

Disciples of the Heinous Path: Exploring Label Structure in Heavy Metal Genres

Maarten van Hees¹[0000-0002-4887-9371], Peter van der Putten¹[0000-0002-6507-6896], and Maarten H. Lamers¹[0000-0003-2672-5475]

Media Technology, Leiden University, P.O. Box 9512, Leiden, The Netherlands
vanheesmaarten@hotmail.com, p.w.h.van.der.putten@liacs.leidenuniv.nl,
m.h.lamers@liacs.leidenuniv.nl

Abstract. Heavy Metal is a popular sub culture, and in itself is highly tribalized, which makes it an interesting domain to research how cultures and sub cultures relate and evolve. To study this, we scrape the Encyclopaedia Metallum heavy metal music archive website to generate a large scale networked data set. Bands are linked through shared musicians, and each band can be labelled with multiple user contributed genres. By applying Word2Vec on genre co-occurrences, and hierarchical network clustering on the band collaboration graph, we gain insight into how music genres relate to each other. While the Word2Vec results show some interesting patterns with regards to the observed clusters, the hierarchical clustering proves to be more inconclusive, partially caused by factors beyond genre that generate the network. From a machine learning point of view, this case is an instance of the more general problem of understanding label structure in networked data.

Keywords: label structure · hierarchical graph clustering · heavy metal

1 Introduction

In case of labeled networked data, not just the relationship between the data and the labels is of interest, but also the structure within the labels itself, hierarchical or otherwise. For example products could be mapped to a product hierarchy, papers and authors to interrelated fields of study and patents to overlapping areas of application.

This phenomenon also occurs in the specific domain of study of this paper, which is Heavy Metal. We would like to understand how the various genres of Heavy Metal relate to each other, as a proxy for how metal cultures and sub-cultures relate and evolve. From a humanities point of view, metal culture is an interesting domain to study cultural evolution in general, given that metal is highly tribalized. To study this, a large dataset was acquired by scraping the metal band archive website Encyclopaedia Metallum [3], which contains bands, albums, artists that are interlinked if there are shared band members, thus containing a collaboration network of metal bands.

We use a variety of methods to understand genre label structure. First as bands can be labelled with multiple user contributed genres, we analyze label co-occurrence through Word2Vec and t-SNE [37]. Second, we analyze the artist-band collaboration network in an attempt to understand genre structure better. By clustering and labeling these graphs with label types such as genre and location we obtain a better understanding of the underlying processes and factors that drive the generation of these collaborations. We then carry out a hierarchical clustering on these graphs to see whether we can uncover relationships across genres, which is complicated by the fact that more factors than genre alone generate the network.

The existing work on understanding label and social network structure is sparse, let alone on heavy metal music and its associated culture. But the methods used in this exploratory case study are not very specific to our domain, and can provide inspiration to anyone interested in understanding label structure in networked data in their own domain of study. Hence this paper is written for a data mining, as well as a digital humanities audience, and should be seen as an exploratory, proof of concept study.

The paper is organized as follows. Firstly, an outline of preliminary and related work is given to provide context (Section 2), followed by a description of the problem statement, the used dataset and methodology in Section 3, which also describes the experimental setup. Section 4 deals with the results of the experiment, Section 5 provides a discussion, followed by Section 6, which provides conclusions and future work.

2 Background

This section outlines available prior literature focusing in the field of Heavy Metal. Next, an overview of existing literature on collaborative networks arising from music and their surrounding communities is given.

2.1 Heavy Metal subculture

Heavy Metal subculture has been the topic of choice for a number of different studies in a wide range of fields. Most studies focus on aspects of Heavy Metal subculture that set it apart from traditional societal norms (e.g. [6]). One often-heard argument against Heavy Metal is that it presumably would lead to depression, nihilism and even suicide, the latter having been studied extensively (e.g. [7][8][4]). When it comes to real world examples, there are certainly a number of excesses present in Heavy Metal that prove the point quite effectively, such as the absolute into said depression, nihilism and suicide that characterized the early Norwegian black metal scene [32]. Even so, while this has been the subject of a number of documentaries [33][34][35], it has evolved into a fairly standard fair subscene. Other studies focus on behavioral aspects [9], importance of geography within collaborative networks [10], sense of community [11], the usage of masks and shadows [12][13], merchandising and creative practices [14][15][16],

philosophy [17], religion in general [18][20] and Biblical exegesis in particular [19]. The relatively tribal nature of the Heavy Metal scene itself means it lends well to studying subcultures contained within subcultures, or genres contained within or related to other genres. Of particular interest is the structure of these genres. We aim to gain more insight in the evolution of the relations between genres, if any. The field of social network analysis in conjunction with Heavy Metal has been relatively unexplored.

