

GENERATIVE TIMBRE SPACES: REGULARIZING VARIATIONAL AUTO-ENCODERS WITH PERCEPTUAL METRICS

Philippe Esling^{1*}, Axel Chemla-Romeu-Santos^{2,1}, Adrien Bitton¹

¹IRCAM, UMR 9912 CNRS - Sorbonne Universite, Paris, France

²Laboratorio d'Informatica Musicale - Universita degli Studi di Milano, Italy

esling@ircam.fr

ABSTRACT

Timbre spaces have been used in music perception to study the perceptual relationships between instruments based on dissimilarity ratings. However, these spaces do not generalize to novel examples and do not provide an invertible mapping, preventing audio synthesis. In parallel, generative models have aimed to provide methods for synthesizing novel timbres. However, these systems do not provide an understanding of their inner workings and are usually not related to any perceptually relevant information. Here, we show that Variational Auto-Encoders (VAE) can alleviate all of these limitations by constructing *generative timbre spaces*. To do so, we adapt VAEs to learn an audio latent space, while using perceptual ratings from timbre studies to regularize the organization of this space. The resulting space allows us to analyze novel instruments, while being able to synthesize audio from any point of this space. We introduce a specific regularization allowing to enforce any given similarity distances onto these spaces. We show that the resulting space provide almost similar distance relationships as timbre spaces. We evaluate several spectral transforms and show that the Non-Stationary Gabor Transform (NSGT) provides the highest correlation to timbre spaces and the best quality of synthesis. Furthermore, we show that these spaces can generalize to novel instruments and can generate any path between instruments to understand their timbre relationships. As these spaces are continuous, we study how audio descriptors behave along the latent dimensions. We show that even though descriptors have an overall non-linear topology, they follow a locally smooth evolution. Based on this, we introduce a method for *descriptor-based synthesis* and show that we can control the descriptors of an instrument while keeping its timbre structure.

1. INTRODUCTION

For the past decades, music perception research has tried to understand the perception of instrumental *timbre*. Timbre is the set of properties that distinguishes two instruments that play the same note at the same intensity. To do so, several studies [1] collected human dissimilarity ratings between pairs of audio samples inside a set of instruments. These ratings are organized by applying MultiDimensional Scaling (MDS), leading to *timbre spaces*, which exhibit the perceptual similarities between different instruments. By analyzing the dimensions of resulting spaces, the studies tried to correlate audio descriptors to the perception of timbre [2]. Although these spaces provided interesting avenues of analysis, they are inherently limited by the fact that ordination techniques (e.g.

MDS) produce a fixed space, which has to be recomputed entirely for any new sample. Therefore, these spaces do not generalize to novel examples and do not provide an invertible mapping, precluding audio synthesis to understand their perceptual topology.

In parallel, recent developments in audio synthesis using *generative models* has seen great improvements with the introduction of approaches such as the *WaveNet* [3] and *SampleRNN* [4] architectures. These allow to generate novel high-quality audio matching the properties of the corpus they have been trained on. However, these models give little cue and control over the output or the features it results from. More recently, *NSynth* [5] has been proposed to synthesize audio by allowing to morph between specific instruments. However, these models still require very large number of parameters, long training times and a large number of examples. Amongst recent generative models, another key proposal is the *Variational Auto-Encoder* (VAE) [6]. In these, a *latent space* is learned that allows both to encode data for analysis, but also to sample from it in order to generate novel content. VAEs address the limitations of control and analysis through this latent space, while remaining simple and fast to learn with a small set of examples. Furthermore, VAEs seem able to disentangle underlying variation factors by learning independent latent variables accounting for distinct generative processes [7]. However, these latent dimensions are learned in an unsupervised way. Therefore, they are not related to perceptual properties, which might hamper their understandability or their use for audio analysis and synthesis.

Here, we show that we can bridge timbre perception analysis and perceptually-relevant audio synthesis by regularizing the learning of VAE latent spaces so that they match the perceptual distances collected from timbre studies. Our overall approach is depicted in Figure 1. First, we adapt the VAE to analyze musical audio content, by comparing the use of different spectral transforms as input to the learning. We show that, amongst the Short-Term Fourier Transform (STFT), Discrete Cosine Transform (DCT) and the Non-Stationary Gabor Transform (NSGT) [8], the NSGT provides the best reconstruction abilities and regularization performances. By training this model on a small database of spectral frames, it already provides a generative model with an interesting latent space, able to synthesize novel instrumental timbres. Then, we introduce a regularization to the learning objective inspired by the t-Stochastic Neighbors Embedding (t-SNE) [9], aiming to enforce that the latent space exhibits the same distances between instruments as those found in timbre studies. To do so, we build a model of perceptual relationships by analyzing dissimilarity ratings from five independent timbre studies [10, 11, 12, 13, 14]. We show that perceptually-regularized latent spaces are simultaneously coherent with perceptual ratings, while being able to synthesize high-quality audio distributions. Hence, we drive the learning of latent spaces to match the topology of given target spaces.

* This work was supported by project MAKIMOno 17-CE38-0015-01 funded by the French ANR and Canadian NSERC (STPG 507004-17) and the ACTOR Partnership funded by the Canadian SSHRC (895-2018-1023).

