by enabling auditors to trace how a user's demographic or behavioral attributes influence intent predictions.

Cross-domain transfer can be organized along a spectrum from fully shared to fully isolated models. Fully shared models maximize parameter efficiency and can transfer common conversational patterns, but they may be vulnerable to negative transfer when domains have strongly divergent feature distributions. Fully isolated models allow precise optimization but cannot generalize to emerging domains or low-resource categories. Intermediate architectures, such as shared encoders with domain-specific adapters, offer a pragmatic compromise. The

same transformer backbone can encode user utterances and dialogue context, while lightweight adapters transform the shared representation into domain-specific intent scores. This modular design also aligns with operational governance because domain owners can update adapters without destabilizing the shared backbone. Continued pretraining and adapter-based fine-tuning strategies have shown that such architectural separation can be implemented without severe performance loss [13].

5. Personalization, Context, and Retrieval

Personalization adds another layer of complexity to cross-domain intent modeling. A conversational search system must infer not only what the user wants but also which aspects of the user's history are relevant to the current request. Historical interaction logs can be powerful signals for preference learning, but they are also noisy and may encode stale preferences. Human-AI interaction guidelines have emphasized that systems should make clear when personalization is being used, allow user control, and avoid overreliance on inferred preferences [17]. In cross-domain settings, these requirements become more challenging because a preference learned in one domain may be misapplied in another. For instance, a user's frequent purchases in a high-price category should not automatically imply a preference for high-price options in unrelated service domains. Personalization models must therefore include domain-aware gating mechanisms that determine which parts of the user profile are transferable.

Context-aware retrieval is especially important for local categorical queries, where the user's intent often depends on location, time, device, and immediate situational constraints. Recent work on AI-driven personalization across domains for local categorical query understanding has highlighted the importance of context-aware retrieval in local service settings [18]. Local intents such as nearby dining, entertainment, or urgent services cannot be resolved by lexical matching alone. The system must infer whether the user is seeking a recommendation, comparing options, or trying to complete a transaction. Cross-domain personalization in this context involves learning shared patterns about urgency and proximity while preserving category-specific constraints. This line of work reinforces the broader observation that context is not a fixed input but a dynamic interpretation process involving the user, the platform, and the service environment.

Retrieval in conversational search must also support incremental and interactive refinement. Users often begin with vague expressions and then provide additional constraints after observing initial results. The system should therefore combine retrieval and clarification in a way that minimizes user burden while maximizing information gain. This requires modeling the expected value of different clarification questions. A cross-domain intent model can inform this process by estimating how much a user's intent distribution would shift under possible answers. In high-risk domains such as health or legal search, clarification becomes particularly important because misunderstanding the intent can have serious consequences. Personalization must be handled cautiously in such settings, and the system should be transparent about the evidence used to refine its interpretation of the user's need.

6. Robustness, Fairness, and Governance

Fairness and robustness are integral to the design of cross-domain intent models. Survey work on bias and fairness in machine learning has identified multiple sources of bias, including historical bias, representation bias, and measurement bias [19]. Conversational search systems are exposed to all of these because they learn from interaction logs, language models, and

user feedback signals that reflect existing social inequalities. A cross-domain intent model may learn spurious associations between certain user attributes and intent categories, especially when those attributes are unevenly distributed across domains. Without careful intervention, the system may systematically misunderstand the intents of underrepresented users or route them toward lower-quality results. Fairness therefore cannot be reduced to optimizing a single aggregate metric; it requires a careful examination of downstream consequences for different user groups.

