

Cross-Domain Transfer Learning for Personalized Recommendation and Dialogue State Prediction

Finn Bush

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
finn.bush270@buffalo.edu

Abstract

Modern human-centered artificial intelligence systems increasingly depend on the ability to model user intent across heterogeneous interaction channels. Personalized recommendation and dialogue state prediction are two tasks that, despite their distinct output spaces, share substantial representational structure related to preference inference, sequential context, and user modeling. Cross-domain transfer learning offers a unifying systems perspective for reusing knowledge across these tasks, but it also introduces architectural complexity, governance challenges, and operational risk. This paper presents a system-level examination of cross-domain transfer learning as applied to personalized recommendation and dialogue state prediction. It considers foundational modeling approaches, architectural design spaces, data governance requirements, fairness constraints, sustainability concerns, and deployment trade-offs. Rather than focusing on a single algorithmic contribution, the analysis emphasizes structural choices such as shared encoders, modular heads, domain adaptation mechanisms, knowledge distillation, and gateway-based orchestration. The paper further discusses how cross-domain transfer can amplify bias, create hidden technical dependencies, complicate privacy compliance, and alter carbon footprints. Policy implications are examined through the lens of trustworthy artificial intelligence, data protection law, and institutional accountability. The discussion highlights that successful cross-domain systems require coordinated governance across data, model, infrastructure, and evaluation layers, not merely improved predictive accuracy.

Keywords

cross-domain transfer learning, personalized recommendation, dialogue state prediction, system architecture, fairness, data governance, sustainable machine learning.

1. Introduction

The expansion of machine learning into interactive services has produced a growing need for models that can infer user preferences and goals across multiple contexts, including product recommendation, content ranking, and task-oriented dialogue. Personalized recommendation systems typically infer long-term and session-level preferences to select items from a potentially vast catalog, while dialogue state prediction systems maintain a structured representation of user objectives, constraints, and expressed preferences during conversation. Although these tasks differ in their output representations and evaluation metrics, they share common underlying challenges: modeling sequential user behavior, handling heterogeneous inputs, adapting to sparse signals, and maintaining coherent user models across time. Cross-domain transfer learning provides a conceptual and engineering framework for reusing knowledge learned in one domain or task to improve performance in another [1]. Early work in transfer learning established taxonomies for understanding distribution shift, feature space

mismatch, and label space divergence, while later surveys have documented the expansion of transfer techniques into deep learning, domain adaptation, and multi-task settings [2].

From a systems perspective, cross-domain transfer learning is not simply an algorithmic strategy. It involves decisions about how data are collected and shared, how representations are stored and served, how models are versioned and updated, and how failures in one domain propagate into another. These concerns are particularly acute when recommendation and dialogue state prediction are combined in a unified assistant, where a user might interact with a conversational interface that both clarifies intent and delivers personalized suggestions. The goal of this paper is to examine the structural, governance, infrastructural, and policy dimensions of such systems. The following sections analyze architectural design spaces, task-specific transfer challenges, fairness and privacy implications, operational sustainability, and institutional governance. In doing so, the paper emphasizes that robust cross-domain systems depend as much on careful socio-technical design as on predictive performance.

2. Background and Problem Framing

Personalized recommendation has traditionally been formulated as a preference estimation problem in which a system predicts how a user would evaluate an item. Latent factor techniques such as matrix factorization have long served as a foundation for collaborative filtering, mapping users and items into shared low-dimensional spaces through observed interaction data [3]. More recent neural approaches have extended this foundation by learning nonlinear interaction functions from user and item embeddings, enabling richer representations of preference structure [4]. Dialogue state prediction, by contrast, has historically been framed as a structured classification or extraction problem over user utterances and system responses. The state representation may include user intent, requested slots, constraints on entity values, and the current stage of a task. In task-oriented systems, state prediction functions as the core cognitive component that informs policy decisions and response generation.

Despite their different historical origins, recommendation and dialogue state prediction both rely on the ability to infer latent user intent from imperfect behavioral signals. In recommendation, a click or purchase is a noisy indicator of preference. In dialogue, a user utterance may be ambiguous, incomplete, or context-dependent. Both domains also face severe sparsity and cold-start problems when new users, new items, or new dialogue domains appear. Cross-domain transfer learning aims to mitigate these problems by sharing knowledge across related domains, such as transferring user preference representations from a mature recommendation domain to a newly launched conversational service, or transferring dialogue state representations across domains with different slot schemas. However, the feasibility of such transfer depends on the alignment of feature spaces, the consistency of label semantics, and the availability of shared user identity or entity linking infrastructure. Without careful system design, transferred knowledge can introduce harmful inductive biases that degrade rather than improve target domain performance.