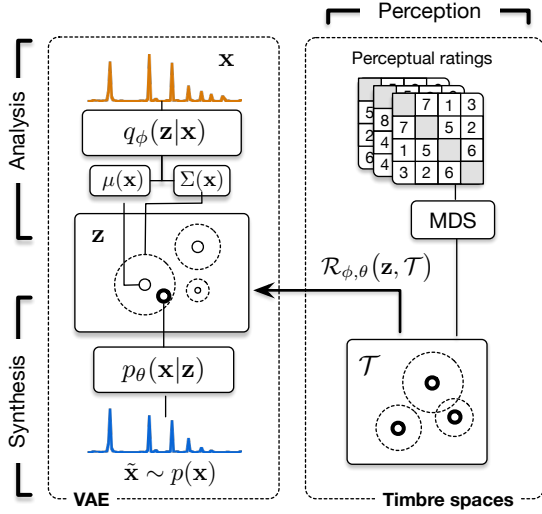


Figure 1: (Left) VAEs can model a spectral frame \mathbf{x} of an audio sample by learning an encoder $q_\phi(\mathbf{z}|\mathbf{x})$ which maps them to a Gaussian $\mathcal{N}(\mu(\mathbf{x}), \sigma(\mathbf{x}))$ inside a latent space \mathbf{z} . The decoder $p_\theta(\mathbf{x}|\mathbf{z})$ samples from this Gaussian to generate a reconstruction $\tilde{\mathbf{x}}$ of the spectral frame. (Right) Perception studies use similarity ratings to construct *timbre spaces* exhibiting perceptual distances between instruments. Here, we develop a regularization $\mathcal{R}(\mathbf{z}, \mathcal{T})$ enforcing that the variational model finds a topology of latent space \mathbf{z} that matches the topology of the timbre space \mathcal{T} .

We demonstrate that these spaces can be used for generating novel audio content, by analyzing their reconstruction quality on a test dataset. Furthermore, we show that paths in the latent space (where each point corresponds to a single spectral frame) provide sound synthesis with continuous evolutions of timbre. We also show that these spaces generalize to novel samples, by encoding a set of instruments that were not part of the training set. Therefore, the spaces could be used to predict the perceptual similarities of novel instruments. Finally, we study how traditional audio descriptors are organized along the latent dimensions. We show that even though descriptors behave in a non-linear way across space, they still follow a locally smooth evolution. Based on this smoothness property, we introduce a method for *descriptor-based path synthesis*. We show that we can modify an instrumental distribution so that it matches a given target evolution of audio descriptors, while remaining perceptually smooth. The source code, audio examples and animations are available on a supporting repository¹.

2. STATE-OF-ART

2.1. Variational auto-encoders

Generative models are a flourishing class of learning approaches, which aim to find the underlying probability distribution of the data $p(\mathbf{x})$ [15]. Formally, based on a set of examples in a high-dimensional space $\mathbf{x} \in \mathbb{R}^{d_x}$, we assume that these follow an un-

known distribution $p(\mathbf{x})$. Furthermore, we consider a set of *latent variables* defined in a lower-dimensional space $\mathbf{z} \in \mathbb{R}^{d_z}$ ($d_z \ll d_x$). These latent variables help govern the generation of the data and enhance the *expressivity* of the model. Thus, the complete model is defined by the joint probability distribution $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$. We could find $p(\mathbf{x})$ through its relation to the posterior distribution $p(\mathbf{z}|\mathbf{x})$ given by Bayes' theorem. However, for complex non-linear models (such as those that we will consider in this paper), this posterior can not be found in closed form.

For decades, the dominant paradigm for approximating $p(\mathbf{x})$ has been *sampling* methods [16]. However, the quality of this approximation depends on the number of sampling operations, which might be extremely large before we have an accurate estimate. Recently, *variational inference* (VI) [15] has been proposed to solve this problem through *optimization* rather than sampling. VI assumes that if the distribution is too complex to find, we could find a simpler approximate distribution that still models the data, while trying to minimize its difference to the real distribution. Formally, VI specifies a family \mathcal{Q} of approximate densities, where each member $q(\mathbf{z}|\mathbf{x}) \in \mathcal{Q}$ is a candidate approximation to the exact $p(\mathbf{z}|\mathbf{x})$. Hence, the inference problem can be transformed into an optimization problem by minimizing the Kullback-Leibler (KL) divergence between the approximation and original density

$$q^*(\mathbf{z}|\mathbf{x}) = \arg \min_{q(\mathbf{z}|\mathbf{x}) \in \mathcal{Q}} \mathcal{D}_{KL}[q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}|\mathbf{x})] \quad (1)$$

The complexity of the family \mathcal{Q} will both determine the quality of the approximation, but also the complexity of this optimization. Hence, the major issue of VI is to choose \mathcal{Q} to be flexible enough to closely approximate $p(\mathbf{z}|\mathbf{x})$, while being simple enough to allow efficient optimization. Now, if we expand the KL divergence that we need to minimize and rely on Bayes' rule to replace $p(\mathbf{z}|\mathbf{x})$, we obtain the following expression

$$D_{KL}[q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}|\mathbf{x})] = \mathbb{E}_{q(\mathbf{z})} [\log q(\mathbf{z}|\mathbf{x}) - \log p(\mathbf{x}|\mathbf{z}) - \log p(\mathbf{z}) + \log p(\mathbf{x})] \quad (2)$$

Noting that the expectation is over $q(\mathbf{z}|\mathbf{x})$ and that $p(\mathbf{x})$ does not depend on it, we can get this term out of the expectation and then observe that the remaining equation can be rewritten as another KL divergence leading to

$$\log p(\mathbf{x}) - D_{KL}[q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}|\mathbf{x})] = \mathbb{E}_{\mathbf{z}} [\log p(\mathbf{x}|\mathbf{z})] - D_{KL}[q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})] \quad (3)$$

This formulation describes the logarithm of the quantity that we want to maximize $\log p(\mathbf{x})$ minus the error we make by using an approximate q instead of p . Therefore, we can optimize this alternative objective, called the *evidence lower bound* (ELBO) as

$$\log p(\mathbf{x}) = D_{KL}[q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}|\mathbf{x})] + ELBO(q). \quad (4)$$

and the KL is non-negative, so $\log p(\mathbf{x}) \geq ELBO(q), \forall q(\mathbf{z})$. Now, to optimize this objective, we will rely on parametric distributions $q_\phi(\mathbf{z}|\mathbf{x})$ and $p_\theta(\mathbf{x}|\mathbf{z})$. Therefore, optimizing our generative model will amount to optimize these parameters $\{\theta, \phi\}$

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi(\mathbf{z})} [\log p_\theta(\mathbf{x}|\mathbf{z})] - D_{KL}[q_\phi(\mathbf{z}|\mathbf{x}) \| p_\theta(\mathbf{z})] \quad (5)$$

We can see that this equation involves $q_\phi(\mathbf{z}|\mathbf{x})$ which *encodes* the data \mathbf{x} into the latent representation \mathbf{z} and a *decoder* $p(\mathbf{x}|\mathbf{z})$,

¹<https://github.com/acids-ircam/variational-timbre>

which generates a data \mathbf{x} given a latent configuration \mathbf{z} . Hence, this whole structure defines the *Variational Auto-Encoder* (VAE), which is depicted in Figure 1 (Left).

