



HAL
open science

Of Spiky SVDs and Music Recommendation

Darius Afchar, Romain Hennequin, Vincent Guigue

► **To cite this version:**

Darius Afchar, Romain Hennequin, Vincent Guigue. Of Spiky SVDs and Music Recommendation. RecSys, ACM, pp.926-932, 2023, 10.1145/3604915.3608850 . hal-04284297

HAL Id: hal-04284297

<https://hal.science/hal-04284297>

Submitted on 14 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Of Spiky SVDs and Music Recommendation

DARIUS AFCHAR and ROMAIN HENNEQUIN, Deezer Research, France

VINCENT GUIGUE, AgroParisTech, MLIA – Sorbonne Universités, France

The truncated singular value decomposition is a widely used methodology in music recommendation for direct similar-item retrieval and downstream tasks embedding musical items. This paper investigates a curious effect that we show naturally occurring on many recommendation datasets: spiking formations in the embedding space. We first propose a metric to quantify this spiking organization's strength, then mathematically prove its origin tied to underlying communities of items of varying internal popularity. With this new-found theoretical understanding, we finally open the topic with an industrial use case of estimating how music embeddings' top-k similar items will change over time under the addition of data.

ACM Reference Format:

Darius Afchar, Romain Hennequin, and Vincent Guigue. 2023. Of Spiky SVDs and Music Recommendation. In . ACM, New York, NY, USA, 10 pages. <https://doi.org/TBD>

1 INTRODUCTION

There is no unique definition of music recommendation, but rather a range of tasks that fall under this name: track sequence recommendation, context-aware recommendation, playlist continuation or generation, similar track, artist, or playlist retrieval [43]. These settings represent the many use cases of recommendation found in the wild (e.g., in a streaming service). Despite this proteiformity, many recommenders leverage a model of similar item retrieval as a basis for their computation. A standard methodology is to rely on some high-dimensional representations of the musical items to recommend. That is why we see more and more work aiming to find general-purpose representations of music tracks (e.g., self-supervision on music spectrograms [47]) to be used in various downstream recommendation scenarios. In the realm of collaborative filtering, the literature is dominated by matrix factorization techniques to create such representations. Compared to other research fields of machine learning, it has been reported in recent years that modern approaches based on neural networks did not bring significant benefit over simpler baselines of factorization and also suffered many reproducibility issues [15, 16, 39, 40]. In this context, the truncated Singular Value Decomposition (SVD) is still a strong contender for factorization. It has demonstrated consistent performances across a variety of datasets [12]. SVD also benefits from rich theoretical literature and has been implemented in efficient ways suitable to an industrial scale (e.g., distributed RSVD [3, 11, 23]). That is why SVD is still widely used in music recommendation.

Nevertheless, despite not requiring much hyperparameter tuning, some blind spots still exist in using SVD as embeddings [32]. For instance, after approximating a dataset M with the low-rank decomposition $U\Sigma V^*$, should we use U , $U\Sigma$, or $U\Sigma^P$ as a basis of similarity? Should the dot product or cosine similarity be used? Many practices for item retrieval rely on experimental results but have yet to find theoretical justifications. In this paper, we submit the reader with a curious geometry that tends to emerge when applying SVD to multiple recommendation datasets: **spikes**. We observe that embedding vectors tend to self-organize along lines that pass through the origin (see Figure 1). This effect

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

Manuscript submitted to ACM

is understudied and goes unnoticed when using the cosine distance for similarity since the normalization squashes all points on a hypersphere [32]. However, we show that leveraging the norm of vectors may be more insightful than it first appears and that they should not be discarded. In detail, we prove that spikes represent communities, and embeddings' norms stem from their varying importance within that community (*i.e.*, intra-popularity). These two latter features are notably prevalent in music data [13, 26] and substantiate the need for theoretical support in this context. As an opening to our results, we show that the norm is strongly indicative of the stability of embeddings, which has industrial implications on how music track representations evolve through time. To our knowledge, this spiking behavior has yet to be leveraged for recommendation before.

The paper proceeds as follows: we first propose a unified metric to quantify this spiking effect, then we mathematically prove the equivalence of spiking geometry with the presence of communities and intra-popularity within each community. We discuss and open our formalization with a practical case of evaluating the stability of music embeddings.

All experiments may be reproduced via our code repository at https://github.com/deezer/spiky_svd.

