

# AUGMENT, DROP & SWAP: IMPROVING DIVERSITY IN LLM CAPTIONS FOR EFFICIENT MUSIC-TEXT REPRESENTATION LEARNING

**Ilaria Manco**

Queen Mary University of London

i.manco@qmul.ac.uk

**Justin Salamon**

Adobe Research

salamon@adobe.com

**Oriol Nieto**

Adobe Research

onieto@adobe.com

## ABSTRACT

Audio-text contrastive models have become a powerful approach in music representation learning. Despite their empirical success, however, little is known about the influence of key design choices on the quality of music-text representations learnt through this framework. In this work, we expose these design choices within the constraints of limited data and computation budgets, and establish a more solid understanding of their impact grounded in empirical observations along three axes: the choice of base encoders, the level of curation in training data, and the use of text augmentation. We find that data curation is the single most important factor for music-text contrastive training in resource-constrained scenarios. Motivated by this insight, we introduce two novel techniques, Augmented View Dropout and TextSwap, which increase the diversity and descriptiveness of text inputs seen in training. Through our experiments we demonstrate that these are effective at boosting performance across different pre-training regimes, model architectures, and downstream data distributions, without incurring higher computational costs or requiring additional training data.

## 1. INTRODUCTION

Music-text embedding models have become a cornerstone of music information retrieval (MIR), facilitating core tasks that underpin music organisation and search, such as music tagging and cross-modal retrieval [8, 11, 14, 26, 27]. At a high level, these are multimodal models that produce aligned audio-text representations by learning to project high-dimensional data from the audio and text modalities onto a lower-dimensional joint representation space whose structure encodes semantic similarity. The canonical learning framework to obtain such embeddings is dual-encoder multimodal contrastive learning, first popularised by CLIP [20] in the image domain, and soon after adopted in most areas of machine perception, including audio [9, 29] and music processing [8, 11, 14].

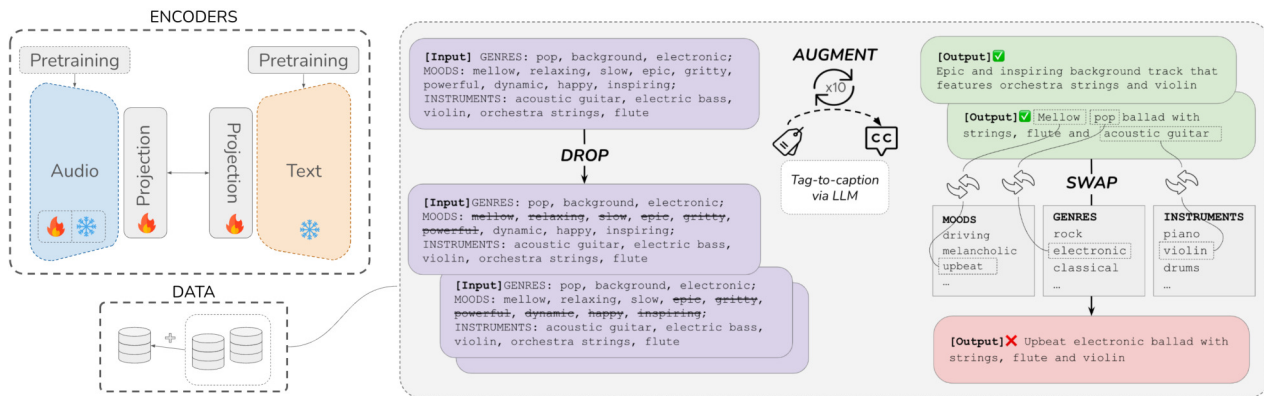
Driven by the empirical success of this framework, a recent line of research has attempted to analyse its inner workings from a theoretical perspective [18, 36] or elucidate which aspects are most responsible for its effectiveness in visual models [34, 35]. However, within the audio domain, our understanding of multimodal contrastive learning remains limited [8], with sparse effort into ablating design choices, or training data- and compute-efficient models. Among prior work that takes a step in this direction, the focus is mostly on comparing backbone models [8, 11, 29], but without considering other important factors such as model initialisation or training data. Additionally, audio-text learning poses specific challenges in the context of music, as the amount of data with aligned audio and text is typically orders of magnitude smaller than in other domains, where large-scale web-crawled data is commonplace. This makes transferring insights from other areas of representation learning particularly challenging.