The VAE objective can be interpreted intuitively. The first term increases the likelihood of the data generated given a configuration of the latent, which amounts to minimize the *reconstruction error*. The second term represents the error made by using a simpler distribution $q_\phi(\mathbf{z} | \mathbf{x})$ rather than the true distribution $p_\theta(\mathbf{z})$. Therefore, this *regularizes* the choice of approximation q so that

$$\mathcal{L}_{\theta,\phi} = \underbrace{\mathbb{E}_{q_\phi(\mathbf{z})} [\log p_\theta(\mathbf{x} | \mathbf{z})]}_{\text{reconstruction}} - \beta \cdot \underbrace{D_{KL}[q_\phi(\mathbf{z} | \mathbf{x}) \parallel p_\theta(\mathbf{z})]}_{\text{regularization}} \quad (6)$$

The first term can be optimized through a usual maximum likelihood estimation, while the second term requires that we define the prior $p(\mathbf{z})$. While the easiest choice is to choose $p(\mathbf{z}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, it also adds the benefit that this term has a simple closed solution for computing the optimization, as detailed in [6]. Here we introduced a weight β to the KL divergence, which leads to the β -VAE formulation [7]. This has been shown to improve the capacity of the model to disentangle factors of variations in the data. However, it has later been shown that an appropriate way to handle this parameter was to perform *warm-up* [17], where the β parameter is linearly increased in the first epochs of training.

Finally, we need to select a family of variational densities \mathcal{Q} . One of the most widespread choice is the *mean-field variational family* where latent variables are independent and are each parametrized by a distinct variational parameter

$$q(\mathbf{z}) = \prod_{j=1}^m q_j(z_j) \quad (7)$$

Therefore, each dimension of the latent space will be governed by an independent Gaussian distribution with its own mean and variance depending on the input data $q_j(z_j) = \mathcal{N}(\mu_j(\mathbf{x}), \Sigma_j(\mathbf{x}))$.

VAEs are powerful representation learning frameworks, while remaining simple and fast to learn without requiring large sets of examples [17]. Their potential for audio applications have been only scarcely investigated yet and mostly in topics related to speech processing such as blind source separation [18] and speech transformation [19]. However, to the best of our knowledge, the use of VAE and their latent spaces to perform musical audio analysis and generation has yet to be investigated.

2.2. Timbre spaces and auditory perception

For several decades, music perception research has tried to understand the mechanisms leading to the perception of *timbre*. Several studies have shown that timbre could be partially described by computing various audio descriptors [13]. To do so, most studies relied on the concept of *timbre spaces* [2], a model that organize audio samples based on perceptual dissimilarity ratings. In these studies, pairs of sounds are presented to subjects that are asked to rate their perceptual dissimilarities inside a given set of instruments. Then, these ratings are compiled into a set of dissimilarity matrices that are analyzed with Multi-Dimensional Scaling (MDS). The MDS algorithm provides a timbre space that exhibits the underlying perceptual distances between different instruments (Figure 1 (Right)). Here, we briefly detail corresponding studies and redirect interested readers to the full articles for more details.

In his seminal paper, Grey [10] performed a study with 16 instrumental sound samples. Each of the 22 subjects had to rate the dissimilarity between all pairs of sounds on a continuous scale from 0 (most similar) to 1 (most dissimilar). This led to the first construction of a timbre space for instrumental sounds. They further exhibit that the dimensions explaining these dissimilarities could be correlated to the spectral centroid, spectral flux and attack centroid. Several studies followed this research by using the same experimental paradigm. Krumhansl [11] used 21 instruments with 9 subjects on a discrete scale from 1 to 9, Iverson et al. [12] with 16 samples and 10 subjects on a continuous scale from 0 to 1, McAdams et al. [13] with 18 orchestral instruments and 24 subjects on a discrete scale from 1 to 16 and, finally, Lakatos [14] with 17 subjects on 22 harmonic and percussive samples on a continuous scale from 0 to 1. Each of these studies shed light on different aspects of audio perception, depending on the aspect being scrutinized and the interpretation of the space by the experimenters. However, all studies have led to different spaces with different dimensions. The fact that different studies correlate to different audio descriptors prevents a generalization of the acoustic cues that might correspond to timbre dimensions. Furthermore, timbre spaces have been explored based on MDS to organize perceptual ratings and correlate spectral descriptors [13]. Therefore, these studies are inherently limited by the fact that

- ordination techniques (such as MDS) produce fixed spaces that must be recomputed for any new data point
- these spaces do not generalize nor synthesize audio between instruments as they do not provide an invertible mapping
- interpretation is bounded to the *a posteriori* linear correlation of audio descriptors to the dimensions rather than analyzing the topology of the space itself

As noted by McAdams et al. [1], critical problems in these approaches are the lack of an objective distance model based on perception and general dimensions for the interpretation of timbral transformation and source identification. Here, we show that relying on VAE models to learn unsupervised spaces, while regularizing the topology of these spaces to fit given perceptual ratings can allow to alleviate all of these limitations.

3. REGULARIZING LATENT SPACE TOPOLOGY

In this paper, we aim to construct a latent space that could both analyze and synthesize audio content, while providing the underlying perceptual relationships between audio samples. To do so, we show that we can influence the organization of the VAE latent space \mathbf{z} so that it follows the topology of a given target space \mathcal{T} . Here, we will rely on the MDS space constructed from perceptual ratings as a target space \mathcal{T} . However, it should be noted that this idea can be applied to any given target space that provides a set of distances between the elements used for learning the VAE space.