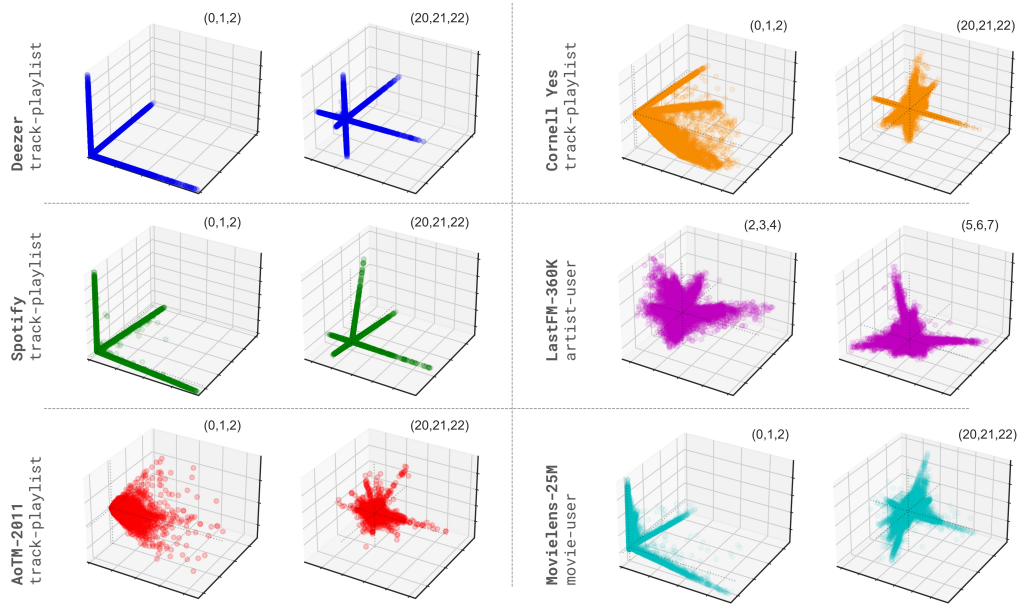


Fig. 1. Naturally emerging "spikes" formations across several independently computed SVDs. Since the embedding vectors are high-dimensional, we can only display some projections of the points on slices of dimensions (indicated above).

2 GEOMETRY OF THE SPIKES

We present and compute SVDs on several datasets and show the emergence of spike formations. We then propose establishing a measure of the "spikiness" of a distribution of embeddings, to quantify the strength of this effect.

2.1 Background on the SVD

We first provide the reader with a brief theoretical background. The *Truncated Singular Value Decomposition* (SVD¹) is a factorization technique that aims to approximate a matrix $M \in \mathbb{R}^{n \times m}$ as the product of three matrices $U \in \mathbb{R}^{n \times f}$,

¹A full-rank decomposition is, of course, intractable to use since it would create embeddings in millions of dimensions. So we abuse notation and drop the T in the acronym. As a reference, the truncated method can also be found under the designation *pureSVD* in the recommenders community.

$\Sigma \in \text{diag}(\mathbb{R}_+^f)$, and $V \in \mathbb{R}^{m \times f}$, such that $\|M - \hat{M}\|_F^2$ is minimal, with $\hat{M} = U\Sigma V^*$ and U and V are orthonormal by columns (i.e., $U^*U = V^*V = I_f$). \hat{M} is a matrix of rank f , with f chosen such that $f \ll \min(n, m)$. \hat{M} is thus said to be a low-rank approximation of M . The "truncated" term refers to the optimal solution for \hat{M} where only the f largest singular values and eigenvectors of M are retained for the decomposition (Eckart-Young-Mirsky theorem [14]).

In the context of music collaborative filtering, several remarks can be made. M is often a large sparse matrix (n and m are frequently counted in millions) and often either denotes track-user, artist-user, or track-playlist interactions. M is often positive, and sometimes additionally binary (e.g., with the formalism of implicit feedbacks [38]). Since the goal is often to leverage \hat{M} for similar item retrieval rather than rebuilding M , the term V is often ignored, and only the similarities between each row of U or $U\Sigma$ are considered. The close-form solution to the SVD is such that U is the dominant eigenvectors of MM^* . Therefore, it does not matter whether we compute the SVD on M or MM^* to retrieve U . With an abuse of notations, we identify M to its symmetrized version MM^* since it does not impact the resulting U .

2.2 Datasets

In order to demonstrate that spiking SVDs are not tied to a particular setting, we use six recommendation datasets: Spotify Million Playlist [45] (Spotify), The Art of the Mix [33] (AoTM-2011), Cornell’s playlists [9], LastFM-360K [7]. We also include the classic MovieLens-25M dataset [24] to hint at an extension of our result to other settings, and a private dataset from *Deezer* containing twelve SVDs computed over one year in an industrial context. This latter dataset will be further exploited in section 4 to evaluate the stability of embedding representations.

