

LEVERAGING KNOWLEDGE BASES AND PARALLEL ANNOTATIONS FOR MUSIC GENRE TRANSLATION

Elena V. Epure
Deezer R&D

Anis Khlif
Deezer R&D

Romain Hennequin
Deezer R&D

research@deezer.com

ABSTRACT

Prevalent efforts have been put in automatically inferring genres of musical items. Yet, the proposed solutions often rely on simplifications and fail to address the diversity and subjectivity of music genres. Accounting for these has, though, many benefits for aligning knowledge sources, integrating data and enriching musical items with tags. Here, we choose a new angle for the genre study by seeking to predict what would be the genres of musical items in a target tag system, knowing the genres assigned to them within source tag systems. We call this a translation task and identify three cases: 1) no common annotated corpus between source and target tag systems exists, 2) such a large corpus exists, 3) only few common annotations exist. We propose the related solutions: a knowledge-based translation modeled as taxonomy mapping, a statistical translation modeled with maximum likelihood logistic regression; a hybrid translation modeled with maximum a posteriori logistic regression with priors given by the knowledge-based translation. During evaluation, the solutions fit well the identified cases and the hybrid translation is systematically the most effective w.r.t. multilabel classification metrics. This is a first attempt to unify genre tag systems by leveraging both representation and interpretation diversity.

1. INTRODUCTION

Music genres have been long studied as semantic dimensions of artists and tracks [8]. Rooted in musicology, music experts have mainly undertaken this endeavour. With digitization of music and prevalence of Internet music consumption, online communities have also shown increasing interest in annotating musical items with genres (e.g. creating folksonomies such as Lastfm). In addition, crowd-sourced, web-based encyclopedias that describe and structure music-related knowledge including genres, have been created and openly disseminated [4, 37, 41].

Apart from ontologically describing musical items, genres are also among the most common attributes of tracks, albums and artists to which the users of music streaming

services relate [21]. Users resort to genres to discover music, create playlists, define their profiles, foster interactions with other users, etc. Hence, being able to correctly infer music genres as metadata is central to such tasks.

Music genre is a challenging concept to model and highly subjective. Past studies [11, 18, 33, 36] convey how difficult it is to agree upon shared definitions and interpretations, even for popular genres. People interpret genres differently, influenced by their culture, personal preferences or acquired musicological knowledge [11, 18, 33]. Genre representations within tag systems vary [31] with respect to: the level of detail (how specialized genres can get); the coverage (which genres are considered); the genre interpretation (*pop/rock* could be distinctly defined and interpreted across sources); how genres are related (*blues rock* is a subgenre of *rock*, but not of *blues* in the MuMu dataset [16, 22]). Divergences also result from the spelling variability (e.g. *alternative rock* vs. *alt. rock*).

The research question we address in this work is: given annotations with genre tag systems of multiple sources, how to infer the equivalent annotations within a target tag system? We refer to this as a translation task, but we do not necessarily seek to translate tags between languages.

When relying only on the definition of the sources and target tag systems, this task could be solved using taxonomy mapping [27, 29]. A taxonomy is a classification schema with concepts organized from general to specialized. The goal of taxonomy mapping is to align the concepts of the source and target taxonomies. Related works integrate commercial catalogues [24, 29], align multi-lingual taxonomies [34, 35, 43] or restructure existing taxonomies [26, 27, 38] in supervised or unsupervised manners. Ontology mapping [23] is a similar task, in which additional relation properties and axioms can be exploited.

A solution focused on taxonomy mapping is nonetheless incomplete as it does not consider the application of the taxonomies in practice, which could reveal divergences in genre interpretation. Thus, we hypothesize that a robust translation is built not only on the definitions of genre tag systems, but also on their use for annotations. In accordance with the terminology of the Automatic Machine Translation domain [9], we call a corpus of items jointly annotated by multiple sources a parallel corpus.

The contribution of the current work is a translation system that effectively leverages knowledge-based and statistical methods for genre translation in three cases:

1. A cold-start case, when genre tag systems of the target



and sources are known, but there is no parallel corpus. We address this case with a Knowledge-Based (KB) system based on taxonomy mapping (Section 3).

2. Many parallel annotations are available allowing to learn mappings between genre interpretations (e.g. when some sources use *alternative rock* the target tends to use *alt. rock* and *indie rock*). To deal with this case, we use a simple linear multilabel classifier, namely a logistic regression model trained with Maximum Likelihood (ML) (Section 4.1).
3. The case in-between when less annotations are available and some target tags may be missing in the parallel corpus. We tackle this scenario with an hybrid Bayesian approach that leverages the KB translation as a prior for the logistic regression model trained with Maximum A Posteriori (MAP). This case, presented in Section 4.2, is the most general. Finding an effective solution for it has multiple positive implications for practice.

We release the code of these methods for reproducibility¹.