2.2 Social networks in music

There have been a few studies using the type of collaborative networks that arise from music genres. A study done by Pablo M Gleiser en Leon Danon [1] focused on the community structure in Jazz. They used a dataset of 128 jazz bands from 1912 to 1940 and found quantitative correlations to an email based social network with regard to typical community degree-distribution. The resulting network also shows correlations between racial segregation, recording location and the community structure.

Teemu Makkonen investigated the Finnish heavy metal scene, exploring the collaboration between music genre heavyweights, evolution of their collaboration networks and whether or not geographical location is important in the development of this collaboration network [5]. While earlier literature [1] suggests that the presence of a local music scene is vital for the development of a band's success, Makkonen's findings suggest otherwise, namely that successful bands are more likely to communicate via global pipelines, rendering the geographical aspect of a collaboration network less important than expected. There are a number of other studies focusing more generally on collaborative networks in music, such as a study on important professions and comparison on different collaborative networks [2], and a study focusing on finding influential nodes in a music based social network, which highlights the importance of meritocracy over topological position of leading actors in this social network.

3 Problem statement and methodology

This section will first provide a description of the problem statement and general approach, followed by the used data source, method of data collection, interpretation and resulting data.

3.1 Problem statement

Our interest lies in the structure and development of Heavy Metal subcultures, and what, if any, relation exists between genres. Secondly, we aim to explore methods for understanding label structure in networked data that are more also applicable outside of our particular domain of interest. Because of the scarce preliminary work, we decided to take an an exploratory approach [21]. One matter that will be explored explicitly is whether or not a social network (i.c. one

of Heavy Metal collaborations) is organized by its labels (i.c. Heavy Metal music genre), and whether we can discover structure across such labels, hierarchical or other.

To gain more insight into labeling structure, we follow two approaches. First, a Word2Vec analysis will be performed using a t-SNE plot in which the similarity across genre labels can be inspected. Then a hierarchical clustering method will be applied to the band collaboration graph to study whether genre label relationships can be derived from the graph.

3.2 Encyclopedia Metallum

The dataset used in this study was obtained by scraping Encyclopedia Metallum (EM) [3], a website devoted to cataloguing Heavy Metal bands, albums and individual artists. EM’s content is entirely generated by its community, which introduces a number of challenges. EM contains, at time of writing, 121,856 bands with associated artists, album releases, album reviews, song lyrics and more. EM has been around since 2002, has a very active user base and has been used sparingly in previous studies, such as a study done by Ingeborg H Aarsand on the usefulness of representing natural utopias in music as a response to environmental crises [22], the influence of Heavy Metal in times of social instability [23] and the aforementioned study on collaborative networks of the Finnish metal scene [5].

3.3 Data collection and pre-processing

The data contained in Encyclopedia Metallum is structured alphabetically. For technical reasons such as a lack of bandwidth and processing power and time constraints, it was decided to include random 50% of the bands within each lettered list, going from A to Z and ending with a special section containing bands with names written in non-Latin characters. Using a custom-built website scraper the data was obtained over the course of several months during spring 2018, interpreted and saved to a relational database.

In order to counter the fuzzy nature of much of the meta-data contained within the dataset, as illustrated in table 1, the data was run through a filtering method based on a stop-word list. Because the bands’ genre labeling is important to this study, the filtering will focus on this particular meta-data attribute.

The filtering method uses a list of stop-words. These stop-words include terms that are irrelevant to the band’s genre. This includes things like ‘(early)’, ‘(later)’, slashes, backslashes, comma’s, and the word ‘metal’. This allows one to create a finite list of genres, decreasing the dimensionality of the genre data. In the case of the Greek band Rotting Christ, after applying the filtering method, this would yield the terms Grindcore, Black, Gothic and Melodic Black.

Bands typically have more than one genre. For example, a band can hypothetically be labeled with the Black Metal and Death Metal genres. Running this band through the filter method would yield two separate genres, Black and Death. However, this leads to some ambiguity, as there are fusion genres present

Table 1. Genre labels from Encyclopaedia Metallum and subsequent filtering.

Band	Genre label	After filtering
Rotting Christ	Grindcore, Black Metal (early), Gothic Metal (mid), Melodic Black Metal (later)	Grindcore, Black, Gothic, Melodic Black
In Flames	Melodic Death Metal (early), Melodic Groove Metal/Alternative Rock/ Metalcore (later)	Melodic Death, Melodic Groove, Alternative Rock, Metalcore
Meshuggah	Technical Groove/Trash Metal (early), Djent (later)	Technical Groove, Trash, Djent
Darkthrone	Death Metal (early), Black Metal (mid), Black/Heavy/Speed Metal (later)	Death, Black, Black, Heavy, Speed

within Heavy Metal, such as Blackened Death Metal. This genre seems pretty similar to Black/Death. However, these fusion genres are treated as a unique genre, as they typically denote a band that does play a fusion of two genres, while Black/Death denotes an evolution from one genre to the next over the course of the band's existence.

3.4 Genre family tree

One way music genres are typically structured in terms of relation to one another is by means of a so-called genre family tree. A genre family tree shows the development of genres into new genres. We decided to obtain such a family tree as to have an a priori label hierarchy which can be compared to whatever relations between genres we may find. While fierce discussion about the structure of this family tree remains [36], the hierarchy with regards to main genres, such as the examples listed before, of metal are widely accepted. There are a number of such genre family trees available online with varying levels of detail. We picked the genre family tree by Bound by Metal [24] as it came in a readily usable format and has a suitable level of detail.

3.5 Genre embedding with Word2Vec

Using Word2Vec in conjunction with t-SNE we create a co-occurrence map of genres using the obtained dataset, stripped of bands that have but one or zero genres. These bands' genres were used as training documents for Word2Vec's neural network, mapping similar genres by their co-occurrence in the documents. The co-occurrence map may shed insight into what genre typically occurs with what other genre, thus opening new avenues of speculation. Word2Vec's neural network approach works in our advantage, as we simply aim to map genre co-occurrence, regardless of semantic meaning, making it more suited for our needs than LSA [40]. The Word2Vec analysis was carried out with the Python3 package Gensim [38], the t-SNE plot was generated with the scikit-learn [39] package.

3.6 Exploring the band collaboration network on genre and locality

For a number of countries, we subsequently create band collaboration networks $G = (V, E)$ with nodes V corresponding to bands, and edges E exist between two nodes if band members have collaborated with these two bands. To visually explore these networks, and understand the underlying processes that generate these networks, we structured and visualized these networks using the ForceAtlas2 algorithm [41], and labeled nodes with genre as well as location. Experiments were carried out with the Gephi tool [27]. See section 4.2 for the results.

3.7 Hierarchical graph clustering of the band collaboration network

Whilst there may be more factors than just genre that drive band collaboration, we wanted to see whether we could evolve genre relationships from the network. To achieve this we carry out a bottom up, hierarchical network clustering on the largest connected component for various countries, and check the genre distribution at different nodes of the hierarchy. The largest connected component was obtained using a standard BFS algorithm [28] and was needed because of the distance measure we used, which is based on shared artists. If there are no shared artists between bands, there is no connection, hence the choice for the largest connected component.

Our hierarchical clustering method uses a ‘clustroid’ distance measure, using edge-weight based on the number of shared artists between bands. This ‘clustroid’ is the most central node in a given cluster, which would intuitively be at the center of a group of networked nodes. Initial distances were calculated using the Floyd-Warshall algorithm [26], which we parallelized [25] to counter the $O(V^3)$ time complexity. Using the largest connected component for the hierarchical clustering has the benefit of making the termination criteria very simple: $C=1$, where C is the number of remaining clusters.

4 Results

This section discusses the results of genre embedding analysis with Word2Vec, the role of genre and locality in the band collaboration graph and the genre hierarchy generated from the graph using hierarchical graph clustering. The data collected c.f. section 3.3 amounted to 63,306 bands and 288,872 artists that were used in the analyses.

4.1 Genre co-occurrence

The Word2Vec method takes in a list of documents from which the co-occurrence vector space is constructed. This list of documents is a band-wise list of genre labels, processed c.f. section 3.3, and stripped of bands that have but one or zero genre labels. This yielded a list of 28,007 band-wise genre co-occurrence documents.

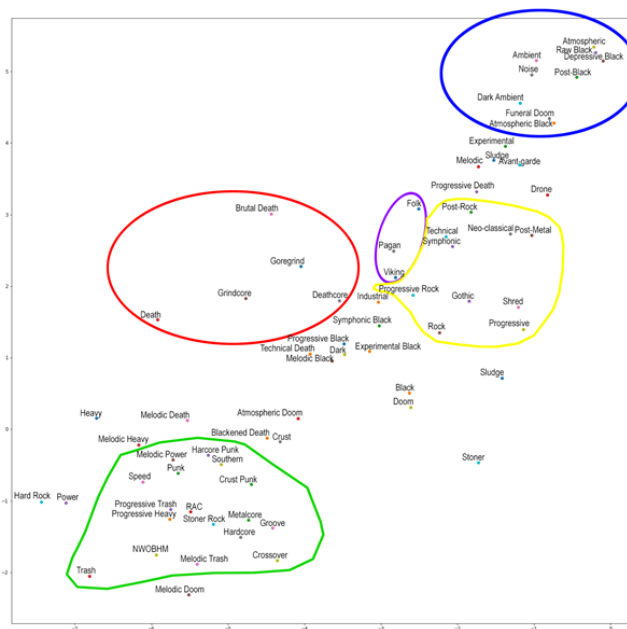


Fig. 1. t-SNE plot of Word2Vec results

The Word2Vec model was implemented using the Python gensim package, ignoring all documents with less than 75 occurrences. Setting the threshold of minimum occurrences for genres as such, eliminated most of the fringe genres and provided a good balance between resolution and readability in the t-SNE plot (fig. 1). The clusters, displayed as colored circles, have been added manually while inspecting the plot.

Several interesting things are noted in this plot from a Heavy Metal domain perspective. Broadly speaking, the plot clusters heavily on style of genre. The blue cluster contains all the ambient genres, which is a slow, melodic style of metal music. The only two outliers in this genre are Funeral Doom and Depressive Black, but these could arguably be part of the ambient genre group as they focus heavily on atmospherics and have a similar style in lyrics. The second cluster worth noting is the red cluster. These are, roughly speaking, the extreme styles of Death Metal. Goregrind, Grindcore and Brutal Death Metal are often used interchangeably as they are fairly similar genres. Deathcore is very much a spin-off and, arguably, a simplification of the Brutal Death Metal genre. The green cluster features all of the genres emerging from the fusion of Punk and Trash, of which the genre Crossover is the most pure mixture of Punk and Trash. One genre that, at first glance, shouldn't be part of this cluster is the RAC genre, which stands for Red-Anarchist and is a sub genre of Black Metal, which does not share much similarity with the typical Punk and Trash genres. However, a case can be made for its inclusion, as RAC is very leftist politically oriented,




Fig. 2. All Colombian bands, colored by genre

very much like the wider Punk movement. The final two clusters are fairly notable in much the same way. The purple cluster are the pagan styles of metal, which focus heavily on folklore, mythology and traditional folk music elements. The yellow cluster are the progressive styles of metal, generally speaking. This is probably the most nebulous cluster, although it is pretty interesting that Shred, Neo-classical and Technical are clustered close together along with Progressive rock and Symphonic. This makes sense: these are all off-shoots from the initial Heavy Metal style of music, incorporating more instruments (Symphonic), experimenting with song structure (Progressive) and extreme technicality (Shred, Technical and Neo-classical).

As stated before, this analysis simply looks at genre co-occurrence, band by band. The clusters found in the plot do point to an almost one dimensional embedding in the data. Bands tends to not mix their genres around much and when they do, it tends to be variations on a certain theme, such as the political themes of the punk genres, the explicit violence and gore of the extreme Death Metal genres and the heavy focus on ambience and depressive lyrics of the ambient genres. This does seem to point towards an organization in genre choice.

4.2 Exploring the band collaboration network

We constructed band collaboration networks for three subsets of the full dataset: all bands from Canada (891 bands), Belgium (363) and Colombia (462). These were selected for their limited network size and the difference in location which may or may not have bearing on the properties of the social network.

To gain some insight into these networks and what associated factors might contribute to their structure, such as genre, locality or others factors, a few experiments using Gephi [27] were conducted. In the force directed collaboration graph of Colombian bands, obtained with Gephi’s inbuilt ForceAtlas2 layout algorithm, the Black Metal cluster of bands in the lower part of the image stands out (fig. 2). There is a strong Black Metal tradition in Colombia, dating back to the early 1980’s which, arguably, was influential in the development of the revered Norwegian scene [29]. This scene was based at the city of Medellin, which can be seen (at least partly) when the graph is colored with the band location (fig. 3).



Fig. 3. Colombian bands



Fig. 4. Canadian bands

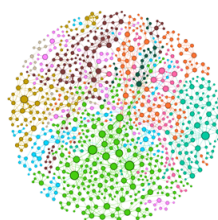


Fig. 5. Canadian bands, Louvain method applied.

The bands that are part of the isolated Black Metal sub-scene are and have been part of the Medellín "Ultra-Metal" movement. Interestingly enough, there is a clear split between the two most represented localities, Medellín and Bogotá. This could be because of the strong scene based around Medellín, and subsequent scenes evolving around that city. Also, since the Medellín scene is relatively unknown outside of Colombia, it could be that Heavy Metal has not taken root outside of the major cities or that local communication lines are more important than regional or provincial communication lines. This would oppose the situation in Finland, as Makonnen observed [5].

The Canadian network has a slightly different look (fig. 4). It does not have the insular community that Colombia has, but a number of interconnected communities that center around high-degree nodes and the major cities. These communities become visible when applying the Louvain method (resolution 3.0) and coloring according to modularity class (fig. 5).

Interestingly enough, the Montreal and Quebec bands are placed into the same community as they are fairly close together and share a language, while bands from Vancouver and Toronto are in their own communities, being at pretty much opposite ends of the country. This again may oppose Makonnen's findings, although his dataset did not take into account regional language differences. It might be interesting to investigate the effect of language on the development of similar social networks.

The Colombian Medellín scene has clearly influenced the structure of the network. This Medellín scene was also one that did not really extend beyond Colombian borders; it was very much a local scene. This could prove to be a method to find the more extreme and ostracized parts of the metal community, such as the National-Socialist and Red Anarchist scenes (the latter combin-

Table 2. Colombia

Genre	%
Black	24.01
Death	19.88
Thrash	18.72
Heavy	8.41
Brutal Death	4.65
Melodic Death	4.02
Power	2.45
Progressive	2.13
Gothic	2.04
Speed	1.77
Sum %	88.08

Table 3. Belgium

Genre	%
Black	16.87
Death	16.37
Thrash	12.87
Heavy	9.39
Doom	4.44
Melodic Death	4.19
Metalcore	3.2
Brutal Death	2.3
Power	2.21
Groove	1.82
Sum %	73.66

Table 4. Canada

Genre	%
Death	17.16
Black	16.71
Thrash	13.35
Heavy	4.4
Melodic Death	4.3
Power	3.63
Progressive	3.57
Grindcore	3.15
Brutal Death	2.75
Technical Death	2.56
Sum %	71.58

ing far-left and environmental philosophies with black metal). The former has received some interest from the scientific community [30][31], yet the latter remains an understudied phenomenon.

The Belgian dataset is very similar in structure to the Canadian set, but will be explored further on.

4.3 Hierarchical clustering results

The explorative analysis clearly demonstrates that multiple factors drive band collaborations: genre, but also locality. Nonetheless, we are interested in whether we can understand genre structure through hierarchical clustering of the band collaboration network. As stated before, the analysis was run on the largest connected component of each country’s band collaboration graph. The hierarchical clustering method yielded a number of histograms which were used to view the convergence of genre within the network.

Tables 2 to 4 show the distribution of the top 10 largest genres within the experimental datasets. All three countries have a fairly similar genre distribution. The Black, Death, Thrash and Heavy genres are most frequent. Something that stands out is the Black Metal frequency for Colombia, which is a fair bit higher than for Belgium and Canada. This makes sense, as there is a strong Black Metal scene in Colombia, as discussed before.

At this stage it is unclear whether or not the actual genre hierarchy is reflected in the structure of the social network. To gain more insights into this, a colored dendrogram was created through the hierarchical clustering procedure. The links of the dendrogram were colored according to the highest distribution of genre. Each cluster has its own genre distribution and an associated histogram, like the ones shown in tables 2 to 4. For each cluster, the highest frequency genre was picked and assigned a color.

In the colored dendrogram for Colombian bands (fig. 6), one thing that stands out is the high density branch in the center of the dendrogram. These are the bands from the aforementioned Medellin scene, contributing a high fraction to the overall Black Metal frequency. However, because of the distribution of genres,

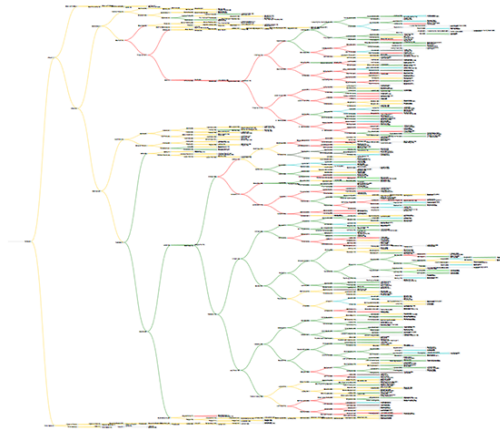


Fig. 6. Hierarchically clustered Colombian bands, colored by highest genre distribution

i.e. the vast over-representation of the genres Black, Death, Trash and Heavy, the dendrogram loses its overview. At some resolution, the overall genre distribution takes shape and from that point onward, moving deeper into the tree, the clusters will have a distribution similar to that of the total tree. It may be more interesting to look at the complete distribution at certain points and balance the data so that it contains an equal representation of genres, even though it may prove challenging to retain the social network structure when balancing the data.

At the local resolution there is some more visible variation in genres, although there is not much to suggest the structure follows the genre hierarchy. This also goes for the Belgian dendrogram. The comparison with the genre hierarchy was done using visual inspection, from which it is unfortunately unclear if the genre hierarchy is present in the hierarchically clustered data. One reason is in the nature of both hierarchies. The hierarchies as seen in figs. 6 and 7 are hierarchies of distribution. This hierarchy is weighted and contains quantitative information with regards to genre distributions. The genre hierarchy itself is not a distribution, but a straightforward hierarchy of tags.

It stands to reason to do the inspection using a discrete computational method. However, it is unclear whether such a method exists or can be modified to work for this particular case. Before a computational method can be applied for comparing these two types of hierarchies, either some filtering needs to occur or a different method of clustering must be attempted, which brings its own challenges.

The dendrogram of Belgian bands in fig. 7 is much less clustered than its Colombian counterpart. Belgium lacks a strong metal scene, unlike Colombia. This means that the structure of the Belgian social network is a lot more connected as there is no isolated scene. This can also be seen when visualizing the Belgian network with Gephi (fig. 8).



Fig. 7. Hierarchically clustered Belgian bands, colored by highest genre distribution

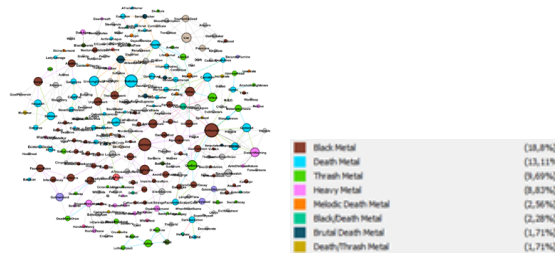


Fig. 8. All Belgian bands, colored by genre.

This is where the hierarchical clustering method, and in particular the coloring, runs into another issue. With networks such as this one, the hierarchical clustering method and coloring does not quite work. There are a fair number of Trash Metal bands in Belgium as shown in fig. 8, but these do not show in the dendrogram. This may happen because these are clustered later along with Black and Death Metal bands, meaning the Trash Metal frequency never gets higher than those of Black and Death metal. Also, on the local resolution, there is nothing that suggests a genre hierarchy has bearing on the social network structure.

5 Discussion

This section will draw some conclusions and reflect on them. The final subsection will provide some directions regarding future work.

5.1 Methodology and data

The exploratory approach taken in this study allowed us to effectively explore the dataset and change directions on the fly. The scope-creep, which is always a problem when the goalposts keep shifting, was kept to a minimum.

The main issue with the dataset was fuzzy meta-data, which we attempted to resolve with the filtering method described earlier. This did remove the timed

element of some of the bands' genres, as the method removed the denotations marking a genre out as something earlier or later in the bands existence. This wasn't the focus of this study, but it could be used in future work.

5.2 Results

A major concern regarding our results relates to the dendograms. Because of the genre distributions being the way they are, the distributions within a cluster quickly take on the final distribution, thus coloring most of the dendograms by the top four most prevailing genres. This made it difficult to reconcile the dendograms with the genre family tree, unfortunately. The Word2Vec analysis did provide some insight however, providing clusters that are heavily style-oriented. This is fairly surprising, since the analysis was done with just a genre co-occurrence list. This points to the possibility that bands, when mixing genres, don't typically travel far off the beaten path, so to speak.

5.3 Future work

One thing that must be addressed is the role of time. A genre hierarchy such as the one used for this study does exist in time; the popularity of genres rise and fall when time progresses and new genres are formed. Therefore, taking the evolution of the social network through time into account when comparing it to the genre hierarchy may yield interesting results. Several challenges are included in this, such as the hierarchical clustering and visualization of the evolving network.

Other potential future work focuses more on the domain of Heavy Metal. This study only used locality and genre of the bands contained within the dataset, ignoring other meta-data attributes. These unused attributes can be interesting for future studies. For example, one could look at lyrical content in combination with locality, to see if there is a connection to local mythology or legends, or any other local particularity. Similarly one could analyze the lyrics to further investigate the more extreme elements within the Heavy Metal subculture.

6 Conclusion

This paper discusses the results of a case study to explore label structure in networked data, in particular to understand how Heavy Metal genres relate to each other. To study this, we scraped information about around 300.000 artists from the Encyclopaedia Metallum web site. As this is a community driven resource without much of a prescribed structure for genres and other information, this is a rich, yet noisy source on metal culture.

First we analyzed label co-occurrences with Word2Vec and t-SNE, and then we performed hierarchical clustering on the band collaboration network to understand to what extent a genre hierarchy could be generated from the network. An exploratory analysis on the network also showed that factors beyond genre

play a role in generating the network, such as locality, which results in confounding factors in the derived genre hierarchy. The approach presented here is not domain specific and could be used for any type of labelled networked data. As such our case study offers an initial step towards a general purpose approach for understanding such data.

References

1. P.M. Gleiser, L. Danon, *Community Structure in Jazz*, *Advanced Complex Systems*, vol 6 issue 4, pp 565-573, 2003.
2. P. Budner, J. Grahl, *Collaboration Networks in the Music Industry*, CoRR, 2016
3. Encyclopedia Metallum, <https://www.metal-archives.com>
4. E. Lacourse, M. Claes, M. Villeneuve, *Heavy Metal Music and Adolescent Suicidal Risk*, *Journal of Youth and Adolescence*, vol 30 issue 3, pp 321-332, 2001
5. T. Makkonen, *North from Here: The Collaboration Networks of Finnish Metal Music Genre Superstars*, *Creative Industries Journal*, vol 10, pp 104-118, 2017
6. R. L. Gross, *Heavy Metal Music: A New Subculture in American Society*, *Journal of Popular Culture*, vol 24 issue 1, pp 119-130, 1990
7. S. Stack, J. Gundlach, J. L. Reeves, *The Heavy Metal Subculture and Suicide*, *Suicide and Life Threatening Behavior*, vol 24 issue 1, pp 15-23, 1994
8. K. R. Scheel, J. R. Westefeld, *Heavy Metal Music and Adolescent Suicidality, An Empirical Investigation*, *Adolescence*, vol 34 issue 134, pp 253-273, 1999
9. J. Arnett, *Heavy Metal Music and Reckless Behavior Among Adolescents*, *Journal of Youth and Adolescence*, vol 20 issue 6, pp 573-592, 1991
10. T. Makkonen, *Tales from the Thousand Lakes: Placing the Creative Network of Metal Music in Finland*, *Environment and Planning A: Economy and Space*, vol 46 issue 7, pp 1586-1600, 2014
11. V. Venkatesh, D. Perri, J. S. Podoshen, K. Urbaniak, *From Pride To Prejudice To Shame, Multiple Faces of the Black Metal Scene Within and Without Online Environments*, *Educational Psychological and Behavioral Considerations in Niche Online Communities*, chapter 12, 2014
12. B. Hickam, *From Shamans to Cyborgs, the Use of Masks and Shadow in Branding Heavy Metal*, *Modern Heavy Metal: Markets, Practices and Cultures conference*, 2015
13. C. Hoad, *He Can Be Whatever You Want Him to Be: Identity and Intimacy in the Masked Performance of Ghost*, *Popular Music*, vol 37 issue 2, pp 175-192, 2018
14. T. Helyar-Cardwell, *Battle Jackets, Craft and Folk Culture: Research through Creative Practice*, *Modern Heavy Metal: Markets, Practices and Cultures conference*, 2015,
15. A. Berg, T. Gulden, V. Kiort af Ornas, N. Pavel, V. Sjøvoll, *The Metal T-shirt: Transmedia Storytelling in Products*, *Modern Heavy Metal: Markets, Practices and Cultures conference*, 2015
16. R. Platz Cortsen, *By The Cake Of The Dark Lord! Metal Cultures in Three Nordic Comics*, *Modern Heavy Metal: Markets, Practices and Cultures conference*, 2015
17. N. Scott, T. OBoyle, *Doom Values: Deceleration Promoting a Philosophy of Progression Through Opposites*, *Modern Heavy Metal: Markets, Practices and Cultures conference*, 2015
18. T. Hjelm, *Controversial Metal and Controversial Religion, Theorising The Connections*, *Modern Heavy Metal: Markets, Practices and Cultures conference*, 2015

