Networked Tastes: Homophily for Music Recommendation

Marta Moscati[1], Xinwei Xu[2] and Markus Schedl[1,3]

[1]Institute of Computational Perception, Johannes Kepler University Linz, Linz, Austria
[2]Social Networks Lab, ETH Zürich, Zürich, Switzerland
[3]AI Lab, Linz Institute of Technology, Linz, Austria

Abstract

Many recommender systems implicitly assume homophily, i.e., that socially connected users have similar tastes. This assumption can be particularly problematic in group recommendation settings, where individuals' preferences can be non-homogeneous. Our study examines the homophily assumption in the music domain. Leveraging techniques from network analysis, we reconsider the definition of music similarity and analyze how users' taste similarity relates to social network structures across multiple levels. We then carry out recommendation experiments to evaluate how incorporating information on user–artist interactions, user–user connections, and artist–artist similarity, affects the quality of music recommendations. Our findings show that homophily operates differently across granularities, with a limited artist-level similarity among connected users, and a higher similarity at the artist-cluster level, and that different user groups benefit differently from the inclusion of social network and artist similarity data in recommendation quality. The results challenge homophily as a uniformly reliably proxy for inferring music preferences and suggest that a multi-granular, multi-layer notion of similarity better accounts for preference heterogeneity in group recommendations settings. Code: https://github.com/hcai-mms/homophily

Keywords

User Modeling, Music Streaming, Social Network Analysis, Music Recommender Systems

1. Introduction

Group recommendation [1] aims at recommending items to a group of users, often by aggregating individuals' preferences. However, individual preferences are rarely independent [2, 3], as they are shaped by individuals' social networks, shared cultural contexts, and collective perceptions of which items are 'similar'. This interdependency partially manifests through the well-documented phenomenon of 'homophily', i.e., the tendency of individuals to associate with others who share similar attributes [4]. Although rarely explicitly stated, such assumption is already used in many RS practices – for instance to aggregate users' preferences in group recommendation scenarios [1] or to address cold start [5, 6, 7]. Yet, recent work has revealed that tie strength may influence how group consensus forms [8] and that similarity in tastes can mediate how social network neighbors influence users' own taste [9, 10], calling into question the assumption of using social connections as uniform proxies for preference similarity. Moreover, perceptions of item similarity (e.g., whether two artists 'sound alike') depend not only on objective features like audio or stylistic characteristics, but also on group norms and community-based expectations [11, 12]. For example, in electronic dance music, acoustically similar tracks may be valued differently based on perceived authenticity and alignment with the genre's non-conformity values [13]. Since cultural items such as movies and music tracks often carry multiple interrelated or nested labels on digital platforms [14], perceived item similarity is often ambiguously defined and highly context-dependent, further complicating how individual preferences can be aggregated at group level. Taking the above considerations into account, aim of this work is to (a) evaluate the homophily assumption in a multi-valued setting where items similarity is evaluated at different granularities and (b) assess how integrating additional information on user–user social networks and on item similarity affects the quality of recommendations. We focus on the music domain, a frequent use-case for group recommendation, and ask:

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden
marta.moscati@jku.at (M. Moscati); xinwei.xu@gess.ethz.ch (X. Xu); markus.schedl@jku.at (M. Schedl)
0000-0002-5541-4919 (M. Moscati); 0000-0001-5740-3327 (X. Xu); 0000-0003-1706-3406 (M. Schedl)

RQ1 How do users' taste profiles relate to their social networks on digital platforms? And under what conditions does the homophily assumption break down? Addressing this question allows identifying scenarios where group-based naive aggregation strategies may harm some members.

RQ2 When jointly modeling user–user social connections, user–artist interactions, and artist–genre connections, do we observe improvements in recommendation quality? How does this affect users with various preference profiles?

2. Related Work

Taste homophily in empirical research. The homophily hypothesis [4] was originally studied in social sciences through survey data to measure offline social networks and cultural preferences (we kindly refer the reader to [15] for a general review). This notion is supported by studies in offline contexts (e.g., [16, 17, 18, 19]). Online settings – where social connections are digitally mediated and vast catalogs enable more fine-grained manifestation of preferences, and where the items often carry multi-valued labels (e.g., spanning multiple genres) – may show different dynamics of taste homophily [20, 9]. Recent empirical studies have shown that, while digital social connections correlate with similar music tastes, findings diverge on the underlying mechanisms and at which granularity homophily manifests. Early work on digital platforms [21] showed that self-perceived shared interests foster the creation of weak ties. In contrast, Bischoff et al. [22] found that the role of taste similarity in predicting friendship links is negligible. Lewis et al. [23] found that among a university student cohort on Facebook, sharing interests in music and movies at the genre level increases the chances of online friendship, but reported little evidence for the diffusion of tastes. At the level of network communities, Bisgin et al. [20] analyzed two large online platforms and documented high within-community similarity in shared top interests. More recently, Zhou et al. [24] documented music listening homophily on a large Chinese social media platform. While Duricic et al. [25] found that socially connected users prefer similarly diverse music, Shakespeare et al. [26] pointed out that the scale at which information is represented can affect measurements of user consumption diversity. The fragmented empirical findings seem to confirm this indication, raising the question of whether taste homophily traditionally established in offline contexts can be translated into online settings; and in turn, how preference granularity and units of analysis (e.g., network dyads vs. communities) shape observed manifestation of taste homophily on digital platforms.

Social network data in Recommmender Systems. Within RS research, social recommendation [5, 27, 28] is often leveraged to infer user preferences. For example, Krohn et al. [6] propose a multi-relational factorization approach learning user–user social graphs jointly with user–item interactions. Zhao et al. [7] leverage social relations as positive social feedback for recommendation. Most recent applications include the use of graph-neural-networks [27]. However, existing approaches overlook two limitations. First, friends' preferences are often uniformly taken as positive signals for user tastes [7, 9]. Surprisingly few studies [29] have examined how the degree of taste similarity varies as a function of connection type or of structural social network features [30, 31]. Second, the predominant focus on recommendation accuracy overlooked other aspects of recommendation quality that might benefit from the inclusion of social network data. This is surprising, especially given the increasing shift towards beyond-accuracy objectives of RS [32] and recent empirical evidence that social connections also serve as important sources for discovering novel content [10].

Addressing the aforementioned gaps, **our work** systematically evaluates the assumption of taste homophily on digital platforms across various levels of analysis (user-level, dyad-level, and network neighborhoods) and at various granularities of taste measurement (individual artists and artist-genre clusters). We then integrate these different layers of information into a heterogeneous information network [33] and assess their impact on music RS and for user groups of different music preferences.

3. Methodology

RQ1. Our analysis of taste homophily is based on measuring preference similarity in music consumption at multiple granularities. First, following prior work [34, 25], we measure taste similarity at the artist-level to avoid underestimating similarity between users listening to the same artists but different songs. We aggregate the user–track playcounts matrix to user–artists playcounts and restrict our analysis to users' top T listened artists by their total playcounts, in line with prior work utilizing users' top artists as positive signals of their preferences [20, 24]. We measure *artist-level taste similarity* using Jaccard index $\frac{|\mathcal{A}_i^T \cap \mathcal{A}_j^T|}{|\mathcal{A}_i^T \cup \mathcal{A}_j^T|}$, with $\mathcal{A}_i^T$ the set of top-T artists for user i. Second, to capture genre-level similarity, we need to address the issue of multi-valued categorization, i.e., artists may belong to multiple genres that are themselves interdependent or hierarchically structured. For instance, "garage rock" and "punk-rock" often co-occur for the same artist, while the first can be considered a sub-genre of the second. Simply treating these genres as discrete, independent labels would miss these relationships. Thus we model the relationship between artists and genres as an undirected bipartite network, where artists and genre labels are treated as nodes, with edges representing artist-genre associations. To preserve the bipartite structure and avoid the limitations of clustering on the one-mode projection of the bipartite network [35, 36], we adopt hierarchical stochastic block-modelling (hSBM) [37] to cluster artists and genre labels simultaneously [38]. This yields artist clusters that map onto multiple genre clusters, and verse versa, uncovering the latent structure emerging from the multi-valued associations between artists and genres while accommodating the inherent fuzziness of genre boundaries [12, 14]. We measure *genre-level taste similarity* using the resulting artist-genre clusters by computing cosine similarity $\frac{\mathbf{p}_i \cdot \mathbf{p}_j}{\|\mathbf{p}_i\| \, \|\mathbf{p}_j\|}$, with $\mathbf{p}_i \in \mathbb{R}^{N_{cluster}}$ user i's normalized playcounts aggregated over all clusters. Finally, we measure users' *taste diversity* [39, 40] as Shannon's entropy [1] over $\mathbf{p}_i$, and for pairs of users, we compute the *relative difference in taste diversity* $\frac{|D_i - D_j|}{(D_i + D_j)/2}$ as further indicator of taste similarity [25]. We also compute *taste mainstreamness* (MS), $\sum_a \sqrt{p_{i,a} \cdot g_a}$, with $p_{i,a}$ user i's playcount fraction and g_a global playcount fraction for artist a [44], segmenting users in three groups: low, medium, and high-MS [45].

We operationalize the homophily hypothesis [4] – that socially connected individuals tend to have similar preferences – at two levels. At the dyadic level, we explore whether directly connected users have more similar tastes than users separated by a larger social distance, measuring social distance as length of shortest path connecting a node pair. Second, at the neighborhood level, we examine whether users tend to have taste profiles that are similar to the average of their network neighbors.

RQ2. Our goal is to understand how information encoded with our hSBM approach impacts music recommendations for users of different mainstreamness. To disentangle the effects of input information vs. model architecture, we base our experiments on GAL-KARS [46], a recently proposed approach based on heterogeneous graphs. Compared to GAL-KARS, instead of leveraging LLMs to extend the CF graph, we design four graphs: including CF data ("I" for interactions), including CF and social data ("IS"), including CF and artist similarity ("IA"), and including all three ("ISA"). Following Spillo et al. [46], we use a graph-encoding approach to learn embeddings for users and artists; these are fine-tuned in a downstream embedding-based RS. For conciseness, we denote the resulting RS with GAS-KARS. We compare GAS-KARS with a downstream RS trained with random embedding initialization. We measure the quality of lists of k recommendations with nDCG at $k = 5$, and in terms of diversity, measured as Shannon's entropy over artist hSBM clusters at $k = 100$, to ensure a larger sample for computing diversity. To evaluate the impact of graph information on group recommendation fairness, we consider a scenario where recommendations are provided to groups of users with different mainstreamness levels (low, medium, high-MS). We compute nDCG and the RecGap [47], i.e., the average disparity between results across groups with varying levels of mainstreamness, to measure the (un)fairness of RS at the group level.

[1] Although other measures of diversity have been used in the music domain [41, 42, 26], previous work by Moscati et al. [43] has shown that results tend to be consistent across definitions.

4. Dataset and Experiment Setup

We base our experiments on the subset of LFM-2b [48] that overlaps with Music4All-Onion [49]. We extract social network information via Last.FM API's endpoint user.getFriends. We restrict to users belonging to the largest weakly connected component (LWCC) of the resulting social network to ensure meaningful computation of structural network metrics. We further apply 5-core filtering to users and items, and restrict to those with at least five listening records. These three filters (LWCC, user core filtering, item core filtering) are applied iteratively until convergence. We construct the user–artist interaction network restricting to users' top T artists, setting $T = 50$, since over 75% of users in our sample have listened to more than 49 distinct artists (IQR in Table 1). We construct the artist–genre bipartite network used for our hSBM cluster analysis based on the user-generated genre labels from Music4All-Onion [49] by assigning to each artist the genres of their tracks. For recommendation, we use the binary user–artist top-T matrix. Table 1 provides the characteristics of our dataset. Following Spillo et al. [46], we extract the graph embeddings with CompGCN [50], tuning the number of layers $\ell \in \{1, 2, 3\}$ based on the validation nDCG and fixing the number of epochs to 15. We use matrix factorization with Bayesian personalized ranking (BPR) [51] as downstream RS, We fix the embedding dimension, which has to match the one of CompGCN, to 64, the learning rate to 0.001, and the maximum number of epochs to 50, selecting the epoch with highest validation nDCG. We report mean and standard error over users, computed on test interactions. We split interactions in a train, validation, and a test set by randomly selecting $80\%, 10\%, 10\%$ of interactions of each user. The 80% training set is also used as the CF ("I"') data for learning the graph embeddings.

Table 1
Descriptive Statistics of the Dataset. U: Users; T: Tracks, A: Artists, G: Genres

User-Item Network		Artist-Genre Network		User-User Network	
U	17,680	Artists	9,339	Following ties	33,027
T	54,873	Genres	715	Avg. in-/out-degree	1.87
U–T int.	9,119,770	Edges	30,171	SD in-/out-degree	4.530 / 1.206
U–A int.	2,476,176	Avg. G per A (SD)	3.23 (2.32)	Reciprocity	0.554
Median A per U (IQR)	96 (49-180)	Avg. A per G (SD)	42.197 (155.792)	Transitivity	0.011

5. Results

RQ1. Taste Homophily. We first examine the overall structural properties of the user-user following social network (Table 1). On average, each user has 1.868 connections (either following or being followed), indicating an extremely sparse network. Reciprocated connections account for 55% of the total. The global transitivity – i.e., the ratio of the number of closed triangles and connected triples – is low (0.011), indicating a very low level of clustering. Addressing RQ1, Figure 1 summarizes how measures of taste similarity relate to the social proximity of nodes – i.e., the shortest path length connecting a node pair. Overall, directly connected nodes (distance$=$ 1) have a higher level of taste similarity than nodes separated by larger social distances, providing evidence for taste homophily at the dyadic level. On average, directly connected user pairs share about $0.07 \cdot 50 = 3.5$ artists in common among their top 50 listened artists (subplot (a)) – quite low in absolute terms; yet their average cosine similarity based on artist-cluster profiles is about 0.45 (subplot (b)), indicating modest similarity. This suggests that connected users may share some genre/artistic style preferences even when they do not listen to the same artists. Both similarity measures decline sharply as social distances increase among node pairs, and then reach the global sample mean. In contrast, while we also observe smaller differences in taste diversity among connected users than the global sample mean (subplot (c)), diversity does not show as clear a distance-dependent trend as the two similarity measures – the fluctuations across distances seem rather random.

Next, Figures 2a and 2b highlight two key findings on taste homophily at the network neighborhood level. We first compare taste diversity – a user-level attribute – between focal users (ego) and their

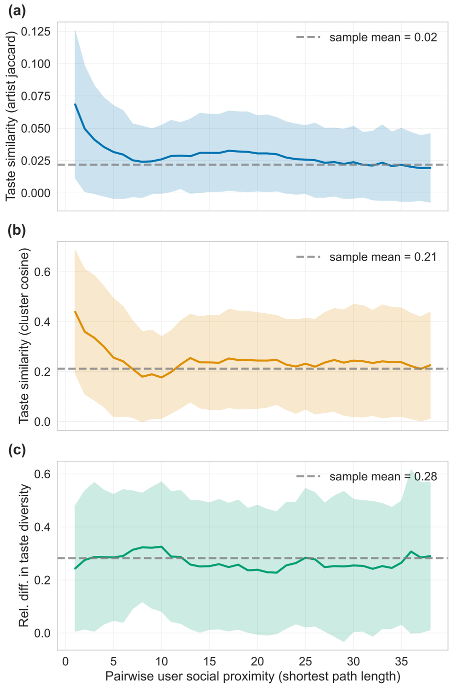

Figure 1: Pairwise taste similarity by proximity in the social network. Solid lines show mean values from 1,000 randomly sampled user pairs per distance bin; shaded areas show standard deviations. Dashed lines show the global mean from 10,000 randomly drawn user pairs.

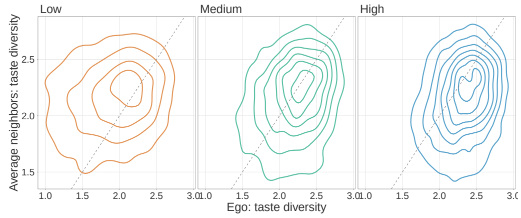

(a) User diversity (x-axis) vs. average neighbor diversity (y-axis). Dashed lines represent $y = x$. Colors indicate users' MS level.

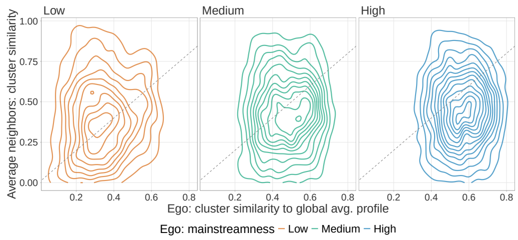

(b) User similarity to global average artist cluster profile (x-axis) vs. average pairwise similarity to network neighbors (y-axis). Dashed lines represent $y = x$. Colors indicate users' MS level.

Figure 2: Neighborhood-level homophily (users vs. social network neighbors), broken down by user mainstreamness (MS) tercile.

average network neighbors, segmented by focal users' mainstreamness (MS). As shown in Figure 2a, low-MS users tend to have neighbors with more diverse tastes, while high-MS users have more diverse tastes than their average neighbors. We then compare users' average taste similarity to their network neighbors (an edge-level attribute) against their similarity to the global average profile as a reference (Figure 2b). On average, low-MS users are more similar to their network neighbors than to the global average profile, while high-MS users are less similar to their network neighbors than the global average. This suggests that taste profiles of social network neighbors provide more valuable signals for inferring the preferences of low-MS users than for medium- or high-MS users.

Table 2

Recommendation performance. I: CF data; S: social connections; A: artists clusters.

	Input	nDCG	Entropy	nDCG $^{\text{Low-MS}}$	nDCG $^{\text{Medium-MS}}$	nDCG $^{\text{High-MS}}$	RecGap ↓
GAS-KARS	ISA	$\mathbf{0.1336}_{0.0003}$	$3.728_{0.004}$	$0.1192_{0.0051}$	$0.1338_{0.0050}$	$\mathbf{0.1290}_{0.0053}$	0.0097
GAS-KARS	IA	$0.1331_{0.0003}$	$3.717_{0.004}$	$\mathbf{0.1244}_{0.0053}$	$0.1360_{0.0052}$	$0.1271_{0.0054}$	**0.0077**
GAS-KARS	I	$0.1329_{0.0003}$	$\mathbf{3.743}_{0.004}$	$0.1179_{0.0051}$	$0.1361_{0.0052}$	$0.1260_{0.0052}$	0.0121
GAS-KARS	IS	$0.1326_{0.0003}$	$3.731_{0.004}$	$0.1208_{0.0052}$	$\mathbf{0.1381}_{0.0053}$	$0.1241_{0.0052}$	0.0115
BPR	—	$0.1288_{0.0003}$	$3.689_{0.004}$	$0.1145_{0.0050}$	$0.1336_{0.0052}$	$0.1200_{0.0052}$	0.0127

RQ2. Impact on Recommendation. Table 2 ranks models in order of decreasing nDCG, with varying input configurations for GAS-KARS. We first observe that adding graph information always improves recommendation accuracy, since all GAS-KARS variants outperform BPR, even when GAS-KARS leverages the same CF data used for BPR. However, pre-training an instance of GAS-KARS comes with substantial time and energy cost and considering the marginal improvement, it might not be cost-efficient. Addressing RQ2 (*Are social-network and artist-genre information useful for recommendation?*), we observe that including artist-genre ("IA" and "ISA") information improves the recommendation accuracy, although the differences in nDCG are not statistically significant with respect to other GAS-KARS

variants. Interestingly, all GAS-KARS variants simultaneously improve recommendation diversity (entropy) compared to the BPR baseline, indicating that recommendations are not only more accurate but also more evenly distributed across artist-genre clusters. This result is insightful, since these two aspects of recommendation quality have often been shown to be in a trade-off relationship. The overall trends hence suggest that leveraging social-network information and modeling artist genres with our proposed bipartite clustering approach might be an effective way to improve one of the two metrics without negatively impacting the other.

Since our previous analysis showed that taste homophily varies by users' mainstreamness, we also evaluate recommendation accuracy by user groups with different MS levels. The nDCG divided by MS groups (Low, Medium, High) partially confirms findings in previous work [45]: low-MS users receive the least accurate recommendations, while, interestingly medium-MS users benefit the most (rather than high-MS as previously reported). As expected, BPR has the highest RecGap, whereas leveraging graph information in all variants of GAS-KARS results in improvements for all MS groups, with the lowest RecGap achieved under the IA condition. Interestingly however, within each MS group, the best-performing condition differs: for low-MS highest under IA, medium-MS under IS, and high-MS under ISA, with the latter also coinciding with the highest nDCG overall. Thus selecting the globally most accurate model would disproportionally benefit high-MS users. Altogether, these results suggest that social connections are not a uniformly reliable proxy for music preferences, and different MS groups benefit differently from the inclusion of different information layers.

6. Conclusions, Limitations, and Future Work

Our study first explored the homophily assumption at multiple levels by including user-, dyad-, and network-level features, while extending the notion of taste similarity to accommodate multi-valued item categorization across granularities. We then assessed the impact of including multi-layer graph-based features characterizing taste homophily on music recommendation quality. Our results show that homophily operates differently across preference granularities and levels of analysis, thereby challenging the use of social connections as uniform signals for preference similarity; moreover, different user groups benefit differently from the inclusion of social network data and artist-genre data. Notably, while low-MS users exhibit stronger homophily with their network neighbors, they benefit the most from including artist-genre data (IA condition). This shows that local network homophily does not straightforwardly translate into recommendation gains, and a multi-layered network approach with group-specific layer weights may better account for user preference heterogeneity.

Several limitations invite future work. First, our social network data were extracted from Last.fm. Future work could consider using data from other sources that better reflect users' offline social ties to see whether our findings generalize across different types of social networks or networks with varying tie strength. Second, while our attempt to address multi-valued categorization revealed artist-genre clustering structure, these clusters could be further validated against objective features such as coherence of within-cluster track content representations or high-quality expert-curated annotations. Third, a further investigation into the temporal dynamics may reveal valuable insights into the evolution of users' tastes and potentially disentangle social selection from social influence. Finally, our recommendation experiments only focused on BPR and GAS-KARS within in a single RS domain (music), future work could examine the generalizability of our results across alternative models and domains.

Acknowledgments

This research was funded in whole or in part by the Austrian Science Funds (FWF): *Intent-aware Music Recommender Systems* (https://doi.org/10.55776/P36413), the Cluster of Excellence *Bilateral Artificial Intelligence* (https://doi.org/10.55776/COE12), and the doc.funds.connect project *Human-Centered Artificial Intelligence* (https://doi.org/10.55776/DFH23), and by the Swiss National Science Foundation (SNSF) (grant No.10001512). For open access purposes, the authors have applied a CC BY public copyright license to any author-accepted manuscript version arising from this submission.

Declaration on Generative AI

While preparing the analysis that produced the reported figures, the authors used Copliot (GitHub Copilot integration with VSCode) for code auto-completion and debugging. Claude Code was used for refactoring the code to separate the homophily from the recommendation part. After using these tools, the authors reviewed and edited the code as needed and take full responsibility for the content produced. No GenAI tool was used for other purposes.

References

[1] J. Masthoff, A. Delic, Group recommender systems: Beyond preference aggregation, in: F. Ricci, L. Rokach, B. Shapira (Eds.), Recommender Systems Handbook, 2022, pp. 381–420.
[2] C. F. Manski, Identification of endogenous social effects: The reflection problem, The review of economic studies 60 (1993) 531–542.
[3] J. Wang, A. Aribarg, Y. F. Atchadé, Modeling choice interdependence in a social network, Marketing Science 32 (2013) 977–997.
[4] M. McPherson, L. Smith-Lovin, J. M. Cook, Birds of a feather: Homophily in social networks, Annual review of sociology 27 (2001) 415–444.
[5] H. Ma, H. Yang, M. R. Lyu, I. King, Sorec: social recommendation using probabilistic matrix factorization, in: Proc. of ACM CIKM, 2008, pp. 931–940.
[6] A. Krohn-Grimberghe, L. Drumond, C. Freudenthaler, L. Schmidt-Thieme, Multi-relational matrix factorization using bayesian personalized ranking for social network data, in: Proc. of ACM WSDM, 2012, pp. 173–182.
[7] T. Zhao, J. McAuley, I. King, Leveraging social connections to improve personalized ranking for collaborative filtering, in: Proc. of ACM CIKM, 2014, pp. 261–270.
[8] F. Barile, J. Masthoff, S. Rossi, A detailed analysis of the impact of tie strength and conflicts on social influence, in: Adjunct Proc. of ACM UMAP, 2017, pp. 227–230.
[9] N. Li, C. Gao, D. Jin, Q. Liao, Disentangled modeling of social homophily and influence for social recommendation, IEEE Transactions on Knowledge and Data Engineering 35 (2022) 5738–5751.
[10] M. Arvidsson, P. Hedström, M. Keuschnigg, Wide social influence and the emergence of the unexpected: An empirical test using spotify data, Sociological Science 12 (2025) 715–742.
[11] J. C. Lena, Meaning and membership: samples in rap music, 1979–1995, Poetics 32 (2004) 297–310.
[12] J. Vlegels, J. Lievens, Music classification, genres, and taste patterns: A ground-up network analysis on the clustering of artist preferences, Poetics 60 (2017) 76–89.
[13] M. Jadidi, H. Lietz, M. Samory, C. Wagner, The hipster paradox in electronic dance music: How musicians trade mainstream success off against alternative status, in: Proc. of AAAI ICWSM, volume 16, 2022, pp. 370–380.
[14] M. Boualami, C. Roth, Investigating musical taxonomy in the era of streaming platforms: Insights from rap music through actual consumption data, Poetics 111 (2025) 102006.
[15] C. Steglich, T. A. Snijders, M. Pearson, Dynamic networks and behavior: separating selection from influence, Sociological methodology 40 (2010) 329–393.
[16] M. H. Selfhout, S. J. Branje, T. F. ter Bogt, W. H. Meeus, The role of music preferences in early adolescents' friendship formation and stability, Journal of Adolescence 32 (2009) 95–107.
[17] A. Edelmann, S. Vaisey, Cultural resources and cultural distinction in networks, Poetics 46 (2014) 22–37.
[18] A. Lomi, C. Stadtfeld, Social networks and social settings: Developing a coevolutionary view, Kölner Zeitschrift für Soziologie und Sozialpsychologie 66 (2014) 395–415.
[19] D. Hachen, C. Wang, B. Sepulvado, O. Lizardo, Generators or diffusers? examining differences in

the dynamic coupling of context and social ties across multiple types of foci, Social Networks 77 (2024) 151–165.
[20] H. Bisgin, N. Agarwal, X. Xu, Does similarity breed connection?-an investigation in blogcatalog and last. fm communities, in: Proc. of IEEE SOCIALCOM, 2010, pp. 570–575.
[21] N. K. Baym, A. Ledbetter, Tunes that bind? predicting friendship strength in a music-based social network, Information, Communication & Society 12 (2009) 408–427.
[22] K. Bischoff, We love rock'n'roll: analyzing and predicting friendship links in last. fm, in: Proc. of ACM WebSci, 2012, pp. 47–56.
[23] K. Lewis, M. Gonzalez, J. Kaufman, Social selection and peer influence in an online social network, Proceedings of the National Academy of Sciences 109 (2012) 68–72.
[24] Z. Zhou, K. Xu, J. Zhao, Homophily of music listening in online social networks of china, Social Networks 55 (2018) 160–169.
[25] T. Duricic, D. Kowald, M. Schedl, E. Lex, My friends also prefer diverse music: homophily and link prediction with user preferences for mainstream, novelty, and diversity in music, in: Proc. of IEEE/ACM ASONAM, 2021, pp. 447–454.
[26] D. Shakespeare, V. Chareyron, C. Roth, Reframing the filter bubble through diverse scale effects in online music consumption, Scientific Reports 15 (2025) 4071.
[27] K. Sharma, Y.-C. Lee, S. Nambi, A. Salian, S. Shah, S.-W. Kim, S. Kumar, A survey of graph neural networks for social recommender systems, ACM Computing Surveys 56 (2024) 1–34.
[28] I. Guy, Social recommender systems, in: F. Ricci, L. Rokach, B. Shapira (Eds.), Recommender Systems Handbook, 2022, pp. 835–870.
[29] X. Wang, S. C. Hoi, M. Ester, J. Bu, C. Chen, Learning personalized preference of strong and weak ties for social recommendation, in: Proc. of ACM WWW, 2017, pp. 1601–1610.
[30] R. Xie, Q. Liu, S. Liu, Z. Zhang, P. Cui, B. Zhang, L. Lin, Improving accuracy and diversity in matching of recommendation with diversified preference network, IEEE Transactions on Big Data 8 (2021) 955–967.
[31] X. Long, C. Huang, Y. Xu, H. Xu, P. Dai, L. Xia, L. Bo, Social recommendation with self-supervised metagraph informax network, in: Proc. of ACM CIKM, 2021, pp. 1160–1169.
[32] T. Duricic, D. Kowald, E. Lacic, E. Lex, Beyond-accuracy: a review on diversity, serendipity, and fairness in recommender systems based on graph neural networks, Frontiers in big data 6 (2023) 1251072.
[33] C. Shi, Y. Li, J. Zhang, Y. Sun, P. S. Yu, A survey of heterogeneous information network analysis, IEEE Transactions on Knowledge and Data Engineering 29 (2016) 17–37.
[34] C. Bauer, M. Schedl, Investigating cross-country relationship between users' social ties and music mainstreaminess, in: Proc. of ISMIR, 2018, pp. 678–686.
[35] T. Zhou, J. Ren, M. Medo, Y.-C. Zhang, Bipartite network projection and personal recommendation, Physical Review E–Statistical, Nonlinear, and Soft Matter Physics 76 (2007) 046115.
[36] Z. P. Neal, A. Cadieux, D. Garlaschelli, N. J. Gotelli, F. Saracco, T. Squartini, S. T. Shutters, W. Ulrich, G. Wang, G. Strona, Pattern detection in bipartite networks: A review of terminology, applications, and methods, PLOS Complex Systems 1 (2024).
[37] T. P. Peixoto, Bayesian stochastic blockmodeling, Advances in network clustering and blockmodeling (2019) 289–332.
[38] M. Gerlach, T. P. Peixoto, E. G. Altmann, A network approach to topic models, Science advances 4 (2018).
[39] B. Ferwerda, M. Schedl, Investigating the relationship between diversity in music consumption behavior and cultural dimensions: A cross-country analysis, in: Proc. of SOAP Workshop co-located with ACM UMAP, volume 1618, 2016.
[40] B. Ferwerda, M. Tkalčič, Exploring online music listening behaviors of musically sophisticated users, in: Adjunct of Proc. of ACM UMAP, 2019, pp. 33–37.
[41] S. F. Way, S. Gil, I. Anderson, A. Clauset, Environmental changes and the dynamics of musical identity, in: Proc. of AAAI ICWSM, 2019.
[42] A. Anderson, L. Maystre, I. Anderson, R. Mehrotra, M. Lalmas, Algorithmic effects on the diversity

of consumption on spotify, in: Proc. of WWW, 2020.
[43] M. Moscati, A. Darius, M. Schedl, S. Bruno, Familiarizing with music: Discovery patterns for different music discovery needs, in: Proc. of ACM UMAP, 2025, pp. 63–73.
[44] M. Schedl, D. Hauger, Tailoring music recommendations to users by considering diversity, mainstreaminess, and novelty, in: Proc. of ACM SIGIR, 2015, pp. 947–950.
[45] D. Kowald, M. Schedl, E. Lex, The unfairness of popularity bias in music recommendation: A reproducibility study, in: Proc. of ECIR, Springer, 2020, pp. 35–42.
[46] G. Spillo, C. Musto, M. Mannavola, M. de Gemmis, P. Lops, G. Semeraro, Gal-kars: Exploiting llms for graph augmentation in knowledge-aware recommender systems, in: Proc. of ACM UMAP, 2025, p. 73–82.
[47] A. B. Melchiorre, N. Rekabsaz, E. Parada-Cabaleiro, S. Brandl, O. Lesota, M. Schedl, Investigating gender fairness of recommendation algorithms in the music domain, Information Processing & Management 58 (2021).
[48] M. Schedl, S. Brandl, O. Lesota, E. Parada-Cabaleiro, D. Penz, N. Rekabsaz, Lfm-2b: A dataset of enriched music listening events for recommender systems research and fairness analysis, in: Proc. of ACM CHIIR, 2022, pp. 337–341.
[49] M. Moscati, E. Parada-Cabaleiro, Y. Deldjoo, E. Zangerle, M. Schedl, Music4all-onion–a large-scale multi-faceted content-centric music recommendation dataset, in: Proc. of CIKM, 2022, pp. 4339–4343.
[50] S. Vashishth, S. Sanyal, V. Nitin, P. Talukdar, Composition-based multi-relational graph convolutional networks, in: Proc. of ICLR, 2020.
[51] S. Rendle, C. Freudenthaler, Z. Gantner, L. Schmidt-Thieme, Bpr: Bayesian personalized ranking from implicit feedback, in: Proc. of UAI, 2009, pp. 452–461.