In this paper we present a deep dive into music-text contrastive learning and its use in text-based music retrieval, adopting a practical perspective and thoroughly investigating the impact of major design choices. In particular, we study the problem of how to train this family of models under different resource-constrained scenarios (with respect to data and compute), and how to meaningfully evaluate them for real-world use. In brief, our contributions are as follows: (i) we systematically compare backbone encoders in parameter-efficient settings, and demonstrate that we can leverage this to enable multilingual support for the first time and without additional training data (Section 3); (ii) we study the trade-off between training dataset size and quality, showing that the impact of data curation outweighs that of scale (Section 4); (iii) building upon these findings, we propose a training recipe, *Augment, Drop & Swap* to construct more effective contrastive views (via *Augmented View Dropout*) and improve model robustness (via *TextSwap*) with no extra computational overhead (Section 5). Incorporating the proposed pipeline within variants of the music-text contrastive framework under different computational constraints, we show that this consistently improves over prior work, establishing a new state-of-the-art on three benchmark datasets. Finally, we conduct the first listening study to evaluate text-based music retrieval, further corroborating our automatic evaluations and underscoring the importance of accounting for distribution gaps when measuring performance.



© I. Manco, O. Nieto, and J. Salamon. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

**Attribution:** I. Manco, O. Nieto, and J. Salamon, "Augment, Drop & Swap: Improving Diversity in LLM Captions for Efficient Music-Text Representation Learning", in *Proc. of the 24th Int. Society for Music Information Retrieval Conf.*, Milan, Italy, 2023.



**Figure 1: Overview** of our approach. We study the role of encoders and data in music-text learning and propose a text augmentation pipeline, *Augment, Drop & Swap*, to increase data diversity and introduce hard negatives during training.

## 2. STUDYING THE DESIGN SPACE OF MUSIC-TEXT EMBEDDING MODELS

We explore two major factors in the design of music-text embedding models: architecture and data. While we acknowledge that there are others, such as training procedure, and alternative designs, we choose to restrict our focus exclusively to these two axes and to dual-encoder models, due to their predominance in the field. In the rest of the paper, we always refer to this family of models when discussing *music-text embeddings* or *music-text models*, and interchangeably use the terms *text* and *language*.

The typical design of a music-text embedding model consists of the following components: two modality-specific base encoders which separately process inputs of the text and audio modality to an intermediate representation space; a fusion or projection module responsible for mapping the intermediate representations to the shared embedding space; and a contrastive loss, through which the model parameters are optimised to encode semantically related audio and text inputs within the same neighbourhood of the embedding space, while pushing apart unrelated items. We provide an overview of this design in Figure 1. While prior works have converged towards standard choices for the last two components, it remains unclear how to reliably choose unimodal encoders among several existing options. We look at this in Section 3, before discussing the role of training data in Section 4.

### 2.1 Our experimental approach

Before delineating our areas of focus, we outline here the standard experimental setup used in our experiments.

**Projection module** We design our experiments to compare variations of the dual-encoder contrastive architecture described above, varying several components, but keeping two fixed throughout: the projection module and the loss. Similarly to [16, 24], we adopt a two-head, two-layer Transformer as our projection module. From a sequence of 256-dimensional embeddings produced by each projection head, we employ the [CLS] token embedding as the global representation for each branch. For ease of

reference, we denote this model architecture by DuET-MC (**D**ual-**E**ncoder **T**ext-**M**usic **C**ontrastive).

**Training** We optimise our network via the multimodal formulation of the InfoNCE loss [19], using cosine similarity between the  $l_2$ -normalised projection embeddings from the audio and text branch as our scoring function, and a temperature parameter of 0.03. As part of our training procedure, we use the Adam optimizer with decoupled weight decay of 0.05, varying our learning rate through a cosine decay schedule from its peak value of  $1e-3$ , after a linear warm-up of 5 epochs. We train on 8 A100 NVIDIA GPUs, with an effective batch size of 1024 or 2048 based on memory requirements, for a maximum of 100 epochs, with early stopping based on the validation loss. Unless otherwise specified, our default training data is a corpus of licensed instrumental music with high-quality, manually curated genre, mood, and instrument tags, which we refer to as MusicTextHQ. For training, we select a subset totalling a duration of 100 hours, and augment tags into captions following our data augmentation strategy described in Section 4.

### 2.2 Evaluation

We evaluate all our models on text-based music retrieval, as this represents the most prominent task for music-text embedding models and has been shown to correlate to performance on other tasks [11, 14]. Retrieval is performed by ranking all audio clips in the dataset by decreasing cosine similarity of their embedding with the embedding of a text query. From this, we compute Recall@ $k$  ( $R@k$ ), the average number of times the target appears within the top- $k$  retrieved items, and Median Rank (MR). To normalise performance scores by the different dataset sizes, we repeat this procedure on random subsets of 500 items, and report the average value for each metric. When reporting a single metric, we always refer to  $R@10$ .