To further specify our problem, we consider a set of audio samples, where each x_i can be encoded in the latent space as \mathbf{z}_i and have an equivalent in the target space \mathcal{T}_i . In order to relate the elements of the audio dataset to the perceptual space, we consider that each sample is labeled with its instrumental class C_i , that has an equivalent in the timbre space. Therefore, we will match the properties of the classes between the latent and target spaces (note that we could use element-wise properties for finer control).

Here, we propose to regularize the learning by introducing the perceptual similarities through an additive term $\mathcal{R}(\mathbf{z}, \mathcal{T})$. This

penalty imposes that the properties of the latent space \mathbf{z} are similar to that of the target space \mathcal{T} . The optimization objective becomes

$$\mathbb{E}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \beta D_{KL}[q_\phi(\mathbf{z}|\mathbf{x}) \parallel p_\theta(\mathbf{z})] + \alpha \mathcal{R}(\mathbf{z}, \mathcal{T}) \quad (8)$$

where α is an hyper-parameter that allows us to control the influence of the regularization. Hence, amongst two otherwise equal solutions, the model is pushed to select the one that comply with the penalty. In our case, we want the distances between instruments to follow perceptual timbre distances. Therefore, we need to minimize the differences between the set of distances in the latent space $\mathcal{D}_{i,j}^z = \mathcal{D}(\mathbf{z}_i, \mathbf{z}_j)$ and the distances in target space $\mathcal{D}_{i,j}^T = \mathcal{D}(\mathcal{T}_i, \mathcal{T}_j)$. Therefore, the regularization criterion will try to minimize the overall differences between these sets of distances. To compute these sets, we take inspiration from the *t-Stochastic Neighbor Embedding* (t-SNE) algorithm [9]. Indeed, as their goal is to map the distances from one (high-dimensional) space into a target (low-dimensional) space, it is highly correlated to our task. However, we can not simply apply the t-SNE algorithm on the latent space as this would lead to a non-invertible mapping. Instead, we aim to steer the learning in a similar way. Hence, we compute the relationships in the latent space \mathbf{z} by using the conditional Gaussian density that i would choose j as its neighbor

$$\mathcal{D}_{i,j}^z = \frac{\exp(-\|\mathbf{z}_i - \mathbf{z}_j\|^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{z}_i - \mathbf{z}_k\|^2 / 2\sigma_i^2)} \quad (9)$$

where σ_i is the variance of the Gaussian centered on \mathbf{z}_i , defined as $\sigma_i = 1/\sqrt{2}$. Then, to relate the points in the timbre space \mathcal{T} , we use a Student-t distribution to define the distances in this space as

$$\mathcal{D}_{i,j}^T = \frac{(1 + \|\mathcal{T}_i - \mathcal{T}_j\|^2)^{-1}}{\sum_{k \neq i} (1 + \|\mathcal{T}_i - \mathcal{T}_k\|^2)^{-1}} \quad (10)$$

Finally, we rely on the sum of KL divergences between the two distributions of distances in different spaces to define our complete regularization criterion

$$\mathcal{R}(\mathbf{z}, \mathcal{T}) = \sum_i \mathcal{D}_{KL}[\mathcal{D}_i^z \parallel \mathcal{D}_i^T] = \sum_i \sum_j \mathcal{D}_{i,j}^z \log \frac{\mathcal{D}_{i,j}^z}{\mathcal{D}_{i,j}^T}$$

Hence, instead of applying a distance minimization a posteriori, we steer the learning to find a configuration of the latent space \mathbf{z} that displays the same distance properties as the space \mathcal{T} , while providing an invertible mapping.

4. EXPERIMENTS

4.1. Datasets

Timbre studies. We rely on the perceptual ratings collected across five independent timbre studies [10, 11, 12, 13, 14]. As discussed earlier, even though all studies follow the same experimental protocol, there are some discrepancies in the choice of instruments, rating scales and sound stimuli. However, here we aim to obtain a consistent set of properties to define a common timbre space. Therefore, we computed the maximal set of instruments for which we had ratings for all pairs. To do so, we collated the list of instruments from all studies and counted their co-occurrences, leading to a set of 12 instruments (Piano, Cello, Violin, Flute, Clarinet, Trombone, French Horn, English Horn, Oboe, Saxophone, Trumpet, Tuba) with pairwise ratings. Then, we normalized the raw

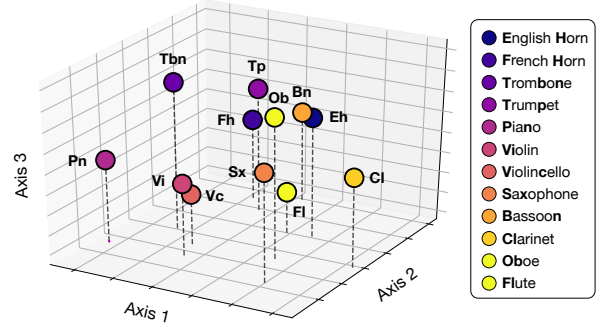


Figure 2: Multi-dimensional scaling (MDS) of the combined and normalized set of perceptual ratings from different studies.

dissimilarity data (keeping all instruments of that study) so that it maps to a common scale from 0 to 1. Finally, we extracted the set of ratings that corresponds to our selected instruments. This leads to a total of 1217 subject ratings for all instruments, amounting to 11845 pairwise ratings. Based on this set of ratings, we compute an MDS space to ensure the consistency of our normalized perceptual space on the selected set. The results of this analysis are displayed in Figure 2. We can see that even though the ratings come from different studies, the resulting space remains very coherent, with the distances between instruments remaining coherent with the original perceptual studies.