3. Architectural Design Spaces for Cross-Domain Transfer

The architecture of a cross-domain transfer system determines what information is shared, when it is shared, and how specialized components interact. Transformer-based attention mechanisms have become the dominant substrate for modeling sequential and contextual information in both recommendation and dialogue systems [5]. Pre-trained bidirectional encoders further provide reusable contextual representations that can be fine-tuned for

downstream tasks, allowing a single network to support multiple prediction heads [6]. In cross-domain systems, a common architectural pattern involves a shared trunk that processes general user behavior or natural language, with separate heads for recommendation scoring and dialogue state extraction. This pattern reduces duplicate computation and allows gradients from multiple tasks to shape a common representation.

However, the shared trunk design also introduces several structural trade-offs. If the shared representation is too tightly coupled, changes to one task may degrade performance in another, a phenomenon often referred to as catastrophic interference or negative transfer. If the representation is too loosely coupled, the system may fail to exploit meaningful cross-task regularities. Engineering solutions include domain-specific adapters, separate normalization layers, gated fusion units, and task-specific attention masks. These mechanisms allow the system to modulate the degree of sharing without rebuilding the entire model for each new domain. From an infrastructure standpoint, the choice between a monolithic shared model and a modular collection of fine-tuned components affects serving latency, memory consumption, update frequency, and fault isolation. A monolithic model may be easier to deploy but harder to audit, while a modular system may improve maintainability at the cost of increased orchestration complexity. Consequently, architectural decisions in cross-domain transfer are inseparable from broader considerations of operational sustainability and governance.

4. Dialogue State Prediction and Transferable Dialogue Representations

Dialogue state prediction has evolved from rule-based slot filling toward data-driven representation learning. The neural belief tracker reformulated state tracking as a decision problem over candidate slot values, using distributed representations to compare an utterance with each candidate rather than relying on predefined lexical patterns [7]. This shift made it easier to transfer components across domains, because the underlying representation could be trained on data from multiple sources. Later architectures introduced external memory and pointer mechanisms to handle the challenge of unseen slot values and large knowledge bases, improving flexibility in multi-domain settings [8]. These methods illustrate a broader trend in dialogue systems: moving from hand-engineered domain ontologies toward learned representations that can be shared, adapted, or constrained at runtime.

Recent work has continued to integrate preference modeling into dialogue state prediction. A dual-stage transformer network with gated fusion for pairwise user preference prediction in dialogue systems has been proposed to address the need for more fine-grained preference signals within conversational interactions [9]. This type of architecture illustrates how dialogue systems increasingly borrow mechanisms from recommendation and ranking, such as pairwise comparison and gated feature fusion, while maintaining the structured state tracking required for task completion. The cross-domain challenge in dialogue state prediction remains significant because different domains may use different slot names, entity types, and dialogue acts. Even when a shared encoder is used, the output heads must often be retrained or adapted for each domain. Structural solutions include ontology-free state tracking, zero-shot slot transfer, and hybrid approaches that maintain an external schema registry. These solutions shift complexity from model training to infrastructure, requiring the system to manage schema mappings, data versioning, and online validation.

5. Personalized Recommendation Systems in Cross-Domain Settings

Large-scale recommendation systems have demonstrated the value of deep neural networks for retrieval and ranking. The YouTube recommendation architecture, for example, uses a

two-stage process in which candidate generation is followed by ranking, allowing the system to handle massive item catalogs while maintaining low latency [10]. Deep learning surveys of recommender systems highlight the importance of representing user-item interactions, side information, temporal dynamics, and sequential patterns in a unified framework [11]. Cross-domain recommendation extends these techniques by leveraging behavior from an auxiliary domain to improve recommendations in a target domain. For instance, a system might transfer user representations learned from music listening to improve movie recommendations, provided that the same user identity can be linked across services.

The system-level implications of cross-domain recommendation are substantial. Data from multiple domains may have different sparsity patterns, temporal resolutions, and noise characteristics. A user who frequently skips songs may still provide strong movie preferences, but the transfer of that behavior must account for domain-specific semantics. Moreover, cross-domain systems often require shared identity graphs, entity resolution pipelines, and consent mechanisms that govern how user data may be combined. In practice, the value of cross-domain transfer can be greatest for cold-start users and new domains, where target-domain interactions are too sparse to support reliable inference. Yet this benefit must be weighed against the risk of negative transfer, in which source-domain patterns mislead the target model. Structural safeguards include domain disagreement detectors, importance weighting, and staged deployment with shadow evaluation before a transferred model is allowed to influence live recommendations.