The Music Information Retrieval (MIR) community has extensively studied the automatic genre annotation of musical items by exploiting the content (e.g. audio, lyrics) [10, 16, 22]. Other genre representations, tackled in [2, 20, 31, 41], create genre graphs from multiple knowledge sources. Yet, to our knowledge, there is no past work translating music genres from one tag system to another (e.g. from Discogs to Wikipedia) by leveraging the diversity of both genre representations and interpretations.

We resort to item annotation to assess the proposed translation methods. To reflect a real-life context [7], we consider a musical item annotated with multiple source tag systems; having multiple labels and not only broad genres such as *rock*, but also very detailed subgenres, which results in predicting among hundreds of possibilities. Lastly, combining multiple tag predictors in a Bayesian framework was done before [10, 39]. However, these works aggregate information from different predictors in the same tag system while we consider several tag systems.

2. NOTATIONS AND PROBLEM FORMULATION

In this work, we denote matrices by bold capital letters, \mathbf{M} ; vectors by bold lower case letters, \mathbf{v} ; the n -th row vector of matrix \mathbf{M} by \mathbf{m}_n ; scalars by italic lower case letters, x ; the coefficient at row i and column j of matrix \mathbf{M} by m_{ij} ; the i -th element of vector \mathbf{v} by v_i . Calligraphic font is used for sets of sets (e.g. \mathcal{S}) and capital letters for sets (e.g. S).

Let \mathcal{D} be a set of tag systems, \mathcal{S} a subset of \mathcal{D} , henceforth referred to as source tag systems, and $T \in \mathcal{D}, T \notin \mathcal{S}$ henceforth referred to as a target tag system. Further, we refer to a tag system as a tag set, but we stress that it may contain broader information such as relations between genres (e.g. taxonomies or ontologies). The research problem we address is: given \mathcal{S}, T and a set of tag annotations (e.g. associated with a given musical item) taken from \mathcal{S} , what would have been the corresponding tag annotations if the tags had been taken from T . We note $S = \cup_{E \in \mathcal{S}} E$ the

union of the source tag systems, and $|S|$ its cardinality.

The approach we adopt consists in defining a translation scoring function $f : \mathcal{P}(S) \rightarrow \mathbb{R}^{|T|}$, where \mathcal{P} denotes partitions over S , that predicts translation scores for every target tag from a set of source tags. Estimating such a scoring function is a standard setting for multilabel classification.

3. KNOWLEDGE-BASED GENRE TRANSLATION

We propose a translation method based on multiple genre taxonomies brought together under a genre graph. Section 3.1 introduces the graph types of concepts and relations and presents the genre taxonomies. In Section 3.2, we show how we create the links between the genre taxonomies using advanced normalization and tokenization. In Section 3.3, we define the translation scoring function f by exploiting the genre graph structure and its relations.

3.1 Building a knowledge-based genre graph

We automatically derive an undirected genre graph by aggregating multiple genre tag systems (e.g. taxonomies, ontologies or social tags), created by either experts or non-experts as in [2, 20, 31, 34]. Its modular design allows to easily integrate new sources through a normalization pipeline that addresses much more variability of genre strings than the existing works [30, 41] (presented in Section 3.2). The knowledge sources used to build the current version of the genre graph are: DBpedia (English, 12443 genres), and Lastfm (327 genres), Tagtraum (296 genres) and Discogs (296 genres)—the taxonomies released in the 2018 MediaVal AcousticBrainz Genre Task [7]. The Discogs genre taxonomy is pre-defined by experts. The Lastfm and Tagtraum genre taxonomies are automatically inferred from social tags with the approach proposed by Schreiber [30], followed by a manual processing [7].

The types of relations between genres vary across sources. In DBpedia, the retrieved types for each genre are: subgenres, origins, aliases—various spellings of the same genre, and derivatives—genres which are influenced by this genre, but could not be considered subgenres. The other knowledge sources contain only subgenre relations.

Each genre tag system becomes a graph by adding a source node that connects all the genre tags as in Figure 1. Then, to connect these decentralized graphs, a normalized graph is produced from all available tags. Each original tag is connected to its normalized form in the normalized graph. The description of how we normalize genres and create the normalized graph is continued in Section 3.2.




Figure 1: Genre graph extract. Dashed edges link the original genre tags to their normalized versions.

¹ <https://github.com/deezer/MusicGenreTranslation>

3.2 Normalizing genre tags

We create a more robust normalization pipeline compared to the related works [29,30,35,41] that, apart from basic tokenization and normalization, also separates words written together (e.g. *poprock* in *pop* and *rock*). The basic tokenization splits tags by non-alphanumeric characters (e.g. "-", "_"). The basic normalization converts tags to lower case and brings tags containing "&", "+" and "n" to the same form (e.g. *d+b*, *drum'n'bass* and *drum and bass*).