**Datasets** In order to robustly measure performance across our experiments, we adopt a multi-dataset evaluation suite comprising three public datasets containing audio tracks paired with human-written captions: YT8M-MusicTextClips (MTC) [16], MusicCaps [2] and Song De-

Dataset	Hours*	Tags	Captions
<i>Training</i>			
LP-MusicCaps [7] (A)	50	Human	Synthetic
MusicTextHQ (B)	100	Human	Synthetic
YT8M-MV [1] (C)	270	Synthetic	Synthetic
<i>Evaluation</i>			
YT8M-MTC [16]	8	-	Human
MusicCaps [2]	8	-	Human
Song Describer [15]	2	-	Human

**Table 1: Overview of the datasets** used in our experiments. \*Hours denotes the audio duration used in training.

Encoder	# Params	Model version
<i>Audio</i>		
HTS-AT [5]	30M	AudioSet <sup>1</sup>
MERT [33]	330M	MERT-v1-330M
<i>Text</i>		
RoBERTa [13]	125M	roberta-base
CLIP-T [20]	151M	clip-vit-base-patch32
T5 [21]	11.3B	flan-t5-xx
mT5 [30]	13B	mt5-xx

**Table 2: Audio and text encoders** we compare in our experiments on the impact of encoder backbones (Section 3).

scriber (SDD) [15]. These all represent out-of-distribution data (see Table 1), with different degrees and types of distribution shifts in both the audio and text modality. For example, MTC and MC both contain 10-second audio clips from YouTube videos, but they differ significantly in their captions, with respect to content, descriptiveness and even text length [15]. Audio in the SDD consists instead of music recordings from the music platform Jamendo [3], while captions describe much longer audio segments.


### 3. THE ROLE OF ENCODER BACKBONES

We experiment with two audio encoders, HTS-AT [5] and MERT [33], and three text encoders, RoBERTa [13], the text encoder from CLIP [20] (CLIP-T), T5 [21] and mT5 [30]. We choose these either because they represent the state of the art in their respective tasks, or because they have been previously used in contrastive audio-text learning, thus allowing for direct comparison with prior work.

#### 3.1 Encoders: initialization and freezing

In this set of experiments, our goal is to study parameter-efficient configurations of existing audio and text encoders, training only a subset of the model weights. The motivation for exploring this setting is threefold: freezing part of the model lowers the memory budget and training time, it avoids catastrophic forgetting [17], and it reduces the risk of overfitting in data-constrained scenarios. To fulfil these requirements, we do not consider end-to-end finetuning, and instead focus on leveraging pre-training, locking the audio and text encoders based on their parameter size.

<sup>1</sup> We use the HTSAT\_AudioSet\_Saved\_6 checkpoint of HTS-AT trained on AudioSet from the official repository.



**Figure 2: Retrieval performance (R@10) of different combinations of audio and text encoders** compared through the lens of our DuET-MC framework.

Specifically, we keep all text encoders frozen, as these all count over 100M parameters, as shown in Table 2, and only train the full audio encoder, both with and without general-audio pre-training when using HTS-AS, due to its smaller size. When using MERT, we keep the encoder frozen but train a learnable aggregator over the hidden states of each layer, implemented as a 1D convolutional layer, to obtain audio representations that capture the different levels of abstraction encoded at different depths of the network [33].

**Results** From Figure 2 we first observe that, under the constraints described above, the overall best configuration is given by MERT and CLIP-T. We attribute this to two main reasons: with regards to the audio branch, the superior performance exhibited by MERT suggests that a larger model capacity and stronger music prior may be beneficial to music-text alignment; with regards to the text branch, while all encoders are characterised by large-scale pre-training, CLIP-T stands out as the only model with multimodal capabilities. Although this is somewhat surprising, as CLIP is pre-trained on image-text pairs, we note that prior work has also shown that it can be successfully transferred to the audio and music domains [6, 16, 28, 32]. Secondly, when using MERT with any of the text encoders considered, we find that we can train less than 1% of the total amount of weights (~ 3M making up the projection and aggregation layers) without loss of performance compared to current state-of-the-art models (shown later in Table 4). This demonstrates that we can successfully align locked text representations to the audio modality through lightweight music-text contrastive learning, confirming that our encoder locking strategy is effective when leveraging powerful music-specific pre-training, in line with similar findings in the visual domain [35]. With regards to the audio branch initialization, comparing the two variants of HTS-AT, we find that general-purpose audio pre-training can give a slight advantage over training from scratch, but this benefit is not consistent across the different text encoders HTS-AT is paired with. In the rest of the paper we fix the encoder configuration to locked MERT + locked CLIP-T in all experiments, unless otherwise specified.