Large pretrained language models have amplified both the capabilities and the risks of conversational search. Critical analyses have shown that such models can reproduce problematic social biases and may create the illusion of understanding even when their outputs are statistically derived from large corpora [20]. In cross-domain systems, the use of shared pretrained encoders means that biases learned from broad web data can propagate into domain-specific intent predictions. This creates a governance challenge because domain owners may not be aware of the full provenance of the shared representation. Auditing mechanisms must therefore span both the shared backbone and the domain-specific adapters. In addition, fairness issues in online advertising and delivery systems have demonstrated that seemingly neutral optimization objectives can lead to discriminatory outcomes when user populations are segmented by sensitive attributes [21]. Cross-domain intent models that personalize ads, recommendations, or local services must be evaluated for such disparate effects.

Governance frameworks for algorithmic accountability can help manage these risks. End-to-end auditing approaches emphasize the importance of documenting model purpose, data sources, design choices, and deployment context, as well as creating channels for external review [22]. In conversational search, such documentation should include the domains covered by the system, the intent taxonomy used, the personalization signals incorporated, and the procedures for handling sensitive queries. Foundation model analyses further suggest that the broad deployment of general-purpose model families requires coordinated evaluation across tasks, domains, and societal impact dimensions [23]. Cross-domain intent models occupy an intermediate position between domain-specific systems and general-purpose assistants. They inherit many of the risks of foundation models while also facing domain-specific compliance obligations.

Robustness is closely related to fairness. A robust cross-domain intent model should maintain reasonable performance when the input distribution shifts, when dialogue patterns change, or when new domains are introduced. Domain adaptation can fail silently if the shared representation overfits to high-resource domains. Evaluation should therefore include stress testing across demographic groups, linguistic registers, and domain shifts. This requires infrastructure for dataset curation, annotation, and longitudinal tracking of model behavior. Governance structures should assign responsibility for these evaluations rather than assuming that model developers will perform them voluntarily. In many cases, external audits and regulatory oversight can provide incentives for more careful system design.

7. Deployment and Operational Sustainability

Deployment of cross-domain conversational search systems requires attention to latency, scalability, data freshness, and system drift. Modular architectures help operational teams make targeted updates to domain adapters, retrieval indexes, and ranking components without retraining the entire pipeline. Learning-to-rank techniques provide a mature operational framework for combining multiple relevance signals in a reproducible manner [24]. In cross-

domain systems, learning-to-rank models can be trained separately for different domains while still consuming shared intent representations. This allows domain-specific business rules to be encoded in the ranking layer without altering the shared semantic encoder. It also facilitates staged rollouts and rollbacks when a domain model degrades.

Data governance is a major operational concern. Cross-domain intent modeling often requires combining data from multiple product areas, each with its own retention policies, privacy requirements, and consent flows. The system architecture must support fine-grained data provenance so that a user's data from one domain is not used in another domain without appropriate authorization. Privacy-aware personalization can be implemented through cryptographic techniques, federated learning, or on-device inference, but these approaches introduce additional latency and complexity. The operational challenge is to balance personalization quality with privacy and regulatory compliance. A shared representation learned from centralized data may improve accuracy, but it also increases the scope of potential harm if a breach occurs. System designers should therefore consider data minimization and purpose limitation as architectural constraints rather than legal afterthoughts.

Sustainability also includes model maintenance. Conversational search systems must be updated as language evolves, new entities appear, and user behavior shifts. Continuous evaluation requires a robust telemetry infrastructure that can detect intent misclassification, retrieval failures, and emerging fairness issues. However, telemetry itself carries privacy risks. Metrics should be designed to capture aggregate system behavior without creating fine-grained behavioral profiles that could be misused. Cross-domain systems face the additional problem that a change in one domain adapter may affect upstream shared components and downstream ranking behavior. Versioned model registries, staged deployment pipelines, and automated regression testing can reduce this coupling risk. Long-term sustainability is therefore not only a technical goal but also an organizational capability.