For the advanced tokenization, we use a modified trie [12] and a probabilistic tokenization built on Wikipedia unigrams [3]. A trie is a tree data structure that efficiently stores and retrieves strings. Each node has a char and a flag to mark if the path from the root to it forms a word. We modify the way we populate the trie as follows. At first, we sort the tokens obtained from the basic tokenization and normalization, ascendingly by length. Then, we add the tokens of DBpedia with less than l letters directly to the trie². For the others, we attempt to split them using the trie and only the unknown words are added to the trie.

The tokenization using the trie is a recursive greedy algorithm that aims at matching the longest possible words in the trie. If a recursion fails, we explore the path with the next best previous word instead. If we assess the split output as incorrect, meaning that it results in too many short words, in short suffixes, or fails to split a large tag, then we use the probabilistic tokenization.

The probabilistic tokenization uses dynamic programming to find the words best maximizing their probability product. The frequency of each word, assuming that they are independently distributed, is approximated using the Zipf's law [44] to $\frac{1}{n \log(N)}$, where n is the word rank [42] and N is the total number of Wikipedia unigrams. We again assess the split output. Some extra conditions are added besides those presented in the previous paragraph: a Wikipedia split is incorrect if there are single letters as middle words and if no word is already contained in the trie³. If this tokenization fails, we add the token as it is.

Finally, we transform the obtained tokens in nodes in the normalized graph (see Figure 1). There are three types of nodes: 1) normalized composed genres (e.g. *altern rock*, *deep house*), 2) concepts which are words that do not represent genres but are part of the name of multiple genres (e.g. *nu* in *nu jazz* and *nu metal*); 3) concept genres which are standalone genres but can be also part of composed genres (e.g. *punk* in *post punk*). If a genre is tokenized, its tokens are sorted and concatenated becoming a composed genre node as in [30] (e.g. *music rock* in Figure 1). This node is then connected to its concept and concept genre nodes.

3.3 Translating Genres through DBpedia Mapping

Using intermediate mapping spaces such as taxonomies or pivot languages has been explored in past works to match multi-lingual [35, 43], multi-cultural [27] or e-commerce [26, 29] taxonomies. Similar to [27], we use DBpedia, an

ontology derived from Wikipedia infoboxes [19] as it has the highest genre coverage and quite high quality. However, to map a genre to DBpedia genres, we avoid using string similarity as it can be very noisy (e.g. *pop* vs. *bop*). Instead, we leverage genre knowledge to create a mapping strategy as we further present. Most related works rely on the structure of taxonomies for mapping the source and target concepts [24, 27, 29, 35, 43]. Our solution uses structural information too, but differently. Specifically, we use the neighbours of the source and target concepts and the structure of the directed DBpedia graph.

We map each genre of the source and target tag system, to the genres of the DBpedia ontology: $B \in \mathcal{D}$. We assume $B \notin \mathcal{S}$ and $B \neq T$. For each input tag system D , with $D \in \mathcal{S}$ or $D = T$, the output of the mapping is a matrix $\mathbf{Z}^D \in \mathbb{R}^{|D| \times |B|}$, where each row represents the relatedness of a genre tag from D to the DBpedia genres. We compute the mapping matrix \mathbf{Z}^D by applying the following steps for each tag $s \in D$:

1. Normalize s with the process described in Section 3.2 (e.g. *Rock/Pop* becomes *pop rock*).
2. Check if the normalized s equals any normalized genre of B . If true, all entries in \mathbf{Z}_D linked to the DBpedia aliases of the found genres are set to 1 and all others to 0 (e.g. *acid house* is mapped to *Acid_house*, with aliases *Acid_(electronic_music)*, *Warehouse_music*, etc.).
3. If the normalized s is not in B , then map it using its context genres in D : compound s with each parent tag in D and check if the normalized compounded tag equals any normalized genre of B (inspired from [43]). If true, proceed as in Step 1. (e.g. *stoner* has parent *rock* in Lastfm; search by *rock stoner* and map it to *Stoner_rock*).
4. If Steps 2 and 3 are unsuccessful, consider two cases:
 - (a) s is a concept genre as defined in Section 3.2. First, retrieve the DBpedia directed subgraph composed of the nodes which contain the normalized s as a substring in their normalized form. Second, map s to the nodes with the highest in-degree centrality [5] in this subgraph. The intuition is that concept genre nodes are more likely fundamental music genres; hence they tend to have many subgenres or related genres. Third, assign to the selected DBpedia genres and their aliases a score of 1 divided by the number of selected nodes, and to the others 0 (e.g. *rock* does not exist as is in DBpedia. To map it, we retrieve all tags that contain it such as *Punk_rock*, *Art_rock*, *Rock_music*, etc. We observe that *Rock_music* is the most connected node in the subgraph with the genres containing *rock*. As only one node is selected, we assign to it and its aliases a score of 1).
 - (b) s is a composed genre as defined in Section 3.2. First, select from the normalized genres in B those that share the greatest number of words with s . Second, select from this list, the genres with the highest number of shared concept genres—if it is 0, then the initial selection is kept unchanged. Third, assign scores as in Step 4(a).