6. Fairness and Bias in Cross-Domain Personalization

Fairness in machine learning has become a central concern in the deployment of personalized systems. Surveys of bias and fairness document multiple sources of harm, including historical bias in training data, representational bias in features, and emergent bias introduced by model optimization [12]. Cross-domain transfer learning can amplify these harms if a source domain contains skewed distributions over user groups, item categories, or behavioral patterns. For example, if a recommendation model trained on a domain with known demographic skew is transferred to a new domain, the resulting representations may carry forward latent stereotypes or suppress minority preferences. Because the target domain may initially have little data of its own, these transferred biases can be difficult to detect through standard accuracy metrics.

Individual fairness frameworks emphasize that similar individuals should receive similar predictions, but defining similarity across domains is nontrivial [13]. Broader conceptual analyses have shown that fairness criteria may conflict, that context matters, and that technical interventions alone cannot resolve normative disagreements [14]. In cross-domain systems, these conflicts become more pronounced because a representation that is fair in one domain may be unfair in another. System designers must therefore treat fairness as a cross-layer property involving data selection, feature engineering, training objectives, evaluation metrics, and post-deployment monitoring. This requires not only algorithmic adjustments but also organizational processes for auditing transferred models, documenting domain assumptions, and establishing appeal mechanisms for affected users. When recommendation and dialogue state prediction are combined, fairness concerns also extend to conversational access: a user from a marginalized group may be denied certain recommendations because the dialogue system failed to accurately track their stated constraints.

7. Operational Infrastructure, Technical Debt, and Sustainability

The operational complexity of machine learning systems has been extensively documented through the concept of hidden technical debt, which includes data dependencies, entanglement, configuration fragility, and feedback loops [15]. Cross-domain transfer magnifies many of these issues because shared components create hidden couplings between otherwise separate product teams. A change in the data pipeline for a source domain can alter the behavior of a target domain model without any direct modification to that model. Similarly, a representation trained on multiple domains may have unclear ownership, making it difficult to assign responsibility when performance degrades. Infrastructure for cross-domain systems should therefore include explicit contracts for data schemas, feature semantics, model versioning, and evaluation baselines.

Sustainability is another critical dimension. Training large neural models for cross-domain transfer can consume significant energy, and the environmental costs of natural language processing have been shown to vary widely depending on model size, hardware, and location [16]. Reporting frameworks for carbon footprints have been proposed to make these costs more transparent [17]. Beyond training, inference at scale can dominate energy consumption in always-on recommendation and dialogue systems. Cross-domain shared encoders may reduce duplicate computation, but they may also require larger models that must be kept in memory. Model reuse can thus be environmentally beneficial or detrimental depending on how it is engineered. Moreover, large pre-trained language models have been criticized for concentrating power, reproducing harmful social biases, and incurring environmental harms [18]. These critiques are directly relevant to cross-domain systems that rely on shared foundation models for both recommendation and dialogue state prediction, because the convenience of reuse must be balanced against broader societal costs.

8. Governance, Privacy, and Foundation Model Policy

The increasing use of foundation models has changed the governance landscape for cross-domain transfer. Foundation models are large models trained on broad data that can be adapted to many downstream tasks, and they raise new questions about capability, homogenization, and accountability [19]. When a single foundation model serves both recommendation and dialogue state prediction, the resulting system inherits both the strengths and the risks of that model. Governance structures must address how such models are evaluated across domains, how their outputs are monitored, and how to manage the potential for systemic errors that affect many users simultaneously.

Data protection law adds further constraints. The European Union General Data Protection Regulation establishes principles of purpose limitation, data minimization, and consent that are particularly relevant when personal data are transferred across domains [20]. A user may have consented to data processing for a specific service but not for cross-domain model training. Legal analyses have also shown that the right to explanation is more limited than often assumed, which complicates efforts to provide meaningful transparency for automated decisions [21]. In cross-domain recommendation and dialogue systems, users may not understand why a particular item was recommended or why the system inferred a particular intent. Governance frameworks should therefore include clear privacy policies, data subject access mechanisms, and technical support for auditing. Institutional oversight boards can help evaluate whether cross-domain data sharing is proportionate and whether the expected benefits justify the risks to user autonomy.

9. Architectural Adaptation, Compression, and Domain Alignment

Several architectural techniques are available to manage the costs and risks of cross-domain transfer. Knowledge distillation allows a larger teacher model to transfer its knowledge to a smaller student model, thereby reducing latency and memory footprint while preserving much of the predictive power [22]. In cross-domain systems, distillation can be used to compress a shared multitask model into specialized domain-level models that are cheaper to serve. This approach can also improve fault isolation, because each domain model can be updated independently after distillation.

Domain adaptation methods provide another set of tools for aligning representations across domains. Adversarial training can encourage a feature extractor to produce representations that are indistinguishable across source and target domains, thereby reducing domain-specific variation [23]. Deep adaptation networks embed alignment objectives directly into the representation learning process, allowing domain discrepancy to be minimized while retaining task-specific discriminative capacity [24]. These techniques are valuable when the source and target domains have overlapping feature spaces but different marginal distributions. However, they also introduce additional optimization complexity and can be difficult to tune in production. System-level governance of these techniques requires careful evaluation of when alignment helps, when it hides important domain-specific structure, and how to detect silent failures in online serving.