We compute SVDs using the same setting for all six datasets. We use the symmetrized version MM^* of the matrix. To curb popularity biases, denoise, and sparsify the matrix, we filter the top- k highest non-null interactions per item. As commonly done in information retrieval, we apply a *positive pointwise mutual information* (PPMI) normalization to the resulting matrix, which further helps with popularity biases and retrieval performances [32] and has been shown to be equivalent to modern skip-gram formulations with negative sampling [31], which is common for music embeddings [3, 6, 10]. We choose $f = 128$ for the decomposition and take $E = U\Sigma$ as the resulting embeddings². The rows of $E = (e_1, \dots, e_n)$ provide a latent representation for each item.

The obtained SVDs are displayed in Figure 1. Since we cannot lay out all 128 dimensions in a single figure, we report curious readers to our repository for a detailed visualization. As a general comment, it seems that the spikiness of embeddings is stronger on dominant dimensions (i.e., eigenvectors associated with the higher singular values and hence the first indexes) and tends to degrade to a blurry cloud of points in the higher indexes. This seems to align with the literature on robust SVD [46] that dominant components first reconstruct underlying structures and then reconstruct noisy perturbations. The *Deezer* and *Spotify* datasets exhibit the cleanest spikes overall.

2.3 Spikiness metric

We could not find an off-the-shelf definition of spikes of points that worked in high-dimension³. In our context, it seems suited to study the embeddings with the highest norms as many candidate peaks for spikes. We can then measure whether many other points with a lower norm are collinear to them⁴. If a small number of peaks are collinear to most of the rest of the distribution, it means that the distribution is spiky overall.

²Any other choice of normalization as $U\Sigma^p$ with $p \in [0, 1]$ simply scales the embedding space differently but does not impact the presence of spikes.

³The closest applicable notion is probably the measure of *kurtosis*. We have tried several experiments with this statistic. However, it turned out to be very sensitive to hyper-parameter choice and does not have a straightforward extension to multi-dimensional variables.

⁴Note that the spikes all seem to pass through the origin.

Table 1. Estimated Spk on several datasets: needed ratio of points that are peaks of spikes to capture $\rho = 50\%$ of the total distribution of n embeddings, with $\cos\theta = 0.9$. The parameter $@f$ denotes several choices of truncation for the SVD. We add a Gaussian sampling as a reference on the improbability of finding collinear vectors in \mathbb{R}^f – even with a θ angle tolerance.

Dataset	Deezer	Spotify	AoTM-2011	Cornell	LastFM-360K	Movielens-25M	$\mathcal{N}(0, 1)$
Spk@128	0.04%	0.14%	18.1%	4.55%	2.42%	7.09%	50% = ρ
Spk@64	150/n < 0.01%	228/n < 0.03%	11.4%	1.59%	0.49%	3.05%	50%
Spk@32	50/n < 0.01%	55/n < 0.01%	2.31%	0.49%	0.14%	1.12%	50%
Spk@16	23/n < 0.01%	23/n < 0.01%	0.93%	0.09%	0.06%	0.02%	49.4%

Since the distributions are noisy, we must instead rely on approximate collinearity. We set some thresholds: θ and ρ . We then denote by e^* a given embedding with a high norm and say that any other vector e belongs to the spike if $\cos(e^*, e) > \cos(\theta)$. We iterate for candidate peaks e^* in descending order of norm until we have captured a ratio ρ of points in the distribution. This can be seen as a greedy heuristic to the partial set cover problem⁵. The obtained number of spikes is divided by n – the total number of embeddings – and denoted Spk. We have $0 < \text{Spk} \leq \rho$, where the lower bound occurs when all vectors form a single spike, and the upper bound means that no two points are collinear.

The results are given in Table 1 for $\cos\theta = 0.9$ and $\rho = 50\%$. Though this measure is computed with a heuristic and constitutes an upper bound to possibly more optimal choices of peaks to cover the distribution, the obtained results seem consistent with our intuitions on the spikiness of each SVD. The *Deezer* and *Spotify* come up as the most spiky. To fully spell out their obtained result, half of the millions of embedding vectors’ directions – that could be arbitrarily anything in 128 dimensions – fall under one of 361 and 1395 spikes. We also confirm that the embeddings are spikier in their first dimensions and then exhibit a sudden increase of Spk when the noise dominates in later dimensions.

We have shown a systematic spiking effect happening over multiple computed SVDs. The result that a significantly low number of spikes may cluster half of the music tracks is highly reminiscent of *spectral clustering* (but with spikes) and hints at the presence of structured groups in the data.

3 MODELING

We next demonstrate a formal equivalence between the spiking structure in the SVD-embedding space and the presence of communities of varying node degrees in the graph associated with \hat{M} . This enables us to get network properties from the embeddings without explicitly building a graph, for which we show a practical impact for recommendations.

3.1 Preliminaries

We define some notations and outline some basic properties to be used in the next section.