8. Policy Implications and Socio-Technical Considerations

Cross-domain user intent modeling raises policy questions that extend beyond technical performance. One central issue is the potential for disparate impact when automated systems make consequential decisions about access to information, services, or economic opportunities. Legal scholarship has argued that data-driven decision-making can produce discriminatory effects even when protected attributes are not explicitly used [25]. In conversational search, discriminatory effects may emerge when intent models systematically misinterpret dialects, fail to recognize nonstandard phrasings, or rely on proxies for socioeconomic status. Policy responses must therefore address both intentional discrimination and structural inequality embedded in training data and evaluation practices.

Transparency is another policy concern. Users may not know when a conversational search system is personalizing results, inferring sensitive attributes, or using cross-domain behavioral history. Meaningful transparency requires more than a privacy policy written in general terms. It should include accessible explanations of how the system interprets intent, why certain results are returned, and how users can contest or correct the system's inferences. Cross-domain systems make this transparency more difficult because an inference made in one domain may affect results in another domain without the user's awareness. Policy frameworks should require clear disclosure of the cross-domain signals and inferred attributes that shape ranking, recommendation, and response generation. Without such disclosure, users cannot meaningfully assess why they are receiving particular results or exercise effective control over the inferences made about them. Cross-domain systems create a distinctive form

of information asymmetry because the user may be aware of only one domain context while the system draws on a profile assembled from many domains. A shopping interaction may be influenced by earlier health-related queries, a local search may be informed by prior travel behavior, and a financial assistant may incorporate signals from social media activity. These secondary uses are not inherently illegitimate, but they require explicit governance because the user is often unable to anticipate or contest them.

A further socio-technical concern is that cross-domain intent modeling is embedded in broader platform infrastructures that are not neutral with respect to social power. Search systems shape how people discover information, access services, and form beliefs. When intent inference is optimized primarily for engagement or commercial conversion, the resulting system may systematically favor certain interpretations of user need while marginalizing others. This is especially consequential for local categorical search, where small differences in intent interpretation can determine whether a user is exposed to one business rather than another, one healthcare option rather than another, or one public service rather than another. The governance of cross-domain intent models must therefore include mechanisms for community input, independent evaluation, and transparent reporting about how the system performs across different user populations and domain contexts.

Accountability frameworks should also address the distributed nature of intent inference. In a modular cross-domain system, a single user intent may be affected by a general language encoder, one or more domain adapters, a personalization layer, and a ranking component. When a harmful inference occurs, it can be difficult to attribute responsibility to any single module. End-to-end auditing can help by requiring documentation of each component, its training data, its intended behavior, and its known limitations. However, auditability alone is insufficient if the system evolves rapidly through continuous learning. Policy mechanisms must therefore require stability in high-risk domains, mandatory review before major model changes, and the preservation of audit trails over meaningful periods. These requirements can be supported by infrastructure such as model registries, dataset cards, and behavior logs, but they also depend on organizational incentives that treat accountability as a core objective rather than a compliance burden.

Finally, cross-domain intent modeling raises important questions about user autonomy and the legitimacy of inferred knowledge. Users often reveal information indirectly through their interactions, and they may not expect that information to be abstracted into an intent profile that travels across domains. Respecting user autonomy requires that systems offer meaningful choices about cross-domain personalization, including the ability to limit transfer between sensitive domains. It also requires that defaults be designed to avoid surprising uses. A health domain intent model should not, by default, incorporate behavioral history from unrelated commercial domains unless the user has been clearly informed and has consented. These design choices have policy implications because they shape whether users can effectively govern their own informational boundaries. Without such safeguards, cross-domain intent modeling may improve short-term relevance at the cost of long-term trust and social legitimacy.