² DBpedia seeds the trie as it has the highest coverage and we set $l=7$.

³ As we already added to the trie short genre and concept tags from multiple sources, we assume the probability of all words to be new is low.

5. For each genre in B associated to s in Steps 1–4, propagate half of the value of its score to its neighbors in B . The intuition is that parent genres or subgenres could be relevant and sometimes specified by other sources.

For each s not mapped in the previous process, we compute its scores by averaging the rows in \mathbf{Z}^D of its related genres in the input taxonomy D (e.g. for *aor* which is not found in DBpedia, we compute the scores by assigning it the scores obtained for *rock*, its parent genre in Discogs). Finally, the relatedness of a source genre $s \in S$ and a target $t \in T$ is computed using cosine similarity between their corresponding rows $\mathbf{s} = \mathbf{z}_s^S$ and $\mathbf{t} = \mathbf{z}_t^T$ in the mapping matrices, \mathbf{Z}^S and \mathbf{Z}^T . We define $\mathbf{W}^{KB} \in \mathbb{R}^{|T| \times |S|}$ such that $w_{ts}^{KB} = \frac{\mathbf{s}^T \mathbf{t}}{\|\mathbf{s}\|_2 \|\mathbf{t}\|_2}$. The translation scoring function is:

$$f_t(\{s_1, s_2, \dots, s_K\}) = \sum_{k=1}^K w_{ts_k}^{KB} = \mathbf{x}^T \mathbf{w}_t^{KB}, \quad (1)$$

where \mathbf{x} is the binary encoded vector of $\{s_1, \dots, s_K\}$.

4. DATA-INFORMED GENRE TRANSLATION

In this section, we consider that a parallel corpus is available and present two statistical approaches: ML that relies only on annotations (Section 4.1), and MAP that leverages the KB results as a prior knowledge (Section 4.2).

4.1 Maximum Likelihood logistic regression

In statistical approaches to the tag translation task, we seek to train a parametric mapping to model the probability $P(\mathbf{y}|\mathbf{x})$ of having a collection of target tags (encoded as a binary vector $\mathbf{y} \in \{0, 1\}^{|T|}$) given the source tags (encoded as a binary vector $\mathbf{x} \in \{0, 1\}^{|S|}$). We assume the independence of the target tags, and only seek to model the conditional probabilities $P(y_t|\mathbf{x})$. This comes down to training $|T|$ binary classifiers, also known as binary relevance. There are more elaborated settings for doing multilabel classification without the target tag independence assumption. We notably also tested classifiers chain [28], but it did not result in significant improvement over the results presented in Section 5.3, while increasing the system complexity. We propose to implement binary relevance with logistic regression [40]. Logistic regression models the probability of having the t -th target tag y_t given the source tags \mathbf{x} and the parameters of the logistic regression $\theta = \{\mathbf{W}, \mathbf{b}\}$, $\mathbf{W} \in \mathbb{R}^{|T| \times |S|}$, $\mathbf{b} \in \mathbb{R}^{|T|}$; as:

$$P(y_t = 1|\mathbf{x}, \theta) = \sigma(\mathbf{w}_t^T \mathbf{x} + b_t) \quad (2)$$

where $\sigma(x) = \frac{1}{1 + \exp(-x)}$. \mathbf{W} is called the weights matrix and \mathbf{b} the bias. Note that, for the statistical approaches, the scoring function f introduced in Section 2 is defined here as $f(\{s_1, s_2, \dots, s_K\}) = (P(y_1 = 1|\mathbf{x}, \theta), \dots, P(y_{|T|} = 1|\mathbf{x}, \theta))$. To train a logistic regression model we maximize the log-likelihood of the

targets, given the source tags, w.r.t. the parameters θ :

$$\mathcal{L} = \log P(\mathbf{Y}|\mathbf{X}; \theta) = \sum_{n=1}^N \mathbf{y}_n^T \log(\sigma(\mathbf{W}\mathbf{x}_n + \mathbf{b})) + (\mathbf{1} - \mathbf{y}_n)^T \log(\mathbf{1} - \sigma(\mathbf{W}\mathbf{x}_n + \mathbf{b})) \quad (3)$$

where N is the size of the parallel corpus; $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T \in \{0, 1\}^{N \times |S|}$; and $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_N]^T \in \{0, 1\}^{N \times |T|}$. In practice the regularization term $\frac{1}{2} \|\mathbf{W}\|_F^2$ is added to \mathcal{L} in the objective, where $\|\cdot\|_F$ denotes the Frobenius norm on matrices, to limit overfitting.

4.2 A unified translation model

While ML logistic regression can be expected to work well with large amounts of parallel annotations, they will not adapt well to settings where no or little parallel data is available. In a real-life scenario, the size of the parallel corpus can range from zero to tens of thousands of samples, which precludes systematically favoring one or the other. Defining a criterion for when to switch from KB to statistical translation is arduous since this criterion would depend on the number of source and target tags as well as on their distribution. Ideally, we would like to have knowledge-based performances when no parallel data is available, and a smooth way to transition towards more data-abundant settings. This leads us to consider the translation table \mathbf{W}^{KB} given by the KB system as a prior in a Bayesian framework, using the MAP [6] objective. Instead of maximizing the likelihood of the target tags, given source tags and parameters, we maximize the posterior probability of the parameters given the source and target tags:

$$P(\theta|\mathbf{x}, \mathbf{y}) \propto P(\mathbf{y}|\mathbf{x}; \theta)P(\theta|\mathbf{x}) = P(\mathbf{y}|\mathbf{x}; \theta)P(\theta). \quad (4)$$

By assuming, for each target tag t a normal distribution for \mathbf{w}_t centered around \mathbf{w}_t^{KB} with a precision matrix $\mathbf{\Lambda} = \lambda^2 \mathbf{I}$ (λ is independent of t), we can write the logarithm of the prior distribution as:

$$\log(P(\theta)) = \frac{\lambda^2}{2} \|\mathbf{W} - \mathbf{W}^{KB}\|^2 + \text{cte}. \quad (5)$$

We also consider a centered Gaussian prior on the bias (corresponding to a l^2 regularization). We then define:

$$\mathcal{L}_{\text{prior}} = \frac{\lambda^2}{2} \|\mathbf{W} - \mathbf{W}^{KB}\|^2 + \nu \|\mathbf{b}\|^2. \quad (6)$$

Using (3), (4) and (6), the final MAP objective becomes $\mathcal{L}_{\text{MAP}} = \mathcal{L} + \mathcal{L}_{\text{prior}}$, where the first term is the loss of Eqn (3), and the second can be seen as a regularization term on the weight matrix \mathbf{W} that penalizes its straying away from the priors. \mathcal{L} depends on the number of training samples, while $\mathcal{L}_{\text{prior}}$ does not. Therefore, \mathcal{L} becomes the predominant term in the loss as the size of the training data grows, leading to an objective function very close to the one of the logistic regression of Section 4.1. Conversely, when little data is available, we can expect the performances to be close or better than those of the KB system.

When a large parallel corpus is available, we can choose λ with grid search on a validation set. This is computationally expensive, and does not adapt well when the parallel corpus is small. For the sake of adaptability, we hereby propose a principled way inspired by [15] to choose λ , that does not require a lot of data while achieving top results. The rationale builds on the limited effective range of the logistic regression parameters. A shift of 5 of w_{ts} in the logit scale can move the probability associated with the target tag t from 0.5 to 0.99 or from 0.01 to 0.5. Hence, we would tend to choose λ in such a way that bigger shifts in the predicted probability of the target tag, which is the result of the added shifts for each annotated source tag, are unlikely. If we note \bar{N}_S the average number of source tags per sample (which can be estimated with a few samples), this would mean restricting the coefficients w_{ts} from shifting by more than $\frac{5}{\bar{N}_S}$. For a normal variable $X \sim \mathcal{N}(\mu, \frac{1}{\lambda})$ we have $P(X \in [\mu - \frac{2}{\sqrt{\lambda}}, \mu + \frac{2}{\sqrt{\lambda}}]) \approx 95\%$, we therefore propose to choose precision λ such that:

$$\frac{2}{\sqrt{\lambda}} \approx \frac{5}{\bar{N}_S} \tag{7}$$

5. EXPERIMENTS

We report the performances of the proposed models on a recording-based tag translation task. This also serves as an indirect evaluation of the DBpedia mapping, which, in a work dedicated to taxonomy mapping, could have been assessed by experts. Due to its novelty, we do not benchmark our work against other genre-related research from MIR.

5.1 Dataset

The dataset used in the experiments was created from the dataset used in the 2018 AcousticBrainz Genre Task, part of the MediaEval benchmarking initiative [7]. The dataset in its original form was aimed at testing the automatic genre annotation from content-based features of musical items in a more challenging setup compared to past works. For each item, annotations from different sources were available, each source taxonomy was much more detailed with hundreds of genres-subgenres, and the overall task was modeled as a multi-label classification. The sources were already introduced in Section 3.1. We further describe how the provided dataset was created. In Discogs, the release annotation was propagated to tracks. In Lastfm and Tagtraum, each track was annotated with music genres and subgenres from the derived taxonomies [7]. We present an overview of the dataset in Table 1.

Dataset	Discogs	Lastfm	Tagtraum
Annotation type	Expert	User	User
Items	1,098,337	686,979	589,584
Number genres	315	327	296

Table 1: Description of the dataset [7].