3.1.1 Spikes in M . As argued in section 2.1, we may assume M to be symmetric since we only seek to retrieve similarities from the "left side" of M . We denote by $S^n(\mathbb{R}_+)$ the space of symmetric matrices of $\mathbb{R}_+^{n \times n}$. Applying the spectral theorem, M is diagonalizable as PDP^* , with P an orthonormal basis of eigenvectors. We rank the eigenvalues (d_i) and corresponding vectors in decreasing order of absolute value. We then have two useful lemmas:

LEMMA 3.1. *Matrix M has a truncated SVD given by $\hat{M} = U\Sigma V^*$ such that:*

$$U = P_{1..f} \quad \Sigma = |D_{1..f,1..f}| = \text{diag}(\sigma_1, \dots, \sigma_f) \quad V = P_{1..f} \text{sign}(D_{1..f,1..f})$$

⁵... which is notoriously NP-complete in high dimension and thus could not be solved exactly with our number of considered points.

where $|\cdot|$ denote the term-by-term absolute value and $\text{sign}(\cdot)$ the sign function, since the SVD requires Σ to be positive.

PROOF. We identify a desired decomposition. Note that the unitary eigenvectors of P are not uniquely defined, and thus, so does the decomposition. \square

Definition 3.2 (Spike formation). A set of vectors (e_1, \dots, e_k) is said to form a spike – as studied in 2.2 – if we may rewrite them as $(\alpha_1 \vec{s}, \dots, \alpha_k \vec{s})$ with \vec{s} a unitary vector and $(\alpha_1, \dots, \alpha_k) \in \mathbb{R}_+^k$ their distance from the origin.

LEMMA 3.3. Let M have a decomposition $\hat{M} = U\Sigma V^*$. If the rows $(u_{i_1}, \dots, u_{i_k})$ of U belong to a spike, then, we have that the rows (i_1, \dots, i_k) of $U\Sigma^P$ and of V also form a spike:

$$u_i \Sigma^P = \alpha_i (\vec{s} \cdot \Sigma^P) = \alpha_i \vec{s}' \quad v_i = u_i \cdot \text{sign}(D_{..f,..f}) = \alpha_i (\vec{s} \cdot \text{sign}(D_{..f,..f})) = \alpha_i \vec{s}''$$

The spiking effect is insensitive to any choice of embedding as $U\Sigma^P$ or V : only the directions of the spikes change. In particular, the vector \vec{s}'' of V remains unitary. In the rest of this section, we arbitrarily fix the embeddings as $U\Sigma$.

3.1.2 Community graph models. The Stochastic Block Model (SBM) is an important model introduced in [25] as a means to generate random graphs according to an underlying structure in communities. Given a set of n nodes indexed as $\llbracket 1, n \rrbracket = \{1, 2, \dots, n\}$, a partition $C : \llbracket 1, n \rrbracket \rightarrow \llbracket 1, K \rrbracket$ of the nodes into K communities, and a matrix $B \in \mathcal{S}^K([0, 1])$ of edge probabilities, the SBM model samples an adjacency matrix $A \in \mathcal{S}^n(\{0, 1\})$ such that $A_{i,j} = A_{j,i} \sim \mathcal{B}(B_{C(i), C(j)})$, with \mathcal{B} the Bernoulli law. On average, we thus obtain a matrix with a constant edge degree within and between any pair of communities of nodes. If there is not a unique way of defining communities [17, 29], it is usually agreed that the edge density within each community must be superior to that with other ones [37], which translates as $B_{i,i} > B_{i,j}$ for all $i \neq j$. This last property is known as *assortivity*. Music data has been shown to be particularly assortive [26].

The *Degree-Corrected SBM* (DCBM) [28] introduces an additional set of parameters $(\alpha_i)_{i=1}^n \in [0, 1]^n$ allowing the edge probability to vary inside communities. Concisely, the DCBM proposes to sample edges between i and j with a probability $\alpha_i \alpha_j B_{C(i), C(j)}$. An example is depicted in Figure 2. In order to make the parameters identifiable, we may renormalize B to such that $B_{i,i} = 1$. For further properties on the DCBM, we refer readers to [28, 30].

3.2 Equivalence of spikes and degree-varying communities

We may finally spell out our main theorem linking recommendation embeddings and modeling of communities:

THEOREM 3.4. Let $M \in \mathcal{S}^n(\mathbb{R}_+)$ have a SVD $\hat{M} = U\Sigma V^*$ as defined in section 2.1. Let $E = U\Sigma = (e_1, \dots, e_n)$ be embeddings that exhibit K distinct spikes within its rows whose directions are given by $\vec{s}_1, \dots, \vec{s}_K$. If row e_i belongs to a spike, let us define $C(i)$ its corresponding spike index. We additionally denote $n' \leq n$ the total number of points belonging to spikes. For ease of notation, we rearrange E and V such that the n' first rows belong to spikes. Then, the submatrix $\tilde{M} = |\hat{M}_{1..n', 1..n'}|$ is proportional to the mean value of a DCBM.