Audio datasets. In order to learn the distribution of instrumental sounds directly from the audio signal, we rely on the Studio On Line (SOL) database [20]. We selected 2,200 samples to represent the 11 instruments for which we extracted perceptual ratings. We normalized the range of notes used by taking the whole tessitura and dynamics available (to remove effects from the pitch and loudness). All recordings were resampled to 22050 Hz for the experiments. Then, as we intend to evaluate the effect of different spectral distributions as input to our proposed model, we computed several invertible transforms for each audio sample. First, we compute the Short-Term Fourier Transform (STFT) with a Hamming window of 40ms and a hop size of 10ms. Then, we compute the Discrete Cosine Transform (DCT) with the same set of parameters. Finally, we compute the Non-Stationary Gabor Transform (NSGT) [8] mapped either on a Constant-Q scale of 48 bins per octave and a Mel scale or ERB scale of 400 bins, all from 30 to 11000 Hz. For all transforms, we only keep the magnitude of the distribution to train our models. We perform a corpus-wide normalization to preserve the relative intensities of the samples (normalizing all distributions by the maximal value found across samples). Then, we extract a single temporal frame from the sustained part of the representation (200 ms after the beginning of the sample) to represent a given audio sample. Finally, the dataset is randomly split across notes to obtain a training (90%) and test (10%) set.

Audio reconstruction. To perform audio synthesis, we consider paths inside the latent space, where each point corresponds to a single spectral frame. We sample along a given path and concatenate the spectral frames to obtain the magnitude distribution. Then, we apply the Griffin-Lim algorithm in order to recover the phase distribution and synthesize the corresponding waveform.

4.2. Models

Here, we rely on a simple VAE architecture to show the efficiency of the proposed method. The encoder is defined as a 3-layer feed-forward neural network with Rectified Linear Units (ReLU) activation functions and 2000 units per layer. The last layer maps to a given dimensionality d of the latent space. In our experiments, we analyzed the effect of relying on different latent spaces and empirically selected latent spaces with 64 dimensions. The decoder is defined in a symmetrical way, with the same architecture and units, mapping back to the dimensionality of the input transform. For learning the model, we use a value of $\beta = 2$, which is linearly increased from 0 to its final value during the first 100 epochs (following the *warmup* procedure [17]). In order to train the model, we rely on the ADAM [21] optimizer with an initial learning rate of 0.0001. In a first stage, we train the model without perceptual regularization ($\alpha = 0$) for a total of 5000 epochs. Then, we introduce the perceptual regularization ($\alpha = 0.1$) and train for another 1000 epochs. This allows the model to first focus on the quality of the reconstruction, and then to converge towards a solution with perceptual space properties. We found in our experiments that this two-step procedure is critical to the success of the regularization.

5. RESULTS

5.1. Latent spaces properties

In order to visualize the 64d latent spaces, we apply a simple Principal Component Analysis (PCA) to obtain a 3d representation. Using a PCA ensures that the visualization is a linear transform of the original space. Therefore, this preserves the real distances inside the latent space. Furthermore, this will allow to recover an exploitable representation when we will use this space to generate novel audio content. The results of learning regularized latent spaces for different spectral transforms are displayed in Figure 3.

As we can see, in VAEs without regularization (small space), the relationships between instruments do not match perceptual ratings. Furthermore, the variance of distributions show that the model rather tries to spread the information across the latent space to help the reconstruction. However, the NSGT provides a better unregularized space with different instrumental distributions already well separated. Now, if we compare to the regularized spaces, we can clearly see the effect of the criterion, which provides a larger separation of distribution. This effect and final result is particularly striking for the NSGT (c), which provides the highest correlation to the distances in our combined timbre space (Figure 2). Interestingly, the instrumental distributions might be shuffled around space in order to comply with the reconstruction objective. However, the pairwise distances reflecting perceptual relations are well matched as indicated by the KL divergence. By looking at the test set reconstructions, we can see that enforcing the perceptual topology on the latent spaces do not impact the quality of audio reconstruction for the NSGT, where the reconstruction provides an almost perfectly matching distribution. In the case of the STFT, we can see that the model is impacted by the regularization and mostly match the overall density of the distribution rather than its exact peak information. Finally, it seems that the DCT model diverged in terms of reconstruction, being unable to reconstruct the distributions. However, we can see that the KL fit to timbre distances is better than the STFT, indicating an overfit of the learning towards the regularization criterion. This generative evaluation is quantified and confirmed in the next section.

Method		$\log p(\mathbf{x})$	$\ \mathbf{x} - \tilde{\mathbf{x}}\ ^2$
Unregularized (NSGT)	PCA	-	2.2570
	AE	-1.2008	1.6223
	VAE	-2.3443	0.1593
Regularized (VAE)	STFT	-1.9237	0.2412
	DCT	4.3415	2.2629
	NSGT-CQT	-2.8723	0.1610
	NSGT-MEL	-2.9184	0.1602
	NSGT-ERB	-2.9212	0.1511

Table 1: Generative capabilities evaluated by the log likelihood and mean quality of reconstructed representations on the test set.

5.2. Generative capabilities

We quantify the generative capabilities from the latent spaces by computing the log likelihood and mean difference between the original and reconstructed spectral representations on the test set. We compare these results for different transforms and without regularization, which are presented in Table 1.

As we can see, the unregularized VAE trained on the NSGT distribution provides a very good reconstruction capacity, and still generalizes very well. This can be seen in its ability to generate spectral distributions from the test set almost perfectly. Interestingly, regularizing the latent space does not seem to affect the quality of the reconstruction at all. It even seems that the generalization increases with the regularized latent space. This could however be explained by the fact that the regularized models are trained for twice as much epochs based on our two-fold procedure.

It clearly seems that NSGTs provide both better generalization and reconstruction abilities, while the DCT seems to provide only a divergent model. This can be explained by the fact that NSGT frequency axis is organized on a logarithmic scale. Furthermore, their distribution are well spread across this axis, whereas STFT and DCT tends to have most of their informative dimensions in the bottom half of the spectrum. Therefore, NSGTs provide a more informative input. Finally, there only seems to be a marginal difference between the results of different NSGT scales. However, for all remaining experiments, we select the NSGT-ERB as it is more coherent with our perceptual endeavor.