9. Conclusion

This paper has examined cross-domain user intent modeling for intelligent conversational search systems from a system-level and socio-technical perspective. The central argument is that cross-domain intent modeling is not simply a predictive modeling task but a complex design space involving representation, architecture, personalization, retrieval, fairness,

deployment, and governance. The paper has shown that user intent in conversational search is dynamically negotiated and cannot be reduced to a fixed label. It has examined architectural trade-offs between shared representation and domain isolation, emphasizing the value of modular designs that combine general intent encoders with domain-specific adapters. It has discussed how personalization and context-aware retrieval are especially important for local categorical queries, where situational factors and user history interact in subtle ways. It has also argued that robustness, fairness, and governance must be treated as first-order design constraints rather than as post-deployment corrections. Cross-domain systems pose distinctive risks because an inference made in one domain may propagate into another, creating forms of informational influence that are opaque to users and difficult to audit. Sustainable conversational search systems therefore require coordinated attention to model architecture, data governance, operational maintenance, and accountability frameworks. The future development of cross-domain intent modeling should not focus solely on improving predictive accuracy. It should also invest in interpretable representations, domain-aware personalization, fairness evaluation, transparent deployment infrastructure, and policy mechanisms that protect user autonomy. Only through such a comprehensive approach can intelligent conversational search systems fulfill their promise without undermining the trust and equity that legitimate information access requires.

References

1. Radlinski, F., & Craswell, N. (2017). A theoretical framework for conversational search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval* (pp. 117-126). ACM.
2. Gao, J., Galley, M., & Li, L. (2019). Neural approaches to conversational AI: Question answering, task-oriented dialogues and social chatbots. *Foundations and Trends in Information Retrieval*, 13(2-3), 127-298.
3. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* 30 (pp. 5998-6008).
4. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 4171-4186). Association for Computational Linguistics.
5. Qu, C., Yang, L., Croft, W. B., Zhang, Y., Trippas, J. R., & Qiu, M. (2019). User intent prediction in information-seeking conversations. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval* (pp. 25-33). ACM.
6. Trippas, J. R., Spina, D., Cavedon, L., Joho, H., & Sanderson, M. (2018). Informing the design of spoken conversational search: Perspective paper. In *Proceedings of the 2018 Conference on Human Information Interaction and Retrieval* (pp. 32-41). ACM.
7. Saracevic, T. (2007). Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II: Nature and manifestations of relevance. *Journal of the American Society for Information Science and Technology*, 58(13), 1915-1933.

8. Ingwersen, P., & Järvelin, K. (2005). *The turn: Integration of information seeking and retrieval in context*. Springer.
9. Zhou, G., Zhu, X., Song, C., Fan, Y., Zhu, H., Ma, X., Yan, Y., Jin, J., Li, H., & Gai, K. (2018). Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 1059-1068). ACM.
10. Baeza-Yates, R., & Ribeiro-Neto, B. (2011). *Modern information retrieval: The concepts and technology behind search* (2nd ed.). Addison-Wesley.
11. Mitra, B., & Craswell, N. (2018). An introduction to neural information retrieval. *Foundations and Trends in Information Retrieval*, 13(1), 1-126.
12. Karpukhin, V., Oğuz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., & Yih, W. (2020). Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (pp. 6769-6781). Association for Computational Linguistics.
13. Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 8342-8360). Association for Computational Linguistics.
14. Conneau, A., Kiela, D., Schwenk, H., Barrault, L., & Bordes, A. (2017). Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 670-680). Association for Computational Linguistics.
15. Chuklin, A., Markov, I., & de Rijke, M. (2015). *Click models for web search*. Morgan & Claypool.
16. White, R. W., & Roth, R. A. (2009). *Exploratory search: Beyond the query-response paradigm*. Morgan & Claypool.
17. Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Paper No. 3, pp. 1-13). ACM.
18. Yu, X. (2026, January). AI-Driven Personalization across Domains for Local Categorical Query Understanding and Context-Aware Retrieval. In *Proceedings of the 2nd International Conference on Artificial Intelligence, Digital Media Technology and Social Computing* (pp. 77-82).
19. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), Article 115, 1-35.
20. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610-623). ACM.
21. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.

22. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (pp. 33-44). ACM.
23. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., ... Liang, P. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.
24. Liu, T.-Y. (2009). Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3), 225-331.
25. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.