PROOF. For all $i \in \llbracket 1, n' \rrbracket$, $e_i = \alpha_i \vec{s}_{C(i)}$. Similarly, $v_i = \alpha_i \vec{s}_{C(i)}'$, thanks to lemma (3.3). Then, for all $i, j \in \llbracket 1, n' \rrbracket^2$,

$$(\hat{M})_{i,j} = (EV^*)_{i,j} = \alpha_i \alpha_j \langle \vec{s}_{C(i)}, \vec{s}_{C(j)}' \rangle \quad |\hat{M}_{1..n', 1..n'}| = K^2 \cdot \text{DCBM} \left((C(i))_{i=1}^{n'} ; B ; \left(\frac{\alpha_i}{K} \right)_{i=1}^{n'} \right)$$

with $K = \max_i \alpha_i$, $B \in [0, 1]^{K \times K}$ the matrix $B_{k,l} = |\langle \vec{s}_k, \vec{s}_l' \rangle|$, and $\frac{\alpha_i}{K} \in [0, 1]$ as required. \square

In practice, there are good reasons to believe that \hat{M} will be (roughly) non-negative and that the absolute value can be dropped in \tilde{M} . When $f = 1$, the Perron-Frobenius theorem [35] ensures the non-negativity, which is also trivially

true when $f \geq \text{rank}(M)$. In between, we may expect negative values to principally arise from punctual approximation artifacts of null values of M , rather than large negative blocks having $\langle \vec{s}_k, \vec{s}_l' \rangle < 0$.




Fig. 2. Illustration of theorem (3.4) with a synthetic matrix with $n = 1000$ and $n' = 3 \times 200$ points forming three spikes. We highlight the three direction vectors $\vec{s}_{\{1,2,3\}}$ in E and their corresponding communities $C_{\{1,2,3\}}$ in \hat{M} (reindexed for clarity).

Interpretation. We have shown that the sub-matrix $\hat{M}_{1..n', 1..n'}$ follows a structure in communities given by a DCBM, as illustrated in Figure 2 with a synthetic example. If we think geometrically, the DCBM is merely a matrix with constant blocks that is multiplied element-wise by the rank-1 matrix $[\alpha_1, \dots, \alpha_{n'}]^* [\alpha_1, \dots, \alpha_{n'}]$. Our interpretation of this theorem is the following. Since M is often sparse (*c.f.*, section 2.2), rebuilding as many communities as possible with rank-1 matrices is intuitively more optimal to minimize $\|M - \hat{M}\|_F^2$ than "spending" many eigenvectors on a single community to fine-tune it with a matrix of higher rank. The widespread emergence of spike formations we uncovered across multiple datasets would thus stem from this general underlying prioritization of rank-1 approximation of communities. This coincidentally results in a DCBM with convenient theoretical properties that we can now use. We leave the detection and analysis of eventual rank > 1 approximation out of the scope of this paper.

Related work. We found theorem (3.4) to be related to a result from community detection that the diagonalization of DCBM matrices leads to the appearance of spikes [27, 30]. Let us stress the difference. We prove the reciprocal (*i.e.*, spikes \Rightarrow DCBM), which is novel and in a more general case since a portion of M is noisy ($n' < n$). In spirit also, the former literature aim to correct this effect – seen as noise – while we instead try to exploit it to gain knowledge on \hat{M} .

3.3 Practical advice

The duality of spike formations with DCBM communities has many convenient properties for recommendation. We highlight some corollaries of our theorem that provide practical advice when using embeddings computed with a SVD.

3.3.1 Fast degree estimation. As a first remark, the $(\alpha)_i$ parameters of our theorem's DCBM may be easily retrieved as the norm $\|e_i\|_2$ of embeddings that belong to a spike. Access to the degree of a node of the denoised matrix \hat{M} without explicitly building a graph may open doors for many benefits. For instance, the topology of a graph may be estimated from its degree distribution (*e.g.*, whether the graph is scale-free), which has been suggested to be indicative of a rather social or information nature of the music graphs [5]. Degree centrality has also been related to popularity biases in music data [8, 44] or used to create hierarchies [1]. More generally, the degree of a node is a strong predictor of its robustness to change in a graph [19, 21]. In section 4, we will further exploit such properties of stability.

3.3.2 Cosine similarity or dot product? The literature on community detection often seeks to curb the effect of the degree heterogeneity of the DCBM (e.g., [18, 28, 36]). One popular method is to normalize the eigenvectors with a spherical clustering [30]. Normalizing on a hypersphere turns out to correspond to the widespread use of cosine similarity in recommendation settings. It effectively removes the $(\alpha)_i$ terms, turning the DCBM into a simple SBM and thus attributing a unique value to all the items of a given community (the cosine similarity between their community direction \vec{s}_i and that of the reference embedding to retrieve items from). From the recommendation lens, both unpopular and popular items will be reweighted with an equal relevance value. This may be useful to improve *discovery* objectives – e.g., improving coverage and serendipity metrics [20], but conversely promotes filter bubbles by design [34, 48].