Language	R@10		
	YT8M-MTC	MusicCaps	Song Describer
English	10.43	<b>16.00</b>	<b>19.00</b>
German	9.90	13.28	18.00
French	11.71	12.32	15.40
Italian	10.43	13.68	15.80
Spanish	<b>11.60</b>	13.48	18.40

**Table 3: Multilingual retrieval performance.**

### 3.2 Supporting retrieval in multiple languages


Due to a lack of data in different languages, music-text modelling has so far exclusively focussed on English. Real-world applications for music-text embeddings, however, can greatly benefit from the support of multiple languages. To address this limitation, we explore the use of pre-trained locked encoders, similarly to Section 3.1, this time adopting mT5 [30], a multilingual text-to-text Transformer model, as our text encoder. To evaluate multilingual performance, we choose a subset of four languages, German, French, Italian and English, and translate our evaluation datasets via GPT3.5-turbo [4]. In Table 3, we show that this approach provides a viable solution to text-based retrieval in multiple languages while using only English text paired with music in training and with only a minor drop in performance compared to English.

## 4. THE ROLE OF TRAINING DATA

Having established best practices with respect to choosing audio and text backbones, we now shift our attention to the training data. As widely acknowledged in the literature [7, 8, 14], a major limitation in training music-language models is the lack of large public datasets with paired audio-text data. To circumvent this issue, a number of works have proposed to employ large language models to augment text data more commonly found in music datasets, such as categorical labels, metadata and tags, into full natural language sentences, corresponding to *pseudo-captions* [7, 10, 16]. In the next section we present our investigation of the impact of tag-to-caption augmentation.

### 4.1 Tag-to-caption augmentation via LLMs

Following [16], we leverage the in-context learning ability of LLMs via few-shot prompting, and adopt a similar approach to augment tags into captions for our training dataset MusicTextHQ. For this, we use BLOOM-176B [23], a competitive open-access LLM trained on responsibly sourced data. Differently from [16], we do not employ synthetic tags, but use tags provided by expert annotators. We compare this to training on LP-MusicCaps-MTT [7] (LP-MusicCaps for short), a dataset obtained via a similar approach, where tags from the MagnaTagATune [12] dataset are augmented into captions via GPT3.5-turbo. To measure the impact of tag-to-caption augmentation, we train three variants of our model on each dataset, varying  $p_{cap}$ , the probability of selecting captions over tags as the text input for each training pair.



**Figure 3: The effect of varying  $p_{cap}$ , the probability of swapping tags with captions. On the y-axis, we show the relative change in performance compared to  $p_{cap} = 0$ .**

**Results** Results are shown in Figure 3, where we compare the effect of gradually shifting from tags to captions in the two training datasets considered. We first note that introducing tag-to-caption augmentation for at least a portion of the training data ( $p_{cap} = 0.5$ ) leads to an improvement regardless of training dataset. Interestingly, unlike in MusicTextHQ, this trend does not extend to the scenario where we replace all text inputs with pseudo-captions ( $p_{cap} = 1$ ) in LP-MusicCaps. In this case, we observe instead a slight degradation in performance on two of the evaluation datasets, compared to using only tags, or using captions half of the time. We posit that this divergence may be due to a gap in label quality between the two training sets, as pseudo-captions in LP-MusicCaps are generated based on sparse labels, with 50% of the items in the dataset paired to only three tags or less, and in many cases without being balanced across categories. Intuitively, this is likely to result in non-descript, or even inaccurate captions, as the LLM generation will be more prone to hallucinations and may therefore deviate substantially from the audio content. In contrast, MusicTextHQ provides strong grounding to the audio content, with multiple expert-provided tags per category (often three or more for *each* tag category). From this, we conclude that, while LLM-enabled text augmentation can provide a valuable strategy for enriching training data, it is not a substitute for adequate data curation, but rather a supplement. This is an important observation, as prior work has also found that specificity in captions is instrumental to effective multimodal contrastive learning [22, 31]. Since LLM-based augmentation, being bounded by the information content in the source data, cannot increase specificity, our results highlight an often overlooked shortcoming of synthetic text.