19. C. N. Davis, *When Marduk Read Romans: Biblical Text as a Hermeneutic in Heavy Metal, Heavy Metal as Biblical Exegesis*, Modern Heavy Metal: Markets, Practices and Cultures conference, 2015
20. I. Abraham, *Innovation and Standardization in Christian Metalcore: The Cultural Influences of Church and Market*, Modern Heavy Metal: Markets, Practices and Cultures conference, 2015
21. P. Doherty, *The Beginner's Guide to Winning the Nobel Prize*, Columbia University Press, 2008
22. I. Husbyn Aarsand, *Imagining Cascadia: Bioregionalism as Environmental Culture in the Pacific Northwest*, Master Thesis, University of Oslo, 2013
23. G. Horetua, *Bullets, Drugs, and Rock and Roll: Colombian Punk Rock, Heavy Metal Culture in a Time of Revolt and Terrorism*, Proquest Dissertations Publishing, 2013
24. Genre hierarchy, Bound By Metal, retrieved from <https://www.boundbymetal.com/en/common/metal-genres-graph>
25. A. Pradhan, G. Mahinhakumar, *Finding All-pairs Shortest Path for a Large-scale Transportation Network Using Parallel Floyd-Warshall and Parallel Dijkstra Algorithms*, Journal of Computing in Civil Engineering, vol 27 issue 3, pp 263-273, 2012
26. R. W. Floyd, *Algorithm 97: Shortest Path*, Communications of the ACM, vol 5 issue 6, p 345, 1962
27. M. Bastian, S. Heymann, J. Mathieu, *Gephi, An Open Source Software for Exploring and Manipulating Networks*, International AAAI Conference on Web and Social Media 8, pp 361-362, 2009
28. K Zuse, *Der Plankalkül*, Gesellschaft für Mathematik und Datenverarbeitung vol 63, 1972
29. C. Iseman, *Medellin Ultra Metal, Invisible Oranges*, 2015, Retrieved from <http://www.invisibleoranges.com/medellin-ultra-metal/>
30. M. Gardell, *Gods of the Blood: the Pagan Revival and White Separatism*, Duke University Press
31. B. Hedge Olsen, *I Am the Black Wizards: Multiplicity, Mysticism and Identity in Black Metal Music and Culture*, Master Thesis, Bowling Green University, 2008
32. M. Phillipov, *Extreme Music for Extreme People? Norwegian Black Metal and Transcendent Violence*, Popular Music History, vol 6 issue 1, pp 1501-163, 2011
33. *Satan Rides the Media*, directed by T. Gruden, 1998
34. *Until The Light Takes Us*, directed by A. Aites and A. Ewell, Variance Films, 2009
35. *Once Upon a Time in Norway*, directed by P. Aasdal and M Ledang, Grenzeløs Productions, 2007
36. D. Weinstein, *Heavy Metal: the Music and It's Culture*, Chapter 2, Da Capo Press, 2000
37. L. van der Maaten, G. Hinton, *Visualizing Data Using t-SNE*, Journal of Machine Learning Research, vol 9, pp 2579-2605, 2008
38. R. Rehurek, P. Sojka, *Software Framework for Topic Modelling with Large Corpora*, Proceedings of LREC 2010 Workshop New Challenges for NLP Frameworks, pp 46-50, 2010
39. Pedregosa et al, *Scikit-learn: Machine Learning in Python*, Journal of Machine Learning Research, vol 12, pp 2825-2830, 2011
40. S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, R. Harshman, *Indexing by Latent Semantic Analysis*, Journal of the American Society for Information Science, vol 41 issue 6, pp 391-407, 1990
41. M. Jacomy, T. Venturini, S. Heymann, M. Bastian, *ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software*, PLoS ONE, vol 9 issue 6, e98679, 2014