Thanks to the decoder and its generative capabilities, we can now directly synthesize the audio corresponding to any point inside the latent space, but also any paths between two given instruments. This allows us to turn our analytical spaces into audio synthesizers. Furthermore, as shown in Figure 5 (Bottom right), synthesizing audio along these spaces lead to smooth evolution of spectral distributions and perceptually continuous synthesis (as discussed extensively in the next section). In order to perform subjective evaluation of the audio reconstruction, generated samples from the latent space are available on the supporting repository.

5.3. Generalizing perception, audio synthesis of timbre paths

Given that the encoder of our latent space is trained directly on spectral distributions, it is able to analyze samples belonging to new instruments that were not part of the original perceptual studies. Furthermore, as the learning is regularized by perceptual ratings, we could hope that the resulting position would predict the perceptual relationships of this new instrument to the existing in-

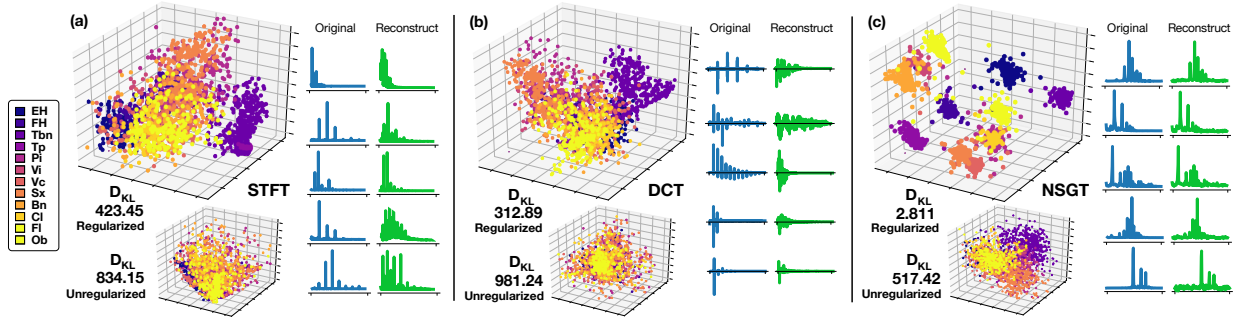


Figure 3: Comparing the regularized VAE latent spaces (large) for the STFT (a), DCT (b) and NSGT (CQT) (c) transforms. For each transform, we plot the corresponding unregularized space (small) and their respective \mathcal{D}_{KL} divergence to the timbre space distances. We plot a set of VAE decoder reconstructions of instrumental spectral frame distributions from the test set directly from the regularized spaces

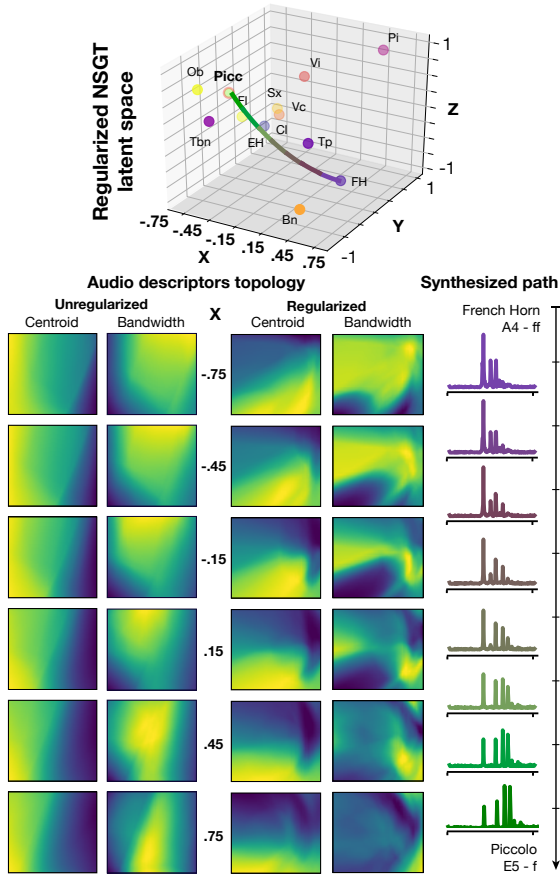


Figure 4: (Top) Projecting new instruments inside the regularized latent space allow to see their perceptual relations to others. (Bottom right) We can generate any path between instruments in the space and synthesize the corresponding perceptually-smooth audio evolution. (Bottom, left) We define 6 equally-spaced projection planes across the x axis and sample points on a 50x50 grid. We reconstruct their audio distribution to compute their spectral *centroid* and *bandwidth*. We compare the resulting descriptor space topology for unregularized (left) and regularized (right) spaces.

struments. This could potentially feed further perceptual studies, to refine timbre understanding. To evaluate this hypothesis, we extracted a set of *Piccolo* audio samples to evaluate their behavior in latent space. We perform the same processing as for the training dataset (Section 4.1) and encode these new samples in the latent space to study the *out-of-domain* generalization capabilities of our model. The results of this analysis are presented in Figure 5 (Top).

Here, we can see that new samples (represented by their centroid for clarity) are encoded in a coherent position in the latent space, as they group with their families, even though they were never presented to the model during learning. However, obtaining a definitive answer on the perceptual inference capabilities of these spaces would require a complete perception experiment, that we leave to future work. Now, as argued previously, one of the key property of the latent spaces is that they provide an invertible non-linear mapping. Therefore, we could thrive on this property to truly understand what are the perceptual relations between instruments based on the behavior of spectral distributions *between* the points in the timbre space. To exhibit this capability, we encode the position in the latent space of a *Piccolo* sample playing an E5-f. Then, based on the position of a *French Horn* playing an A4-ff, we perform an interpolation between these latent points to obtain the path between these two instruments in latent space. We then sample and decode the spectral distributions at 6 equally spaced positions along the path, which are displayed in Figure 5 (Right). As we can see, the resulting audio distributions demonstrate a smooth evolution between timbral structures of both instruments. Furthermore, the resulting interpolation is clearly more complex than a linear change between one structure to the other. Hence, this approach could be used to understand more deeply the timbre relationships between instruments. Also, this provides a model able to perform perceptually-relevant synthesis of novel timbres, while sharing the properties of multiple instruments.