### 4.2 Training data: size vs quality

Next, we ask whether simply increasing dataset size can emphasise the benefits of tag-to-caption augmentation. To scale up the size of our training data, we include YT8M-MV, a subset of the YouTube8M dataset [1] tagged as *music video*, as an additional dataset to our training pool. For this, we follow [16] and employ tags from an automatic



music tagger and pseudo-captions generated following the same procedure described in Section 4.1. For simplicity, we refer to LP-MusicCaps, MusicTextHQ and YT8M-MV as *Dataset\_A* (or simply A), *Dataset\_B* (B) and *Dataset\_C* (C), ordered by size as shown in Table 1. We also consider combining the two biggest datasets (B + C) and all three together (A + B + C). We note that each dataset differs not only in size, but also in audio and label quality.

**Results** In Figure 4 we showcase results from training on the datasets described above. Notably, we find that scaling dataset size does not consistently result in an improvement, signalling that the gap in quality between datasets can eclipse their size difference. Although we observe that combining all datasets yields better performance, likely due to overall increased diversity in the training data, the difference is not proportionate to the rise in training cost necessary to scale up. Instead, our results underscore the importance of data curation as a more efficient way of boosting performance, confirming that constructing a subset of highly curated examples, with descriptive and accurate captions, more positively contributes to learning in the contrastive setting [22].

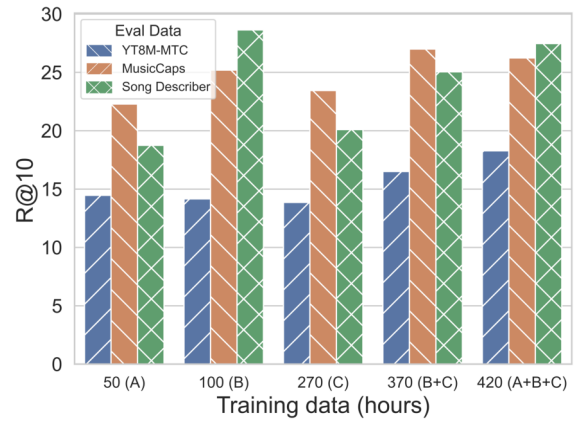
### 5. IMPROVING DIVERSITY VIA TEXT AUGMENTATIONS

Having established that augmenting high-quality tags into captions offers a useful and inexpensive strategy to enrich training data, we explore this further and propose two augmentation-based techniques aimed at increasing data diversity and model robustness.

#### 5.1 Augment, Drop & Swap

**Augmented View Dropout** First, building upon the tag-to-caption strategy described in Section 4.1, we explore text augmentation with the goal of constructing more effective views for contrastive learning, following the principle that optimal views should minimise mutual information between paired items while retaining a high degree of semantic alignment [25]. To this end, we propose *Augmented View Dropout*, where, for each item in our dataset, we randomly sample a subset of the tags, balanced by category (genre, mood, instrumentation) and produce a set of 10 different captions. Each can be thought of as a complementary, but partial view of the associated music track, as we mask a subset of all the ground-truth tags to produce each view. At training time, views are randomly sampled, effectively resulting in a further form of data augmentation.

**Hard negatives via TextSwap** Finally, we tackle another important challenge in contrastive learning, hard negative sampling, and propose to also address this through the lens of text augmentation, via a technique which we call TextSwap. In order to increase the rate of hard negatives beyond the natural rate found in the dataset, we create partially perturbed versions of the captions by stochastically swapping genre, mood or instrument keywords with alternative descriptors from a predefined dictionary (e.g. “a mellow *pop* track” becomes “a mellow *hip-hop* track”).



**Figure 4:** Retrieval performance across models trained on datasets that differ in size and annotation quality.

During training, for each positive pair, we then select a random subset of the negative captions in a batch and replace them with hard negatives by applying TextSwap once per descriptor category. This is illustrated in Figure 1, where we provide a visual guide for the full *Augment, Drop & Swap* pipeline. We hypothesise that the presence of hard negatives is particularly critical in later stages of training, once the model has already acquired basic features, and learning on “easy” negatives has saturated. Based on this, we follow a curriculum learning approach and linearly increase the probability of applying TextSwap from 0 to 15% over the course of 20 epochs, after a warm-up period of 5.

#### 5.2 Experiments

**Ablations** In this set of experiments we examine the effect of each of the three components in our augmentation pipeline: tag-to-caption augmentation, Augmented View Dropout and TextSwap. We look at two scenarios: one where we want to measure their contribution in training two variants of our parameter-efficient DuET-MC framework, each with different degrees of audio pre-training and finetuning and locked text encoders, and one where we relax our computational requirements and explore whether our proposed method can be usefully applied in finetuning a general purpose audio-text embedding model (CLAP [29]), with limited paired music data.