Although the data was already split between development, validation and test [7], we brought several modifi-

cations to accommodate the translation task. We created a large dataset comprising the original development and validation data. In order to assess a notion of confidence on the computed metrics, we resorted to K-Fold cross validation [14]. For each possible target, we splitted the data in 4 folds using stratified sampling. First, we filtered out the items which were not annotated in the target tag system. Then, we used an altered version of the iterative stratification algorithm in [32] in order to ensure that the proportion of items for each target label was roughly the same across folds. Following [13], we added the constraint that items belonging to the same artist had to be assigned to the same fold. For that, we used MusicBrainz artist ids retrieved from the recording ids provided in the MediaEval data.

5.2 Evaluation setup


The presented models output a score for each target tag that relates to the confidence of this tag being used in the target annotation. We evaluated these outputs with a ranking metric called Area Under the receiver operating characteristic Curve (AUC), as commonly done in multilabel classification [22]. The (*macro*) averaging is over target tags and measures the ability of the system to rank higher a positive tag than a negative one. Specifically, shifting the values in a column by the same factor (or changing the values of \mathbf{b} in the logistic regression) does not change the AUC macro score, being in that sense, unaffected by item popularity.

We evaluated the logistic regression models on each fold and trained on the three others. We uniformly subsampled the training data to simulate low data availability and chose subsampling factors as powers of 2 between 2^{-13} and 1. Consequently, for the smallest subsampling factors, some source and target tags may not be present in the training data. We used scikit-learn [25] implementation for ML logistic regression, with L-BFGS as the solver. We wrote a Tensorflow [1] implementation of MAP logistic regression. The Adam optimizer [17] was used, with a learning rate of 0.5. We trained the model for 500 epochs with batches of size 100000 or with the full training set if there were less samples. λ was chosen using Eqn 7.

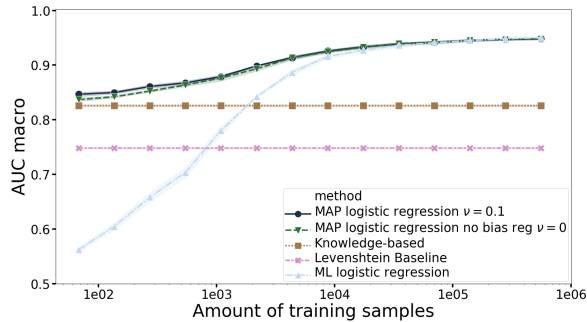
5.3 Results

Figure 2 illustrates how the ML translation eventually outperforms the KB model when enough data is available, while the latter performs much better when little data is available. The MAP translation successfully builds on the KB translation to yield the best results across the whole data availability spectrum. A simple baseline based on tag Levenshtein distance is also shown. Using only a source instead of two (e.g. only Lastfm) led to the same kind of behavior. While we currently proposed one method to obtain the KB translation table, we could also imagine it replaced by an expert-created one, if desired.

The fact that the MAP logistic regression performs consistently well on all the translation tasks is favorable evidence towards the choice of λ given in Eqn 7 as a good default, which we also confirmed using a grid search. Furthermore, we see that the MAP logistic regression lever-



(a) Translation from Lastfm and Tagtraum to Discogs



(b) Translation from Discogs and Tagtraum to Lastfm

Figure 2: AUC scores for the Levenshtein models per amount of training samples (log-scale). The width of the area around the lines marks the standard deviation computed on folds. KB yields a constant value as it is data-independent.

ages even low amounts of training data to improve over the KB model, and even more so when applying regularization on the bias. We further explain this effect by analyzing how the AUC scores compare on a per-tag basis.

Figure 2 shows that MAP logistic regression with bias regularization can achieve better AUC scores than the KB translation, even on target tags absent during training. We argue that this is due to the regularization term on the bias that enables to learn negative correlations between tags. When no bias regularization is used, the optimal set of parameters for a tag t missing from the training data is: $(\mathbf{w}_t^*, b_t^*) = (\mathbf{w}_t^{KB}, -\infty)$. Indeed, we see in Figure 3 that the MAP results are very close to those of the KB model. This is not true anymore with bias regularization. The gradient of the cost function w.r.t. \mathbf{w}_t can be written as:

$$\frac{\partial \mathcal{L}_{\text{MAP}}}{\partial \mathbf{w}_t} = [\hat{s}\sigma(w_{ts} + b_t) + \lambda^2(w_{ts} - w_{ts}^{KB})]_{1 \leq s \leq |S|} \quad (8)$$

where \hat{s} is the number of times (possibly 0) the source tag s appears in the training set. We therefore see that as w_{ts} gets closer to w_{ts}^{KB} , the term $\hat{s}\sigma(w_{ts} + b_t)$ will start to outweigh the second. Applying a gradient step will tend to decrease w_{ts} , away from w_{ts}^{KB} , and even more so when \hat{s} is large (popular tags) and b_t is close to 0 (controlled by the regularization term), hence the negative correlation.

Finally, it is worthwhile to mention that the AUC metric relies on occurrences and is thus arguably biased towards statistical methods. We end this section by taking a qualitative look at how statistics modified the similarities between source and target tags, in particular for those with




Figure 3: AUC scores for tags that were not in the training set, subsampled by a factor of 2^{-12} with Lastfm as target.

very different KB and ML AUC scores. These differences fall under four explanations:

- **Annotation noise:** Statistical models learn a very high similarity between the Discogs tag *italo-disco* and the Lastfm tag *classicalbritishheavymetal*. Both indeed often co-occur in the data, but are ontologically unrelated.
- **The target tag does not have a suitable equivalent in the source taxonomies.** Some Latin and Caribbean music genres like *cumbia*, *fado*, *rocksteady* or *forró* are present in Discogs but are not in Lastfm or Tagtraum. Thus, the mapping to DBpedia, described in Section 3.3, fails.
- **The considered tag is highly ambiguous.** Take the example of the tag *classical*. Besides the identical counterparts, knowledge-based translation tables also indicates relatedness to some subgenres of *jazz*. However, the specific translation task on which we evaluate appears to be more biased towards an understanding of *classical* that relates to subgenres of *metal* and *electronic* music (*symphonicmetal*, *germanmetal*, *postmodernelectronicpop*).
- **The existing genre representations are incomplete or noisy.** For instance, *baroque* has a counterpart in each taxonomy, but no direct link with *classical* in DBpedia. Statistical models find high correlation in the data between those two tags and so achieve better AUC scores.

6. CONCLUSION

In this work, we investigated the translation of tags from various source tag systems to a common target tag system. We show that the availability of large amounts of data advantages statistical methods over the knowledge-based one in terms of multilabel classification metrics. Moreover, the proposed hybrid method consistently outperforms both other methods on the whole range of data availability.

Although we did not address multi-language tag systems, both the knowledge-based approach that uses a mapping through the multilingual DBpedia, and the data-informed approach that only takes advantages of parallel annotations and is then insensitive to language, should be able to handle it. As future work, we aim to gather multilingual music genre datasets in order to confirm this claim. We also aim to exploit more thoroughly the genre graph we created by adding more knowledge sources and generating genre representations as node embeddings. We also consider modelling the tag annotation noise, such as missing or spurious tags, or tag bombing, in order to filter it out.

7. REFERENCES

- [1] Martín Abadi and et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Manel Achichi, Pasquale Lisena, Konstantin Todorov, Raphaël Troncy, and Jean Delahousse. Doremus: A graph of linked musical works. In *International Semantic Web Conference*, pages 3–19, 2018.
- [3] Derek Anderson. Word ninja, 2019. Software available at <https://github.com/keredson/wordninja>.
- [4] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In *The semantic web*, pages 722–735. Springer, 2007.
- [5] Jrgen Bang-Jensen and Gregory Z. Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer Publishing Company, Incorporated, 2nd edition, 2008.
- [6] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006. pages 30–31.
- [7] Dmitry Bogdanov, Alastair Porter, Julián Urbano, and Hendrik Schreiber. The mediaeval 2017 acousticbrainz genre task: content-based music genre recognition from multiple sources. In *MediaEval 2017 AcousticBrainz*, 2017.
- [8] David Brackett. *Categorizing sound: genre and twentieth-century popular music*. Univ of California Press, 2016.
- [9] Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. In *International Conference on Learning Representations*, 2018.
- [10] Emanuele Coviello, Riccardo Miotto, and Gert R.G. Lanckriet. Combining content-based auto-taggers with decision-fusion. In *Conference of the International Society on Music Information Retrieval*, pages 705–710, 2011.
- [11] Alastair JD Craft, Geraint A Wiggins, and Tim Crawford. How many beans make five? the consensus problem in music-genre classification and a new evaluation method for single-genre categorisation systems. In *Conference of the International Society on Music Information Retrieval*, pages 73–76, 2007.
- [12] Rene De La Briandais. File searching using variable length keys. In *Western Joint Computer Conference, IRE-AIEE-ACM ’59 (Western)*, pages 295–298, New York, NY, USA, 1959. ACM.
- [13] Arthur Flexer. A closer look on artist filters for musical genre classification. In *Conference of the International Society on Music Information Retrieval*, pages 341–344, 2007.
- [14] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*. Springer series in statistics New York, 2001. pages 241–249.
- [15] Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, and Yu-Sung Su. A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2(4):1360–1383, 2008.
- [16] Romain Hennequin, Jimena Royo-letelier, and Manuel Moussallam. Audio based disambiguation of music genre tags. In *Conference of the International Society of Music Information Retrieval*, pages 645–652, 2018.
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Jin Ha Lee and J Stephen Downie. K-pop genres: A cross-cultural exploration. In *Conference of the International Society on Music Information Retrieval*, pages 529–534, 2013.
- [19] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6:167–195, 2015.
- [20] Pasquale Lisena, Konstantin Todorov, Cécile Cecconi, Françoise Leresche, Isabelle Canno, Frédéric Puyrenier, Martine Voisin, Thierry Le Meur, and Raphaël Troncy. Controlled vocabularies for music metadata. In *Conference of the International Society on Music Information Retrieval*, pages 424–430, 2018.
- [21] Michael Mandel, Douglas Eck, and Yoshua Bengio. Learning tags that vary within a song. In *International Society for Music Information Retrieval Conference, ISMIR 2010*, pages 399–404, 08 2010.
- [22] Sergio Oramas, Oriol Nieto, Francesco Barbieri, and Xavier Serra. Multi-label music genre classification from audio, text and images using deep features. In *Conference of the International Society on Music Information Retrieval*, pages 23–30, 2017.
- [23] Lorena Otero-Cerdeira, Francisco J. Rodríguez-Martínez, and Alma Gómez-Rodríguez. Ontology matching: A literature review. *Expert Systems with Applications*, 42(2):949–971, 2015.
- [24] Panagiotis Papadimitriou, Panayiotis Tsaparas, Ariel Fuxman, and Lise Getoor. Taci: Taxonomy-aware catalog integration. *IEEE Transactions on Knowledge and Data Engineering*, 25(7):1643–1655, July 2013.
- [25] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning

- in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [26] Simone Paolo Ponzetto and Roberto Navigli. Large-scale taxonomy mapping for restructuring and integrating wikipedia. In *International Joint Conference on Artificial Intelligence*, IJCAI’09, pages 2083–2088, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.
- [27] Natalia Prytkova, Gerhard Weikum, and Marc Spaniol. Aligning multi-cultural knowledge taxonomies by combinatorial optimization. In *International Conference on World Wide Web*, WWW ’15 Companion, pages 93–94, New York, NY, USA, 2015. ACM.
- [28] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Machine learning*, 85(3):333, 2009.
- [29] Steven S. Aanen, Damir Vandic, and Flavius Frasin-car. Automated product taxonomy mapping in an e-commerce environment. *Expert Systems with Applications*, 42(3):1298–1313, 2015.
- [30] Hendrik Schreiber. Improving genre annotations for the million song dataset. In *Conference of the International Society of Music Information Retrieval*, pages 241–247, 2015.
- [31] Hendrik Schreiber. Genre ontology learning: Comparing curated with crowd-sourced ontologies. In *Conference of the International Society on Music Information Retrieval*, pages 400–406, 2016.
- [32] Konstantinos Sechidis, Grigorios Tsoumakas, and Ioannis Vlahavas. On the stratification of multi-label data. In *European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 145–158, Berlin, Heidelberg, 2011. Springer-Verlag.
- [33] Mohamed Sordo, Oscar Celma, Matin Blech, and Enric Guaus. The Quest for Musical Genres: Do the Experts and the Wisdom of Crowds Agree? In *Conference of the International Society on Music Information Retrieval*, pages 255–260, 2008.
- [34] Robyn Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. In *AAAI Conference on Artificial Intelligence*, pages 4444–4451, 2017.
- [35] Dennis Spohr, Laura Hollink, and Philipp Cimiano. A machine learning approach to multilingual and cross-lingual ontology matching. In *International Semantic Web Conference*, pages 665–680, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [36] Bob L. Sturm. Classification accuracy is not enough. *Journal of Intelligent Information Systems*, 41(3):371–406, 2013.
- [37] Aaron Swartz. Musicbrainz: A semantic web service. *IEEE Intelligent Systems*, 17(1):76–77, 2002.
- [38] Tobias Swoboda, Matthias Hemmje, Mihai Dascalu, and Stefan Trausan-Matu. Combining taxonomies using word2vec. In *ACM Symposium on Document Engineering*, DocEng’16, pages 131–134, New York, NY, USA, 2016. ACM.
- [39] Brian Tomasik, Joon Hee Kim, Margaret Ladlow, Malcolm Augat, Derek Tingle, Rich Wicentowski, and Douglas Turnbull. Using regression to combine data sources for semantic music discovery. In *Conference of the International Society on Music Information Retrieval*, pages 405–410, 2009.
- [40] Strother H. Walker and David B. Duncan. Estimation of the probability of an event as a function of several independent variables. *Biometrika*, 54:167–178, 1967.
- [41] Jun Wang, Xiaoou Chen, Yajie Hu, and Tao Feng. Predicting High-level Music Semantics using Social Tags via Ontology-based Reasoning. In *Conference of the International Society on Music Information Retrieval*, pages 405–410, 2010.
- [42] Eric W Weisstein. Statistical rank, 2019. From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/StatisticalRank.html>.
- [43] Tianxing Wu, Guilin Qi, Haofen Wang, Kang Xu, and Xuan Cui. Cross-lingual taxonomy alignment with bilingual biterm topic model. In *AAAI Conference on Artificial Intelligence*, pages 287–293, 2016.
- [44] G. K. Zipf. *Human behavior and the principle of least effort*. Cambridge, MA, Addison-Wesley, 1949.