5.4. Topology of audio descriptors

Here, we analyze the topology of signal descriptors across the latent space. As the space is continuous, we do so by sampling uniformly the PCA space and then using the decoder to generate audio samples at a given point. Then, we compute the audio descriptors of this sample. In order to provide a visualization, we select 6 equally-distant planes across the x dimension, at $\{-0.75, -0.45, -0.15, 0.15, 0.45, 0.75\}$, which define an uniform 50x50

grid between $[-1, 1]$ on other dimensions. We compare the results between unregularized or regularized NSGT latent spaces in Figure 5 (Bottom left) for the *spectral centroid* and *spectral bandwidth*. Animations of continuous traversals of the latent space are available on the supporting repository. As we can see, the audio descriptors behave following overall non-linear patterns for both unregularized and regularized latent spaces. However, they still exhibit locally smooth properties. This shows that our model is able to organize audio variations. In the case of unregularized spaces, the organization of descriptors is spread out in a more even fashion. The addition of perceptual ratings to regularize the learning seems to require that this space is organized with a more complex topology. This could be explained by the fact that, in the unregularized case, the VAE only needs to find a configuration of the distributions that maximizes their reconstruction. Oppositely, the regularization requires that instrumental distances follow the perceptual dissimilarity ratings, prompting the need for a more complex relationship between descriptors. This might underline the fact that linear correlations between MDS dimensions and audio descriptors is insufficient to truly understand the dimensions related to timbre perception. However, the audio descriptors topology overall still provide locally smooth evolutions. Finally, a very interesting observation comes from the topology of the centroid. Indeed, all perceptual studies underline its correlation to timbre perception, which is partly confirmed by our model (by projecting on the y axis). This tends to confirm the perceptual relevance of our regularized latent spaces. However, this also shows that the relation between centroid and timbre might not be linear.

5.5. Descriptor-based synthesis

As shown in the previous section, the audio descriptors are organized in a smooth locally linear way across the space. Furthermore, as discussed in Section 5.1, we have seen that the instrumental distributions are grouped across spaces depending on perceptual relations. Based on these two findings, we hypothesize that we can find paths inside these spaces that modify a given audio distribution to follow a target descriptor, while remaining perceptually smooth. Hence, we propose a simple method for perceptually-relevant *descriptor-based path synthesis* presented in Algorithm 1.

Based on the latent space \mathbf{z} (with corresponding encoder q and decoder p) and a given origin spectrum \mathbf{x}_0 , the goal of this algorithm is to find the succession of spectral distributions that match a given target evolution $\mathbf{t} \in \mathbb{R}^N$ for a descriptor d . First, we find the position of the origin distribution in latent space \mathbf{z}_0 and evaluate its descriptor value \mathbf{d}_0 (lines 1-4). Then for each point i , we compute the descriptor values \mathbf{D}_i in the neighborhood of the current latent point (lines 6-10) by decoding their audio distributions. Note that the neighborhood is defined as the set of close latent points, and its size directly defines the complexity of the optimization. Then, we select the neighboring latent point \mathbf{z}_i that provides the evolution of descriptor closest to the target evolution $t[i]$ (lines 11-14). Finally, we obtain the spectral distribution $S[i]$ by decoding the latent position \mathbf{z}_i . The results of applying this algorithm to a given instrumental distribution is presented in Figure 5.

Here, we start from the NSGT distribution of a Clarinet-Bb playing a G#4 in *fortissimo*. We apply our algorithm twice from the same origin point, either on a descending target shape for the spectral centroid (top), or an ascending log shape for the spectral bandwidth (bottom). In both cases, we plot the synthesized NSGT distributions at different points of the optimized path, and

Algorithm 1: Descriptor-based path synthesis

Data: space \mathbf{z} , encoder $q_\phi(\mathbf{z}|\mathbf{x})$, decoder $p_\theta(\mathbf{x}|\mathbf{z})$
Data: origin spectrum \mathbf{x}_0 , target series $\mathbf{t}_{1..N}$, descriptor d
Result: spectral distrib. $S \in \mathbb{R}^{N \times F}$

```

1 // Find origin position in latent space
2  $\mathbf{z}_0 = q_\phi(\mathbf{x}_0)$ 
3 // Evaluate origin descriptor
4  $\mathbf{d}_0 = \text{evaluate}(\mathbf{x}_0, d)$ 
5 for  $i \in [1, N]$  do
6   // Latent 3-d neighborhood of current point
7    $\mathbf{N}_i = \text{neighborhood}(\mathbf{z}_{i-1})$ 
8   // Sample and evaluate descriptors
9    $\mathbf{X}_i = q_\phi(\mathbf{N}_i)$ 
10   $\mathbf{D}_i = \text{evaluate}(\mathbf{X}_i, d)$ 
11  // Compute difference to target
12   $\Delta_i = \|(\mathbf{D}_i - \mathbf{d}_{i-1}) - (t[i] - t[i-1])\|^2$ 
13  // Find next latent point
14   $\mathbf{z}_i = \text{argmin}(\Delta_i)$ 
15  // Decode distribution
16   $S[i] = p_\theta(\mathbf{z}_i)$ 
17 end
```

the neighboring descriptor space. As we can see, the resulting descriptor evolution closely match the input target in both cases. Furthermore, we can see by visual inspection of the spectrum evolution, that the corresponding distributions are indeed sharply modified to match the desired descriptors. Interestingly, the optimization of different target shapes on different descriptors lead to widely different paths in the latent space. However, the overall timbre structure of the original instrument still seems to follow a smooth evolution. Here, we note that the algorithm is quite rudimentary, and could benefit from more global neighborhood information, as witnessed from the slightly erratic local selection of latent points.