**Results** We present our ablations on the proposed pipeline in Table 4, where we also compare to two audio-text contrastive baselines, CLAP [29] and TTMR [8], trained on general-purpose audio and music respectively. The table displays three different settings to which we apply our proposed pipeline: (1) training the audio encoder from scratch (shown in the HTS-AT + CLIP-T configuration), (2) training only 1% of the parameters in our locked audio-text encoder (MERT + CLIP-T), and (3) fine-tuning the full model on music, following general audio-text pre-training (CLAP-FT). From this, we observe that, while the vanilla version of DuET-MC (trained only on tags) exhibits at best comparable performance to the baselines, each additional component in our pipeline lifts performance across all model configurations, pre-training regimes and finetuning strategies. Among these, tag-to-caption augmentation

Model	Tag-to-caption	Augmented View Dropout	TextSwap	YT8M-MTC		MusicCaps		Song Describer		Avg R@10 ↑
				R@10 ↑	MR ↓	R@10 ↑	MR ↓	R@10 ↑	MR ↓	
<i>Baselines</i>										
CLAP [29]	-	-	-	11.9	80	40.3*	17*	19.8	53	24.0*
TTMR [8]	-	-	-	11.6	79	9.6	115	16.5	57	12.6
DuET-MC	✗	✗	✗	8.5	103	12.2	82	15.3	53	12.0
(HTS-AT + CLIP-T)	✓	✗	✗	8.0	104	13.4	76	14.1	57	11.8
	✓	✓	✗	<u>9.4</u>	<u>93</u>	15.1	<u>65</u>	<u>19.6</u>	49	<u>14.7</u>
	✓	✓	✓	<u>9.4</u>	<u>93</u>	<u>15.8</u>	66	17.4	<u>48</u>	14.2
DuET-MC	✗	✗	✗	10.8	82	18.3	56	20.2	45	16.4
(MERT + CLIP-T)	✓	✗	✗	11.7	69	21.3	41	23.4	36	18.8
	✓	✓	✗	13.4	65	<u>24.9</u>	36	<b>27.7</b>	32	22.0
	✓	✓	✓	<u>14.5</u>	<u>62</u>	24.6	<u>34</u>	<u>27.3</u>	<b>29</b>	<u>22.1</u>
CLAP-FT	✗	✗	✗	14.2	63	38.8*	18*	20.8	38	24.6*
	✓	✗	✗	14.6	61	42.3*	15*	23.5	34	26.8*
	✓	✓	✗	<b>16.3</b>	<b>55</b>	41.6*	16*	24.5	36	27.3*
	✓	✓	✓	15.7	57	<b>43.5*</b>	<b>14*</b>	26.3	<u>31</u>	<b>28.5*</b>

**Table 4: Ablations.** For each model, subsequent rows show the effect of introducing an additional step in our proposed *Augment, Drop & Swap* pipeline. We highlight best results for each model (underlined) and amongst all models (bold). \* denotes values that may be inflated due to in-distribution bias.

and Augmented View Dropout emerge as the most influential, while the benefits of TextSwap are more prominent for model configurations where encoders have higher levels of pre-training, hinting at the necessity to increase the complexity of negatives later in training. This suggests that our *Augment, Drop & Swap* recipe provides a data-efficient strategy to improve music-text modelling under a variety of model configurations, at no additional computational cost. Importantly, this trend generalises across evaluation datasets, suggesting that it is beneficial to model robustness, and demonstrates that the lack of large-scale paired data in the music domain can be alleviated through augmentation-based techniques which enhance data quality instead of quantity. Finally, comparing retrieval scores of different family of models (TTMR, CLAP and DuET-MC), we note consistent differences between datasets, with CLAP-based models invariably showing a significant jump in performance on the MusicCaps dataset compared to MTC and SDD. We hypothesise that this may be a result of in-distribution bias, since there are several instances of non-music or noisy, low-quality recordings in MC. Since CLAP is trained to recognise everyday sounds, this points at a smaller shift from its training distribution, compared to SDD and MTC, which are exclusively composed of music recordings. We posit that further mismatches in the training and test distributions exist along the text dimension and investigate this through human evaluation.