On the flip side of the coin, measuring the internal popularity of a node within a community – i.e., its “importance” [18] – may be helpful to navigate between communities while displaying their most representative items. This setting is, in particular, more suited for artist similarity and cold-start applications [41, 42]. There, a dot product will always return the item with the highest norm of the community as top-similar, but may also merge popular items from nearby communities in the retrieved items. The dot product thus seems more suited to propose a *panopticon view* of the many available items, but with the adverse risk of being biased by popularity and thus making long-tail items harder to access.

As always, this choice of hyper-parameter depends on the task at hand and desired criteria to optimize.

3.3.3 Tuning of f . The Spk metric we proposed in 2.3 may constitute a promising *offline metric* on optimally choosing the truncation parameter f in the SVD, i.e., finding when recommendation performances are saturated despite adding more embedding dimensions. In the light of theorem (3.4), we can now reinterpret our metric as an approximation of K , where we have implicitly assumed that $\hat{n}' = \rho n$ points should be considered as being modeled by a DCBM. We have additionally used the empirical heuristic that noise was prevalent in the lower norm values of E . When the chosen \hat{n}' comes close to the ground-truth n' , and the hypothesis starts to fail, noise starts to permeate, and the number of spikes grows linearly with the number of points (similar to the baseline Normal distribution where collinear points are rare).

In Table 1, this effect was striking between *Deezer* and *Spotify*, where the Spk metric was proportional between the two datasets for low values of f , then exploded at $f = 128$ for the latter. We conclude that $f = 64$ for Spotify and $f = 128$ for Deezer constitutes an appropriate choice to capture clean communities in the embeddings. This metric makes the assumption to treat non-spikes as uninformative noise, which dismisses potential higher-ranking approximations that may hold further information on the data. However, we leave this last aspect for future work.

4 OPENING: INDUSTRIAL USE-CASE

We open our results with an industrial use case to illustrate the few previously highlighted properties. Specifically, designing an industrial recommender system involves many challenges (e.g., scalability constraints) that can render adopting a complex machine learning model cumbersome [2]. Insights like the one we propose can thus be valuable to gain knowledge to better tune and evaluate recommenders leveraging simple item similarity models, such as the SVD, without changing the computation pipeline. We had access to a history of size twelve of music track embeddings from *Deezer*, spanning between May 2022 and May 2023, computed on user data with a truncated SVD as described in 2.2. These SVDs contain 1.5 million common music tracks and have $f = 128$.

Music data is particularly prone to exhibiting community effects – linked to its many historical and cultural groundings [13, 26], but it is also ever-changing. New trends emerge every day (e.g., *lo-fi* music, *phonk*), others vanish (e.g., *emo*, *jumpstyle*), or are sometimes revived (e.g., fans of Kate Bush thanks to *Stranger Things*’ OST), or shape one of many




Fig. 3. Stability through time of the top-500 of *Deezer* embeddings starting in May 2022 and compared with eleven successive embeddings within five-week steps until May 2023. We condition the recommendation on a partition of the embedding norms (indicated on the right). Both the cosine similarity and dot product are considered. 95% confidence intervals are displayed.

types of evolution [4]. It is thus natural to wonder how relevant some embeddings computed at a given date will remain in the following months under the change of the underlying data.

In section 3.3, we stressed that the estimated degree is considered a good predictor of the robustness of a node in a graph. We thus propose to link both notions and show that embeddings of higher norms have better stability in time than their lower-norm counterpart. Using the oldest SVD of May 2022 as a reference, we partition the embeddings equally into five according to their norm. We sample 1000 points in each partition and compute their top-500 similar items in each SVDs – using either the cosine similarity or the dot product – among the almost 300.000 track candidates in each partition. We finally compute the mean Jaccard index (IoU) [22] to quantify the relative similarity between the top items of the reference and subsequent compared date.

The results are given in Figure 3. We can see that conditioning our results on several norm ranges reveals strikingly different dynamics that could not have been inferred solely from a mean analysis. The dot product notably displays the most disparate curves where low-norm embeddings quickly become irrelevant while higher-norm ones remain over 80% similar after almost a year. Note that our point is not to state that change is necessarily undesirable in the embeddings, though we could expect well-established music genres to remain stable (*e.g.*, jazz music tracks). Instead, we envision this type of conditioned analysis as a means to better evaluate music recommenders by discriminating between various dynamics of embeddings – and thus some expected associated properties.

5 CONCLUSION

Embeddings are commonly used in music recommendation. Yet, there are mostly so in a black-box manner, and many practices rely on empirical knowledge. In this paper, we have proven that the widespread emergence of spikes in the embeddings of music items was tied to an underlying graph structure of degree-varying communities. Our insights are particularly relevant for music data, where such structure is found naturally due to cultural and historical groundings. Reframing a community detection model through the recommendation lens enabled us to formulate novel practical insights for music embeddings (*e.g.*, cosine or dot product, stability) that are better grounded in theory. We hope to see more work bridging the gap between graph theory and recommendation. Our future work includes the analysis of rank > 1 approximation effects on embeddings.

REFERENCES

- [1] Darius Afchar, Romain Hennequin, and Vincent Guigue. 2022. Learning Unsupervised Hierarchies of Audio Concepts. In *23rd International Society for Music Information Retrieval Conference (ISMIR)*.
- [2] Darius Afchar, Alessandro Melchiorre, Markus Schedl, Romain Hennequin, Elena Epure, and Manuel Moussallam. 2022. Explainability in music recommender systems. *AI Magazine* 43, 2 (2022), 190–208.
- [3] Léa Briand, Guillaume Salha-Galvan, Walid Bendada, Mathieu Morlon, and Viet-Anh Tran. 2021. A semi-personalized system for user cold start recommendation on music streaming apps. In *27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2601–2609.
- [4] Piotr Bródka, Stanisław Saganowski, and Przemysław Kazienko. 2013. GED: the method for group evolution discovery in social networks. *Social Network Analysis and Mining* 3 (2013), 1–14.
- [5] Pedro Cano, Oscar Celma, Markus Koppenberger, and Javier M. Buldú. 2006. Topology of music recommendation networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 16, 1 (01 2006). <https://doi.org/10.1063/1.2137622>
- [6] Hugo Caselles-Dupré, Florian Lesaint, and Jimena Royo-Letelier. 2018. Word2vec applied to recommendation: Hyperparameters matter. In *12th ACM Conference on Recommender Systems (RecSys)*. 352–356.
- [7] Óscar Celma. 2010. *Music Recommendation and Discovery in the Long Tail*. Springer. <http://ocelma.net/MusicRecommendationDataset/lastfm-360K.html>
- [8] Oscar Celma and Pedro Cano. 2008. From hits to niches? or how popular artists can bias music recommendation and discovery. In *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition*. 1–8.
- [9] Shuo Chen, Josh L Moore, Douglas Turnbull, and Thorsten Joachims. 2012. Playlist prediction via metric embedding. In *18th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD)*. 714–722. https://www.cs.cornell.edu/~shuochen/lme/data_page.html
- [10] Zhiyong Cheng, Jialie Shen, Lei Zhu, Mohan S Kankanhalli, and Liqiang Nie. 2017. Exploiting Music Play Sequence for Music Recommendation. In *IJCAI*, Vol. 17. 3654–3660.
- [11] Paul G Constantine and David F Gleich. 2011. Tall and skinny QR factorizations in MapReduce architectures. In *2nd international workshop on MapReduce and its applications*. 43–50.
- [12] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of Recommender Algorithms on Top-n Recommendation Tasks. In *4th ACM Conference on Recommender Systems (RecSys)*. 39–46. <https://doi.org/10.1145/1864708.1864721>
- [13] Ian Cross. 2001. Music, cognition, culture, and evolution. *Annals of the New York Academy of sciences* 930, 1 (2001), 28–42.
- [14] Carl Eckart and Gale Young. 1936. The approximation of one matrix by another of lower rank. *Psychometrika* 1, 3 (01 Sep 1936), 211–218. <https://doi.org/10.1007/BF02288367>
- [15] Maurizio Ferrari Dacrema, Simone Boglio, Paolo Cremonesi, and Dietmar Jannach. 2021. A troubling analysis of reproducibility and progress in recommender systems research. *ACM Transactions on Information Systems (TOIS)* 39, 2 (2021), 1–49.
- [16] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *13th ACM conference on recommender systems (RecSys)*. 101–109.
- [17] Santo Fortunato. 2010. Community detection in graphs. *Physics reports* 486, 3-5 (2010), 75–174.
- [18] Chao Gao, Zongming Ma, Anderson Y. Zhang, and Harrison H. Zhou. 2018. Community Detection in Degree-Corrected Block Models. *The Annals of Statistics* 46, 5 (2018), 2153–2185. <https://www.jstor.org/stable/26542860>
- [19] Johannes Gasteiger, Stefan Weissenberger, and Stephan Günnemann. 2019. Diffusion improves graph learning. *Advances in neural information processing systems (NeurIPS)* 32 (2019).
- [20] Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *4th ACM conference on Recommender systems (RecSys)*. 257–260.
- [21] Simon Geisler, Daniel Zügner, and Stephan Günnemann. 2020. Reliable graph neural networks via robust aggregation. *Advances in Neural Information Processing Systems (NeurIPS)* 33 (2020), 13272–13284.
- [22] Derek Greene, Donal Doyle, and Pádraig Cunningham. 2010. Tracking the evolution of communities in dynamic social networks. In *2010 international conference on advances in social networks analysis and mining*. IEEE, 176–183.
- [23] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* 53, 2 (2011), 217–288.
- [24] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)* 5, 4 (2015), 1–19. <https://grouplens.org/datasets/movielens/25m/>
- [25] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social networks* 5, 2 (1983), 109–137.
- [26] Kurt Jacobson and Mark Sandler. 2009. Musically meaningful or just noise? an analysis of on-line artist networks. In *Computer Music Modeling and Retrieval. Genesis of Meaning in Sound and Music: 5th International Symposium*. Springer, 107–118.
- [27] Jiashun Jin. 2012. Fast network community detection by SCORE. *The Annals of Statistics* 43 (11 2012). <https://doi.org/10.1214/14-AOS1265>
- [28] Brian Karrer and Mark EJ Newman. 2011. Stochastic blockmodels and community structure in networks. *Physical review E* 83, 1 (2011), 016107.
- [29] Sune Lehmann. 2014. *Community Detection, Current and Future Research Trends*. Springer New York, New York, NY, 214–220. https://doi.org/10.1007/978-1-4614-6170-8_27
- [30] Jing Lei and Alessandro Rinaldo. 2015. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics* 43, 1 (2015), 215–237.

- [31] Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems (NIPS)*, Vol. 27.
- [32] Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the association for computational linguistics (ACL)* 3 (2015), 211–225.
- [33] B. McFee and G. R. G. Lanckriet. 2012. Hypergraph models of playlist dialects. In *13th International Symposium for Music Information Retrieval (ISMIR)*. <https://brianmcfree.net/data/aotm2011.html>
- [34] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. penguin UK.
- [35] Oskar Perron. 1907. Zur theorie der matrices. *Math. Ann.* 64, 2 (1907), 248–263.
- [36] Tai Qin and Karl Rohe. 2013. Regularized spectral clustering under the degree-corrected stochastic blockmodel. *Advances in neural information processing systems (NIPS)* 26 (2013).
- [37] Filippo Radicchi, Claudio Castellano, Federico Cecconi, Vittorio Loreto, and Domenico Parisi. 2004. Defining and identifying communities in networks. *Proceedings of the national academy of sciences* 101, 9 (2004), 2658–2663.
- [38] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *25th Conference on Uncertainty in Artificial Intelligence (UAI '09)*. 452–461.
- [39] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. 2020. Neural collaborative filtering vs. matrix factorization revisited. In *14th ACM Conference on Recommender Systems (RecSys)*. 240–248.
- [40] Steffen Rendle, Li Zhang, and Yehuda Koren. 2019. On the Difficulty of Evaluating Baselines: A Study on Recommender Systems. *CoRR* abs/1905.01395 (2019). arXiv:1905.01395
- [41] Guillaume Salha, Stratis Limnios, Romain Hennequin, Viet-Anh Tran, and Michalis Vazirgiannis. 2019. Gravity-inspired graph autoencoders for directed link prediction. In *28th ACM international conference on information and knowledge management (CIKM)*. 589–598.
- [42] Guillaume Salha-Galvan, Romain Hennequin, Benjamin Chapus, Viet-Anh Tran, and Michalis Vazirgiannis. 2021. Cold start similar artists ranking with gravity-inspired graph autoencoders. In *15th ACM Conference on Recommender Systems (RecSys)*. 443–452.
- [43] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. 2018. Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval* 7 (2018), 95–116.
- [44] Tobin South, Matthew Roughan, and Lewis Mitchell. 2020. Popularity and centrality in Spotify networks: critical transitions in eigenvector centrality. *Journal of Complex Networks* 8, 6 (2020), cnaa050.
- [45] Spotify. 2020. *Spotify Million Playlist*. <https://research.atspotify.com/2020/09/the-million-playlist-dataset-remastered/>
- [46] Gilbert W Stewart. 1998. *Perturbation theory for the singular value decomposition*. Technical Report.
- [47] Luyu Wang, Pauline Luc, Yan Wu, Adria Recasens, Lucas Smaira, Andrew Brock, Andrew Jaegle, Jean-Baptiste Alayrac, Sander Dieleman, Joao Carreira, et al. 2022. Towards learning universal audio representations. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 4593–4597.
- [48] Yuan Cao Zhang, Diarmuid Ó Séaghdha, Daniele Quercia, and Tamas Jambor. 2012. Auralist: introducing serendipity into music recommendation. In *5th ACM international conference on Web search and data mining (WSDM)*. 13–22.