6. CONCLUSION

Here, we have shown that regularizing VAEs with perceptual ratings provides timbre spaces that allow for high-level analysis and audio synthesis directly from these spaces. The organization of these perceptually-regularized latent spaces prove the flexibility of these systems, and provides a latent space from which generation of novel audio content is straightforward. These spaces allow to extrapolate perceptual results on new sounds and instruments without the need to collect new measurements. Finally, by analyzing the behavior of audio descriptors across the latent space, we have shown that even though they follow a non-linear evolution, they still exhibit some locally smooth properties. Based on these, we introduced a method for descriptor-based path synthesis that allow to synthesize audio that match a target descriptor shape, while retaining the timbre structure of instruments. Future work on these latent spaces would be to perform perceptual experiments to confirm their perceptual topology.

7. REFERENCES

- [1] Stephen McAdams, Bruno L. Giordano, Patrick Susini, Geoffroy Peeters, and Vincent Rioux, “A meta-analysis of

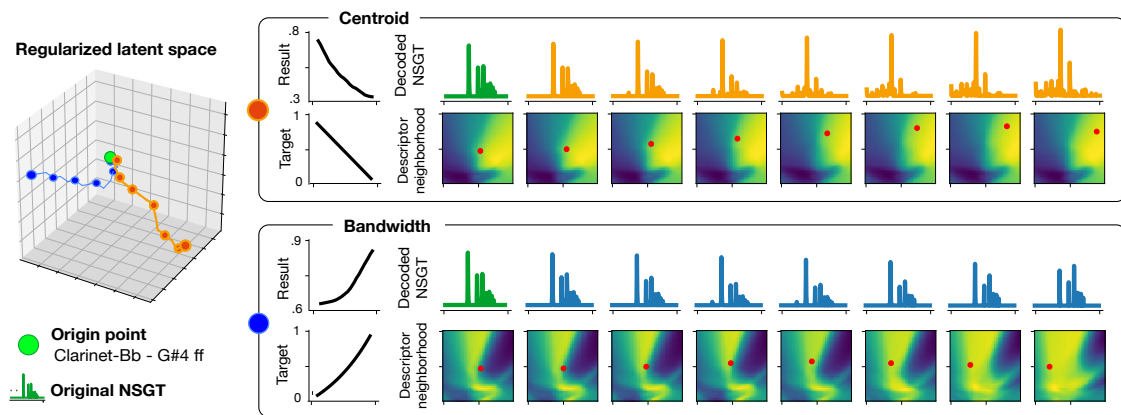


Figure 5: *Descriptor-based synthesis*. Given an origin point in latent space (Clarinet-Bb G#4 ff), we apply our algorithm either on a descending target shape for the spectral centroid (top), or an ascending log shape for the spectral bandwidth (bottom). In both cases, we plot the decoded NSGT distributions and neighboring descriptor space information along the optimized path

- acoustic correlates of timbre dimensions,” *Journal of the Acoustical Society of America*, vol. 120, no. 5, 2006.
- [2] John M Grey and John W Gordon, “Perceptual effects of spectral modifications on musical timbres,” *The Journal of the Acoustical Society of America*, vol. 63, no. 5, 1978.
 - [3] Aaron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
 - [4] Soroush Mehri, Kundan Kumar, Ishaan Gulrajani, Rithesh Kumar, Shubham Jain, Jose Sotelo, Aaron Courville, and Yoshua Bengio, “SAMPLRN: An unconditional end-to-end neural audio generation model,” *ICLR Conference*, 2017.
 - [5] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Douglas Eck, Karen Simonyan, and Mohammad Norouzi, “Neural audio synthesis of musical notes with wavenet autoencoders,” *arXiv preprint:1704.01279*, 2017.
 - [6] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” *ICLR Conference*, 2014.
 - [7] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner, “beta-vae: Learning basic visual concepts with a constrained variational framework,” *ICLR Conference*, 2016.
 - [8] Peter Balazs, Monika Dörfler, Florent Jalliet, Nicki Holighaus, and G Velasco, “Theory, implementation and applications of nonstationary gabor frames,” *Journal of computational and applied mathematics*, vol. 236, no. 6, 2011.
 - [9] Laurens van der Maaten and Geoffrey Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
 - [10] John M Grey, “Multidimensional perceptual scaling of musical timbres,” *the Journal of the Acoustical Society of America*, vol. 61, no. 5, pp. 1270–1277, 1977.
 - [11] Carol L Krumhansl, “Why is musical timbre so hard to understand,” *Structure and perception of electroacoustic sound and music*, vol. 9, pp. 43–53, 1989.
 - [12] Paul Iverson and Carol L Krumhansl, “Isolating the dynamic attributes of musical timbre,” *The Journal of the Acoustical Society of America*, vol. 94, no. 5, pp. 2595–2603, 1993.
 - [13] Stephen McAdams, Suzanne Winsberg, Sophie Donnadieu, Geert De Soete, and Jochen Krimphoff, “Perceptual scaling of synthesized musical timbres: Common dimensions, specificities, and latent subject classes,” *Psychological research*, vol. 58, no. 3, pp. 177–192, 1995.
 - [14] Stephen Lakatos, “A common perceptual space for harmonic and percussive timbres,” *Perception & psychophysics*, vol. 62, no. 7, pp. 1426–1439, 2000.
 - [15] Christopher M Bishop and Tom M Mitchell, “Pattern recognition and machine learning,” 2014.
 - [16] Keith Hastings, “Monte carlo sampling methods using markov chains and their applications,” *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
 - [17] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther, “How to train deep variational autoencoders and probabilistic ladder networks,” *arXiv preprint arXiv:1602.02282*, 2016.
 - [18] Jen-Tzung Kuo and Kuan-Ting Chien, “Variational recurrent neural networks for speech separation,” *INTERSPEECH 2017*.
 - [19] Wei-Ning Hsu, Yu Zhang, and James Glass, “Learning latent representations for speech generation and transformation,” *arXiv preprint arXiv:1704.04222*, 2017.
 - [20] Guillaume Ballet, Riccardo Borghesi, Peter Hoffmann, and Fabien Levy, “Studio online 3.0: An internet “killer application” for remote access to ircam sounds and processing tools,” *Journee Informatique Musicale (JIM)*, 1999.
 - [21] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint:1412.6980*, 2014.