**Are metrics aligned with human preference?** We recruit 35 participants to evaluate DuET-MC, CLAP and TTMR in a head-to-head pairwise comparison. Participants are presented with up to 24 text prompts, where each is a caption taken from one of the three evaluation datasets, and are asked to choose which one of two music tracks best aligns to the description. Through this qualitative evaluation, we find that DuET-MC does substantially better than TTMR, losing against it only 30.9% of the times, and largely mirroring our findings from Section

5.2. Surprisingly, the win and tie rate vs CLAP drops instead to 37.3% and 38.5% respectively. Looking at the breakdown of scores by dataset, this advantage in CLAP is predominantly observed on MC and MTC, while DuET-MC outperforms CLAP on SDD. Interestingly, DuET-MC is preferred or considered equivalent to the ground truth 38.9% of the times on MTC compared to 15.4 and 17.9% on the other two datasets. This points to significant differences in the level of alignment between caption and audio in the different datasets, signalling that evaluating on several datasets is paramount to understanding real-world performance. Additionally, it leads to an observation that complements our automatic evaluation in Table 4: the discrepancy between DuET-MC’s performance on MTC compared to MC and SDD may be ascribed to a higher degree of *vagueness* in MTC captions, which, as revealed through our qualitative evaluation, admit instead alternative matching tracks to those in the ground truth.

## 6. CONCLUSIONS

In this work we presented *Augment, Drop & Swap*, a training recipe for efficient music-text representation learning informed by our findings on training music-text contrastive models in resource-constrained scenarios. Through our experiments, we provide a practical guide to this family of models, and foreground their real-world use by focusing on multilingual support, computationally efficient techniques, and cross-dataset evaluation. Showing that data curation has a significant effect at modest data scales, we design each step in our pipeline to tackle specific aspects of the text used in training, such as descriptiveness and specificity, via data augmentations, leading to views that are more effective in multimodal contrastive learning. Through automatic and qualitative evaluations, we show the usefulness of our approach and reveal insights on the relation between measured performance and distribution shifts in the test data.

## 7. REFERENCES

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. “Youtube-8m: A large-scale video classification benchmark”. In: *arXiv preprint arXiv:1609.08675* (2016).
- [2] Andrea Agostinelli et al. “Musiclm: Generating music from text”. In: *arXiv preprint arXiv:2301.11325* (2023).
- [3] Dmitry Bogdanov, Minz Won, Philip Tovstogan, Alastair Porter, and Xavier Serra. “The MTG-Jamendo Dataset for Automatic Music Tagging”. In: *Machine Learning for Music Discovery Workshop, International Conference on Machine Learning (ICML)*. 2019.
- [4] Tom B Brown et al. “Language Models are Few-Shot Learners”. In: *Advances in Neural Information Processing Systems*. 2020.
- [5] Ke Chen, Xingjian Du, Bilei Zhu, Zejun Ma, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. “HTS-AT: A Hierarchical Token-Semantic Audio Transformer for Sound Classification and Detection”. In: *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022.
- [6] Tianyu Chen, Yuan Xie, Shuai Zhang, Shaohan Huang, Haoyi Zhou, and Jianxin Li. “Learning Music Sequence Representation From Text Supervision”. In: *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2022, pp. 4583–4587.
- [7] SeungHeon Doh, Keunwoo Choi, Jongpil Lee, and Juhan Nam. “LP-MusicCaps: LLM-Based Pseudo Music Captioning”. In: *Proceedings of the 24th International Society for Music Information Retrieval (ISMIR) Conference*. Milan, 2023.
- [8] SeungHeon Doh, Minz Won, Keunwoo Choi, and Juhan Nam. “Toward Universal Text-to-Music Retrieval”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022.
- [9] Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. “CLAP Learning Audio Concepts from Natural Language Supervision”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2023, pp. 1–5.
- [10] Josh Gardner, Simon Durand, Daniel Stoller, and Rachel Bittner. “LLark: A Multimodal Instruction-Following Language Model for Music”. In: *Proceedings of the 41st International Conference on Machine Learning (ICML)*. 2024.
- [11] Qingqing Huang, Aren Jansen, Joonseok Lee, Ravi Ganti, Judith Yue Li, and Daniel P. W. Ellis. “MuLan: A Joint Embedding of Music Audio and Natural Language”. In: *23rd International Society for Music Information Retrieval Conference (ISMIR 2022)*. 2022.
- [12] Edith Law, Kris West, Michael Mandel, Mert Bay, and J. Stephen Downie. “Evaluation of algorithms using games: The case of music tagging”. In: *Proceedings of the 10th ISMIR Conference*. 2009.
- [13] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. “Roberta: A robustly optimized bert pretraining approach”. In: *arXiv preprint arXiv:1907.11692* (2019).
- [14] Iliaria Manco, Emmanouil Benetos, Elio Quinton, and Gyorgy Fazekas. “Contrastive Audio-Language Learning for Music”. In: *23rd International Society for Music Information Retrieval Conference (ISMIR 2022)*. 2022.
- [15] Iliaria Manco et al. “The Song Descriptor Dataset: a Corpus of Audio Captions for Music-and-Language Evaluation”. In: *Machine Learning for Audio Workshop at NeurIPS 2023*. 2023.
- [16] Daniel McKee, Justin Salamon, Josef Sivic, and Bryan Russell. “Language-guided music recommendation for video via prompt analogies”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023, pp. 14784–14793.
- [17] Sanket Vaibhav Mehta, Darshan Patil, Sarath Chandar, and Emma Strubell. “An Empirical Investigation of the Role of Pre-training in Lifelong Learning”. In: *Journal of Machine Learning Research* 24.214 (2023), pp. 1–50.
- [18] Ryumei Nakada, Halil Ibrahim Gulluk, Zhun Deng, Wenlong Ji, James Zou, and Linjun Zhang. “Understanding Multimodal Contrastive Learning and Incorporating Unpaired Data”. en. In: *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*. PMLR, 2023, pp. 4348–4380.
- [19] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. “Representation learning with contrastive predictive coding”. In: *arXiv preprint arXiv:1807.03748* (2018).
- [20] Alec Radford et al. “Learning Transferable Visual Models From Natural Language Supervision”. In: *International Conference on Machine Learning*. PMLR, 2021.