10. Policy Implications and Institutional Governance

Policy frameworks for trustworthy artificial intelligence emphasize principles such as human agency, privacy, transparency, accountability, and environmental well-being [25]. Applying these principles to cross-domain transfer systems requires more than abstract commitment. It requires institutional mechanisms for model documentation, impact assessment, incident reporting, and user recourse. A cross-domain system that combines recommendation and dialogue state prediction operates at the intersection of personal data processing, persuasive technology, and conversational interaction, which increases its potential impact on user autonomy. Therefore, organizations should adopt a lifecycle approach in which risks are assessed before data sharing, during model training, and after deployment.

Institutional governance should also address the asymmetry between users and system operators. Users may not know that their behavior in one domain is being used to improve predictions in another. Even when consent is obtained, the implications may be difficult to grasp. Governance mechanisms such as model cards, data nutrition labels, and algorithmic impact assessments can improve transparency. External auditing can provide an additional layer of accountability, especially when cross-domain transfer affects vulnerable populations or high-stakes decisions. Ultimately, the policy challenge is to create structures that allow beneficial reuse of knowledge while preventing the concentration of unaccountable power. This requires collaboration among engineers, legal scholars, designers, and affected communities, not merely technical optimization.

11. Conclusion

Cross-domain transfer learning offers a promising path toward more coherent, data-efficient, and responsive systems for personalized recommendation and dialogue state prediction. By sharing representations across tasks and domains, such systems can reduce cold-start difficulties, improve conversational understanding, and support more natural user experiences. However, the benefits of transfer are accompanied by significant structural, operational, and governance challenges. Architectural decisions about shared encoders, domain heads, and

adaptation mechanisms shape not only predictive performance but also maintainability, fairness, and privacy. Operational concerns such as hidden technical debt, energy consumption, and model auditing become more complex when multiple domains interact. Fairness risks can be amplified by transferred biases, and data protection law imposes meaningful constraints on cross-domain data reuse.

Addressing these challenges requires an interdisciplinary perspective that treats technical architecture and institutional governance as complementary. Shared representations are not neutral assets; they carry assumptions from their source domains and can propagate errors at scale. System designers must therefore invest in robust evaluation, staged deployment, domain documentation, and user-facing transparency. Policymakers and institutional leaders must establish oversight processes that go beyond accuracy metrics and include fairness, sustainability, privacy, and accountability. Only through such integrated design can cross-domain transfer learning fulfill its potential in building personalized, context-aware systems that are both effective and socially responsible.

References

1. Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345-1359.
2. Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., & He, Q. (2021). A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1), 43-76.
3. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30-37.
4. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T.-S. (2017). Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web* (pp. 173-182). International World Wide Web Conferences Steering Committee.
5. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998-6008.
6. 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, Volume 1 (Long and Short Papers)* (pp. 4171-4186). Association for Computational Linguistics.
7. Mrkšić, N., Ó Séaghdha, D., Wen, T.-H., Thomson, B., & Young, S. (2017). Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1777-1788). Association for Computational Linguistics.
8. Wu, C.-S., Madotto, A., Hosseini-Asl, E., Xiong, C., Socher, R., & Fung, P. (2019). Global-to-local memory pointer networks for task-oriented dialogue. In *International Conference on Learning Representations*.
9. Shih, K., Deng, Z., Chen, X., Zhang, Y., & Zhang, L. (2025, May). DST-GFN: A Dual-Stage Transformer Network with Gated Fusion for Pairwise User Preference Prediction in Dialogue Systems. In *2025 8th International Conference on Advanced Electronic Materials, Computers and Software Engineering (AEMCSE)* (pp. 715-719). IEEE.

10. Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198). Association for Computing Machinery.
11. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1), Article 5, 1–38.
12. 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.
13. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (pp. 214–226). Association for Computing Machinery.
14. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68). Association for Computing Machinery.
15. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
16. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650). Association for Computational Linguistics.
17. Lacoste, A., Luccioni, A., Schmidt, V., & Dandres, T. (2019). Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*.
18. 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). Association for Computing Machinery.
19. 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*.
20. European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, L119, 1–88.
21. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99.
22. Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.

23. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., & Lempitsky, V. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1), 2096–2030.
24. Long, M., Cao, Y., Wang, J., & Jordan, M. I. (2015). Learning transferable features with deep adaptation networks. In *Proceedings of the 32nd International Conference on Machine Learning* (pp. 97–105). PMLR.
25. High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission.