Improving Music Collaboration Recommendations Using TikTok Social Connections

Alejandro Fernandez-Sanchez[0009−0000−9824−684X], Pedro J. Navarro[0000−0001−8367−2934], and Fernando Terroso-Saenz[0000−0002−1921−1137]*

Universidad Politécnica de Cartagena member of European University of Technology EUT+, Division of Systems and Electronic Engineering, Department of Information and Communication Technologies, Technical School of Telecommunication Engineering. Plaza del Cronista Isidoro Valverde 30202 Cartagena (Spain).
*Corresponding author
alejandro.fernandezs@edu.upct.es
{pedroj.navarro,fernando.terroso}@upct.es

Abstract. In today's music industry, collaborations among musicians are increasingly important because they help artists reach new audiences and markets. In that sense, social ties, as reflected in online social networks, are strong predictors of future musical collaborations. In this context, this paper presents a graph-based recommender system for musical collaborations. The main hypothesis of this work is that explicitly including social relationships extracted from TikTok in a heterogeneous graph neural network will improve the prediction of future artist collaborations. To test this idea, we built a heterogeneous knowledge graph with artist, track, and tag nodes, and formulated collaboration prediction as a link prediction task. The proposed architecture uses graph convolution layers to learn artist representations from musical, behavioral, and social data, followed by dense layers to refine the embeddings. We compared two models: one based on MusicBrainz and Last.fm only, and another that also includes TikTok follow relationships and artist activity features. The results show that the TikTok-enhanced model consistently outperforms the baseline across different popularity levels and prediction horizons, improving precision, F1-score, accuracy, and AUC. These findings suggest that TikTok provides complementary social information that helps identify plausible collaborations.

Keywords: Recommender system · Music Collaborations · Heterogeneous graph neural networks · TikTok data.

1 Introduction

The past two decades have witnessed a marked shift in the organization of musical production and consumption: collaborations between artists have become a central tactic for audience expansion, music genre hybridization and enduring market positioning [15].

Concurrently, the promotional logic of the industry has become increasingly data-driven to pair artists whose audiences partially overlap [11]. In that sense, several works claim that social ties are stronger predictors of future collaboration than mere music genre co-occurrence, and document a pattern of *complementary homophily* where collaborators share interests but differ in stylistic tendencies [16]. These findings suggest that online social networks (OSNs), like X or TikTok, can capture underlying social relations that precipitate creative partnerships [8].

From a methodological standpoint, previous works in the music domain have already proposed the usage of Heterogeneous Knowledge Graphs (HKGs) and Graph Neural Networks (GNNs) to integrate structural topology, multi-modal data and behavioral signals into unified embeddings for different tasks [5, 22]. However, a limited number of proposals are explicitly designed to suggest artist-to-artist collaborations based on graph and link-prediction approaches [18, 6]. Besides, one important limitation observed in these previous artist-centered works is the fact that they did not explicitly encode the social ties among artists in their HKGs.

Building upon the limitations identified above, we formulate the following hypothesis: *the explicit integration of social relationships derived from an OSN like TikTok into a GNN-based recommendation framework for artist collaboration will enhance its effectiveness.* By incorporating *follower–following* interactions from this platform, together with established music data repositories such as *MusicBrainz*[1], the resulting HKG is expected to capture not only stylistic compatibility but also socio-audience dynamics that influence successful partnerships. Under this assumption, we expect that modeling these social ties will improve link-prediction performance for unseen collaborations and generate recommendations.

The rest of this paper is organized as follows. Section 2 reviews current trends and techniques used in analyzing music streaming data. In section 3, we describe the datasets and the internal design of the recommender system, while section 4 evaluates its performance. Finally, section 5 concludes with the main findings and outlines directions for future research.

2 Related Work

Analyzing streaming and social music data has become a central theme in music recommendation and discovery research. In this context, several studies have proposed graph-based models that capture the complex relationships among listeners, tracks, and artists [5, 21]. For example, hybrid architectures combining convolutional feature extractors with recurrent modules have been used to fuse temporal signals with content-based descriptors for personalized playlists and context-dependent suggestions [10]. Other works have explored compact transformer variants and MLP-based models for music recommendation from questionnaires or extracted audio/text features [12].

[1] https://musicbrainz.org/

Other contributions have adapted large-scale graph recommendation architectures to music data. In particular, PinSage-style models and their variants have been tailored to playlist co-occurrence or streaming-session graphs, often outperforming classical collaborative-filtering baselines [2]. Complementary approaches combine audio-derived descriptors with graph topology. For instance, GNNs trained with metric or triplet losses over artist graphs have proven useful for related-artist recommendation tasks [4].

From a data-modeling perspective, structured semantic knowledge has also been incorporated to enrich music recommenders. Music knowledge graphs have improved both the accuracy and explainability of recommendations by enabling semantically coherent suggestions and explicit relational justification [1]. In addition, multi-task frameworks that jointly learn graph embeddings and recommendation objectives have shown consistent gains over single-task baselines [13].

The impact of TikTok on the music industry has been examined from multiple perspectives. Some studies analyze how platform-driven virality influences musical production and composition [17], while others investigate user interaction and consumption patterns on TikTok [20]. In the context of recommendation systems, prior work has proposed methods for suggesting background music for short videos before publication [3]. Research on TikTok as a facilitator of collaboration has also shown that the platform can support distributed creativity through its *Duet* feature and related interaction formats [9, 14].

Finally, our work follows the graph-based link-prediction formulation used in recent music-collaboration recommendation studies [6]. However, unlike that prior work, our framework incorporates TikTok as an additional social source and therefore captures artist-to-artist social proximity beyond purely musical relations. This distinction is important because social ties can reflect forms of complementarity that are not strictly musical but remain highly relevant for successful collaboration [16].

3 Materials and Methods

In this section, the technical and information sources along with methodological procedures employed in the design and development of the recommender are detailed, including the description of the Heterogeneous Knowledge Graph (HKG) to model musical collaborations, the data feeds used to populate it and the architectural design of the GNN models to develop the recommender.

3.1 Description of the Datasources

This study takes three datasources to create the recommender, namely the MusicBrainz database, the Last.fm platform and the TikTok OSN. We provide a brief description of each of them next.

The MusicBrainz Database (MBDB) is a large-scale, open repository of music metadata built on top of the PostgreSQL relational database system[2]. It contains rich core information about music-related entities (artists, release groups and so forth) as well as a wide range of relationships connecting these entities. For the purposes of this study, three primary entity types were selected from MBDB: *Artist*, to capture information related to musicians; *Track*, to represent their musical works; and *Tag*, to incorporate additional descriptive metadata such as musical genres. A variety of features were extracted for each of these entities, along with the relationships connecting them to compose the HKG as it is put forward in sec. 3.2

Last.fm (LFM) is a music recommendation service that operates by tracking users' listening habits. While MBDB is suitable for establishing historical and explicit artistic relationships, such as joint authorship or direct participation in recordings, it lacks quantitative information about artists' actual impact and audience reception. In this regard, Last.fm provides direct popularity metrics (unique listeners and total plays) and it makes it possible to uncover latent affinities between musicians through users' cross-listening patterns. These data were obtained through its public API[3].

TikTok (TT) served as a crucial source of dynamic and trend-aligned data for the HKG, providing indicators of artists' social ties that complement the static metadata of MBDB and the popularity patterns of LFM. Access to these insights required the use of the *TikTok Research API*[4] due to its restrictive data policies. To connect TikTok profiles with MBDB entities, Wikidata[5] was used as an interoperability bridge by linking TikTok user identifiers with MBDB identifiers.

3.2 Definition of the Heterogeneous Knowledge Graph

The core of the proposed system is a Heterogeneous Knowledge Graph (HKG). This structure enables the representation of different types of entities, as well as the multiple ways in which they are interconnected. Fig. 1 presents a visual representation of the schema. As illustrated in the figure, the proposed graph is composed of three main types of nodes: *Artist*, *Track*, and *Tag*.

The *Artist* node constitutes the central entity of the graph. Each node represents a musician, band, or unified musical entity (e.g., a symphony orchestra). The attributes associated with each artist, including the data source from which they were obtained and a brief description, are detailed in Table 1. As shown in this table, each *Artist* node is characterized by four distinct features extracted

[2] https://musicbrainz.org/
[3] https://www.last.fm/api/
[4] https://developers.tiktok.com/products/research-api/
[5] https://www.wikidata.org/

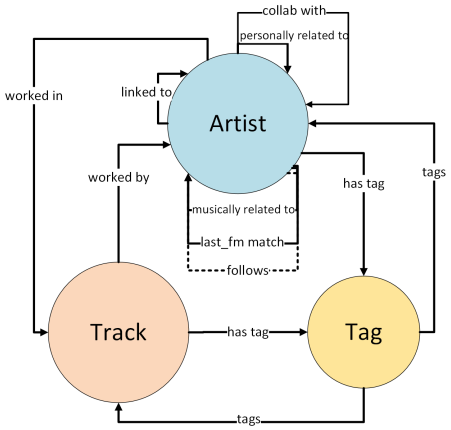


Fig. 1. Representation of the HKG capturing the relationships and attributes of artists, tracks, and genres. The *follows* relationship depicted with a dashed line has been directly extracted from the TikTok platform.

from MBDB, namely the *begin date* specifies the starting point of the artist's career, while *gender* describes the artist's gender. The *type* attribute categorizes the artist (e.g., solo performer, group, or other entity). The *end date* captures the time at which their career concluded, when applicable. Furthermore, the attributes *listeners* and *playcount* extracted from LFM represents the popularity level of the artist. Last, four different attributes were extracted from TT to capture the activity of the artist in such OSN (*likes_count* and *video_count*) along with his social connectivity level in such OSN (*follower_count* and *following_count*).

Table 1. Attributes of the *Artist* node in the HKG.

Attribute	Data Source	Description
`begin_date`	MBDB	Starting year (birth or formation) of the identity.
`end_date`		Ending year (death or dissolution) of the identity.
`gender`		Gender: Male, Female, Other, Not applicable, or Non-binary.
`type`		Type of artist: Person, Group, Other, Character, Orchestra, or Choir.
`listeners`	LFM	Number of Last.fm users who listen to the artist.
`playcount`		Number of artist play counts on Last.fm.
`likes_count`	TT	Total number of *likes* on TT.
`video_count`		Total number of videos published on TT.
`follower_count`		Total number of followers on TT.
`following_count`		Total number of accounts followed by the artist on TT.

On the other hand, the *Track* node establishes a common reference axis that allows artists to be unambiguously linked to their works and their corresponding musical attributes within the HKG, as described in Table 2. These attributes

enable the representation of the release date (*year* and *month*) of each track in the market along with its popularity level in the Last.fm platform by means of the *listeners* and *playcount* features.

Table 2. Attributes of the *Track* node in the HKG.

Attribute	Data Source	Description
`year`	MBDB	Release year.
`month`		Release month.
`artist_count`		Total number of artists credited on the track.
`listeners`	LFM	Number of Last.fm users who listen to the track.
`playcount`		Number of track play counts on Last.fm.

The *Tag* node acts as a connecting element that groups artists and tracks under shared concepts (e.g., genres, moods, and countries). Unlike the other node types, the *Tag* node does not possess any attributes other than a single *label* attribute that identifies the specific concept or category represented by the tag.

Moreover, the HKG in Fig. 1 includes several types of edges, from artist–artist or artist–song relationships to tagging relationships which link artists or tracks to semantic descriptors. In this regard, Table 3 briefly describes the meaning of each edge type. As we can see, the edges extracted from MBDB or LFM represent either purely artistic connections (e.g., `musically_related_to`), listener-based similarities (e.g., `last_fm_match`), or close-family relationships (e.g., `personally_related_to`). However, the `follows` edges extracted from TT allow us to extend these connections by also capturing ties among musicians that are not directly related to artistic or family-based similarities but to social interests or common friendships.

Table 3. Types of edges in the HKG.

Type	Data Source	Name	Description
Artist–Artist	MBDB	`COLLAB_WITH`	Connects two artists who collaborated on a song
		`MUSICALLY_RELATED_TO`	Represents professional musical ties (e.g., band members)
		`PERSONALLY_RELATED_TO`	Models personal or biographical relationships (e.g., family, marriage)
		`LINKED_TO`	Groups miscellaneous artist connections (e.g., teacher–student, founder)
	LFM	`LAST_FM_MATCH`	Indicates similarity between artists based on listener behavior
	TT	`FOLLOWS`	Indicates that one artist's official TikTok account follows another's
Artist–Song	MBDB	`WORKED_IN`	Directed link from Artist to Track indicating participation in a song.
		`WORKED_BY`	Inverse link from Track to Artist, required for bidirectional message passing.
Tagging	MBDB	`HAS_TAG`	Link from Artist/Track to Tag indicating that the entity is assigned that tag.
		`TAGS`	Inverse link from Tag to Artist or Track.

As a result, we compound an instantiation of the HKG comprising 15,808 *artists*, 828,125 *tracks* and 23 *tags*.

3.3 Problem Formulation

At this stage, the objective of the recommender can be cast as an edge prediction task. Specifically, given a HKG, $\mathcal{G}^t_{\mathcal{M}}$, constructed according to the schema introduced in Sec. 3.2 and representing a set of musicians $\mathcal{M}$ within a time interval t, the goal is to learn a mapping function $\mathcal{P}$ such that

$$\mathcal{P}(\mathcal{G}^t_{\mathcal{M}}) \rightarrow \mathcal{C}^{\Delta}_{\mathcal{M}} \tag{1}$$

where $\mathcal{C}^{\Delta}_{\mathcal{M}}$ denotes the set of collaborations (`collab_with` relationships) established among musicians in $\mathcal{M}$ during the future time window $(t, t + \Delta]$.

3.4 Description of the Graph Neural Network Architecture

To realize the mapping function defined previously, we employed the GNN architecture depicted in Fig. 2. The model operates on the HKG $\mathcal{G}^t_{\mathcal{M}}$ and processes it through a four-layer pipeline specifically designed to handle the diversity of node and relation types present in the graph.

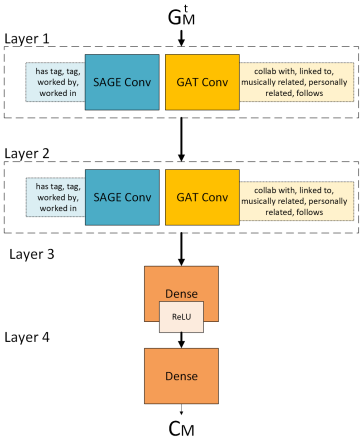


Fig. 2. Layer structure of the GNN for the recommendation task.

The architecture is composed of two heterogeneous graph convolution layers followed by a two-layer fully connected network that refines the learned artist representations. Each convolutional layer applies different operators depending on the relation type, enabling the model to capture both structural and semantic heterogeneity.

For artist-to-artist relationships (e.g., `collab with`, `linked to`, `follows`), a Graph Attention mechanism [19] is employed. This approach assigns adaptive importance weights to neighboring nodes via a multi-head attention scheme,

allowing the model to focus on the most relevant connections when computing node embeddings (ε_v):

$$\varepsilon_v = \sigma\left(\frac{1}{K}\sum_{k=1}^{K}\sum_{u\in\mathcal{N}_v}\alpha_{vu}^{k}W^{k}\varepsilon_u\right) \tag{2}$$

where α_{vu}^{k} represents the attention coefficient of neighbor u to node v in the k-th attention head, and W^k denotes the corresponding weight matrix.

For the remaining heterogeneous relations involving different node types, a GraphSAGE-based aggregation strategy [7] is used. This operator updates node representations $h_v^{k,\zeta}$ by combining their features with those of their neighbors:

$$h_v^{k,\zeta} = \sigma\left(W_\zeta \cdot \text{MEAN}\left(\{h_v^{k-1,\zeta}\} \cup \{h_u^{k-1,\zeta} \mid u \in \mathcal{N}_v^{\zeta}\}\right)\right) \tag{3}$$

where $\mathcal{N}_v^{\zeta}$ denotes the set of neighboring nodes of v connected via edges of type ζ, W_ζ is the weight matrix and σ represents a non-linear activation function.

After convolution, the resulting embeddings are processed by two dense layers with a ReLU activation, followed by L2 normalization to ensure unit-length representations. This normalization improves numerical stability and supports similarity-based computations.

The final output of the GNN is an embedding matrix $\mathcal{E}_{\mathcal{M}}^{\Delta} \in \mathbb{R}^{|\mathcal{M}|\times\rho}$, where each artist is represented by a ρ-dimensional vector. From this, a predicted collaboration matrix is obtained as:

$$\hat{\mathcal{C}}_{\mathcal{M}}^{\Delta} = \text{sigmoid}\left(\mathcal{E}_{\mathcal{M}}^{\Delta}(\mathcal{E}_{\mathcal{M}}^{\Delta})^{T}\right) \tag{4}$$

This matrix $\hat{\mathcal{C}}_{\mathcal{M}}^{\Delta}$ encodes the likelihood of collaboration between artist pairs, forming the basis for the recommendation output.

4 Evaluation of the Proposal

This section presents our evaluation approach which focuses on how the integration of TT data allowed the system to provide more more accurate collaboration recommendations.

4.1 Evaluation Methodology

In order to properly compare the improvement level provided by the TT data integration, we trained two different GNN instances following the architecture described in sec. 3.4 namely, $\text{GNN}_{mb,lfm,tt}$ and $\text{GNN}_{mb,lfm}$. The first one considers all the edges and node attributes from MBDB, LFM and TT described in sec. 3.2 whereas $\text{GNN}_{mb,lfm}$ includes neither the `follows` edge nor the artist's attributes extracted from TT (Table 1).

Furthermore, we trained the two GNN models using three different samples of the HKG based on the artists' `popularity` attribute. This attribute was computed as the Principal Component Analysis (PCA) of the `listeners`

and `playcount` attributes (Table 1). Specifically, each sample comprised artists (and their associate tracks, tags and edges) within a given quantile of this new attribute: 10, 50, and 100%. This approach allowed us to evaluate the recommender's accuracy when identifying matches among artists spanning different popularity ranges. Moreover, we considered three different time horizons, $\Delta = 1, 3$, and 5 years, to assess the suitability of the proposal for recommending long-term collaborations. We also used a 80-20% split for the train and evaluation datasets.

Last, the performance of the proposed model is assessed using standard evaluation metrics, including Precision, Recall, F1-score, Accuracy (ACC), and the Area Under the ROC Curve (AUC).

4.2 Results Discussion

Table 4 shows the metric scores considering the evaluation parameters described above. This table confirms that adding TikTok information consistently improves the quality of the collaboration recommender. Across all popularity quantiles and time horizons, $GNN_{mb,lfm,tt}$ outperforms $GNN_{mb,lfm}$ in Accuracy, Precision, F1-score and AUC, which supports the idea that the TikTok layer contributes complementary signals rather than redundant information. This result is coherent with the role of TikTok in the proposed HKG: unlike MusicBrainz and Last.fm, which mainly capture artistic and listening-based relations, TikTok adds explicit social ties through the *follows* edges and artist activity features, helping the model reflect broader social proximity among musicians.

The improvement is especially clear for the smaller popularity samples, where the collaboration problem is more difficult because fewer listening interactions are available. For the 10% quantile, the proposed model reaches very strong results for every horizon, such as an F1-score of 0.9778 and an AUC of 0.9947 for $\Delta = 1$, while the model without TikTok remains clearly behind, with 0.9011 and 0.7991, respectively. Similar gains appear for $\Delta = 3$ and $\Delta = 5$, which indicates that the TikTok data helps the recommender not only in the short term but also when predicting collaborations further in advance.

A closer look at the metrics also shows that the TikTok-enhanced model usually increases Precision without sacrificing Recall, leading to a more balanced and reliable recommendation output. This is important because, in this task, recommending a plausible collaboration is more useful than simply finding many possible links. In that sense, TikTok appears to help the model filter better candidates by capturing social affinity, shared audience exposure, and cross-genre proximity, which are all relevant in real collaboration scenarios. Overall, the results support the main hypothesis of the paper: integrating TikTok data strengthens the graph representation and leads to more accurate collaboration recommendations.

Table 4. Results of $\text{GNN}_{mb,lfm,tt}$ and $\text{GNN}_{mb,lfm}$ based on the selected popularity quantile of musicians and the time horizon Δ (in years). The best value for each (metric, Δ, Pop. Quantile) tuple is shown in bold.

Δ	Pop. Quantile	Model	ACC	Precision	Recall	F1-score	AUC
1	100%	$\text{GNN}_{mb,lfm,tt}$	**0.9185**	**0.9309**	0.9042	**0.9173**	**0.9627**
		$\text{GNN}_{mb,lfm}$	0.9053	0.8881	**0.9275**	0.9074	0.9508
	50%	$\text{GNN}_{mb,lfm,tt}$	**0.9327**	**0.9207**	0.9469	**0.9336**	**0.9747**
		$\text{GNN}_{mb,lfm}$	0.9103	0.8612	**0.9783**	0.9160	0.9059
	10%	$\text{GNN}_{mb,lfm,tt}$	**0.9773**	**0.9566**	**1.0000**	**0.9778**	**0.9947**
		$\text{GNN}_{mb,lfm}$	0.8915	0.8276	0.9890	0.9011	0.7991
3	100%	$\text{GNN}_{mb,lfm,tt}$	**0.9074**	**0.9075**	0.9072	**0.9074**	**0.9620**
		$\text{GNN}_{mb,lfm}$	0.8695	0.8145	**0.9570**	0.8800	0.9230
	50%	$\text{GNN}_{mb,lfm,tt}$	**0.9539**	**0.9408**	0.9688	**0.9546**	**0.9876**
		$\text{GNN}_{mb,lfm}$	0.9267	0.8831	**0.9837**	0.9307	0.9138
	10%	$\text{GNN}_{mb,lfm,tt}$	**0.9719**	**0.9467**	**1.0000**	**0.9726**	**0.9894**
		$\text{GNN}_{mb,lfm}$	0.8806	0.8089	0.9966	0.8930	0.8021
5	100%	$\text{GNN}_{mb,lfm,tt}$	**0.9099**	**0.9025**	0.9190	**0.9107**	**0.9665**
		$\text{GNN}_{mb,lfm}$	0.8772	0.8150	**0.9760**	0.8883	0.8971
	50%	$\text{GNN}_{mb,lfm,tt}$	**0.9385**	**0.9193**	0.9614	**0.9399**	**0.9808**
		$\text{GNN}_{mb,lfm}$	0.9189	0.8657	**0.9917**	0.9244	0.9144
	10%	$\text{GNN}_{mb,lfm,tt}$	**0.9622**	**0.9300**	**1.0000**	**0.9635**	**0.9768**
		$\text{GNN}_{mb,lfm}$	0.8819	0.8109	0.9961	0.8940	0.7976

4.3 Recommendation examples

Table 5 illustrates representative collaboration recommendations generated by the system for different time horizons. The results include both same-genre pairs (e.g., Bruno Mars and Lady Gaga) and cross-genre collaborations (e.g., Taylor Swift and HAIM), reflecting the model's ability to capture stylistic and social compatibility. These examples show that the system produces realistic and diverse suggestions aligned with actual successful collaborations.

Table 5. Examples of musical collaborations suggested by MUSYNERGY for different time horizons Δ.

Δ	Artist A	Primary genre A	Artist B	Primary genre B	Track title	Track Genres
1	Bruno Mars	Pop	Lady Gaga	Pop	Die With A Smile	Pop, Rock
	ROSÉ	Pop	Bruno Mars	Pop	APT.	Rock, Pop
3	KAROL G	Latin	Shakira	Pop	TQG	Pop, Latin Countries, Latin
	Ice Spice	Hip-Hop	PinkPantheress	Pop	Boy's a Liar Pt. 2	Pop, Electronic, Hip-Hop
5	Megan Thee Stallion	Hip-Hop	Cardi B	Hip-Hop	WAP	Pop, Hip-Hop
	Taylor Swift	Pop	HAIM	Rock	No Body, No Crime	Rock, Pop, Country, English Countries

5 Conclusions and Future Work

This paper has presented a graph-based framework for recommending musical collaborations by integrating data from MusicBrainz, Last.fm, and TikTok. The

results validate the hypothesis stated in the introduction, showing that the inclusion of TikTok-derived social connections improves the performance of the collaboration prediction task. The proposed model consistently outperforms the version without TikTok across all evaluation settings, confirming that social ties provide complementary and valuable information.

These findings highlight the relevance of TikTok as a source of contextual data, capturing social proximity and audience dynamics that are not reflected in traditional music datasets. This is particularly beneficial for less popular artists, where additional signals help improve recommendation quality. Overall, the approach demonstrates that combining artistic and social information leads to more accurate and realistic collaboration suggestions.

As future work, we plan to incorporate additional contextual information to further enrich the model. In particular, analyzing song lyrics may provide useful semantic insights into artistic compatibility. Other extensions include integrating audio features and additional temporal or social signals.

Source Code Respository The source code of the project can be found at https://github.com/aleferu/musynergy_tiktok

Acknowledgments. This work was supported by Agencia Estatal de Investigación (AEI) of Spain under Grant PLEC2021-007609.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bertram, N., Dunkel, J., Hermoso, R.: I am ALL EARS: Using open data and knowledge graph embeddings for music recommendations. Expert Systems with Applications **229**, 120347 (2023)
2. Bevec, M., Tkalčič, M., Pesek, M.: Hybrid music recommendation with graph neural networks. User Modeling and User-Adapted Interaction **34**, 1891–1928 (2024)
3. Chen, J., He, L., Liu, H., Yang, Y., Bi, X.: Background music recommendation on short video sharing platforms. Information Systems Research **35**(4), 1890–1908 (2024)
4. Di Francesco, A.G., Giampietro, G., Spinelli, I., Comminiello, D.: GATSY: Graph attention network for music artist similarity. In: 2025 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE (2025)
5. Fan, W., Ma, Y., Li, Q., He, Y., Zhao, E., Tang, J., Yin, D.: Graph neural networks for social recommendation. In: Proc. of the World Wide Web Conference (WWW) (2019)
6. Fernandez-Sanchez, A., Navarro, P.J., Terroso-Saenz, F.: MUSYNERGY: A framework for music collaboration discovery based on neural networks and graph analysis. Entertainment Computing **55**, 101033 (2025)
7. Hamilton, W.L., Ying, R., Leskovec, J.: Inductive representation learning on large graphs. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 1025–1035 (2017)

8. Hayat, T.: Creative webs: decoding the creativity potential of twitter followers. Frontiers in Psychology **15 - 2024** (2024). https://doi.org/10.3389/fpsyg.2024.1294838
9. Kaye, D.B.V.: Please Duet This: Collaborative music making in lockdown on TikTok. Networking Knowledge: Journal of the MeCCSA Postgraduate Network **15**(1) (Feb 2022)
10. Korzeniowski, F., Oramas, S., Gouyon, F.: Artist similarity with graph neural networks. In: Proc. of the 22nd Int. Society for Music Information Retrieval Conf. (ISMIR) (2021)
11. Lee, M., Kim, H.J.: A collaborative filtering model incorporating media promotions and users' variety-seeking tendencies in the digital music market. Decision Support Systems **174**, 114022 (2023). https://doi.org/https://doi.org/10.1016/j.dss.2023.114022
12. Lin, G., Meng, Z., Wang, D., Long, Q., Zhou, Y., Xiao, M.: GUME: Graphs and user modalities enhancement for long-tail multimodal recommendation. In: Proc. of the 33rd ACM Int. Conf. on Information & Knowledge Management (CIKM) (2024)
13. Liu, X., Yang, Z., Cheng, J.: Music recommendation algorithms based on knowledge graph and multi-task feature learning. Scientific Reports **14**, 2055 (2024). https://doi.org/10.1038/s41598-024-52463-z
14. O'Toole, K.: Collaborative creativity in TikTok music duets. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. pp. 1–16 (2023)
15. Schneider, L., Meyer, A., Rossi, P.: Come together, right now? an empirical study of collaborations in the music industry. Management Science **69**(11), 12345–12367 (2023). https://doi.org/10.1287/mnsc.2023.4743
16. Settles, B., Dow, S.: Let's get together: the formation and success of online creative collaborations. In: Proceedings of the SIGCHI conference on human factors in computing systems. pp. 2009–2018 (2013)
17. Šuput, B.R.: TikTok and sound: Changing the ways of creating, promoting, distributing and listening to music. INSAM Journal of Contemporary Music, Art and Technology (9), 51–73 (2022)
18. Terroso-Saenz, F., Soto, J., Muñoz, A., Roose, P.: PRESTO: a recommender of musical collaborations based on heterogeneous graph neural networks. International Journal of Interactive Multimedia and Artificial Intelligence **9**(6), 28–37 (2026)
19. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. In: International Conference on Learning Representations (ICLR) (2018)
20. Vizcaíno-Verdú, A., De-Casas-Moreno, P., Tirocchi, S., et al.: Online prosumer convergence: Listening, creating and sharing music on YouTube and TikTok. Communication & Society **36**(1), 151–166 (2023)
21. Wei, Y., Wang, X., Nie, L., He, X., Hong, R., Chua, T.: MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In: Proc. of the 27th ACM Int. Conf. on Multimedia (MM) (2019)
22. Zhou, X., Zhou, H., Liu, Y., Zeng, Z., Miao, C., Wang, P.: Bootstrap latent representations for multi-modal recommendation. In: Proc. of the Web Conference (WWW) (2023)