- [21] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. “Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer”. In: *Journal of Machine Learning Research* 21.140 (2020), pp. 1–67.
- [22] Shibani Santurkar, Yann Dubois, Rohan Taori, Percy Liang, and Tatsunori Hashimoto. “Is a Caption Worth a Thousand Images? A Study on Representation Learning”. In: *The Eleventh International Conference on Learning Representations*. 2023.
- [23] Teven Le Scao, Angela Fan, and Christopher Akiki. *BLOOM: A 176B-Parameter Open-Access Multilingual Language Model*. 2023.
- [24] Didac Suris, Carl Vondrick, Bryan Russell, and Justin Salamon. “It’s Time for Artistic Correspondence in Music and Video”. en. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2022, pp. 10554–10564.
- [25] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. “What Makes for Good Views for Contrastive Learning?” In: *Advances in Neural Information Processing Systems*. Vol. 33. 2020, pp. 6827–6839.
- [26] Minz Won, Keunwoo Choi, and Xavier Serra. “Semi-Supervised Music Tagging Transformer”. In: *International Society for Music Information Retrieval Conference (ISMIR)*. 2021.
- [27] Minz Won, Justin Salamon, Nicholas J. Bryan, Gautham J. Mysore, and Xavier Serra. “Emotion Embedding Spaces for Matching Music to Stories”. In: *International Society for Music Information Retrieval Conference (ISMIR)*. 2021.
- [28] Ho-Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello. “Wav2CLIP: Learning Robust Audio Representations from Clip”. In: *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2022, pp. 4563–4567.
- [29] Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. “Large-Scale Contrastive Language-Audio Pretraining with Feature Fusion and Keyword-to-Caption Augmentation”. In: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2023, pp. 1–5.
- [30] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. “mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer”. In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, 2021, pp. 483–498.
- [31] Yihao Xue, Siddharth Joshi, Dang Nguyen, and Baharan Mirzasoleiman. “Understanding the Robustness of Multi-modal Contrastive Learning to Distribution Shift”. In: *The Twelfth International Conference on Learning Representations*. 2024.
- [32] Dongchao Yang, Jianwei Yu, Helin Wang, Wen Wang, Chao Weng, Yuexian Zou, and Dong Yu. “DiffSound: Discrete Diffusion Model for Text-to-Sound Generation”. In: *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 31 (2023), pp. 1720–1733.
- [33] Li Yizhi, Ruibin Yuan, Ge Zhang, Yinghao Ma, Xingran Chen, Hanzhi Yin, Chenghao Xiao, Chenghua Lin, Anton Ragni, Emmanouil Benetos, et al. “MERT: Acoustic Music Understanding Model with Large-Scale Self-supervised Training”. In: *The Twelfth International Conference on Learning Representations*. 2023.
- [34] Haoxuan You, Luowei Zhou, Bin Xiao, Noel Codella, Yu Cheng, Ruochen Xu, Shih-Fu Chang, and Lu Yuan. “Learning visual representation from modality-shared contrastive language-image pre-training”. In: *European Conference on Computer Vision*. Springer, 2022, pp. 69–87.
- [35] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. “Lit: Zero-shot transfer with locked-image text tuning”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2022, pp. 18123–18133.
- [36] Qi Zhang, Yifei Wang, and Yisen Wang. “On the generalization of multi-modal contrastive learning”. In: *International Conference on Machine Learning*. PMLR. 2023, pp. 41677–41693.