

CROSS-MODAL VARIATIONAL INFERENCE FOR BIJECTIVE SIGNAL-SYMBOL TRANSLATION

Axel Chemla–Romeu-Santos^{1,2}, Stavros Ntalampiras¹, Philippe Esling², Goffredo Haus¹, Gérard Assayag¹

¹ Laboratorio d’Informatica Musicale (LIM)
UNIMI, Milano, Italy
axel.chemla@unimi.it
stavros.ntalampiras@unimi.it
goffredo.haus@unimi.it

² IRCAM - CNRS UMR 9912
Sorbonne Université, Paris, France
esling@ircam.fr assayag@ircam.fr

ABSTRACT

Extraction of symbolic information from signals is an active field of research enabling numerous applications especially in the Musical Information Retrieval domain. This complex task, that is also related to other topics such as pitch extraction or instrument recognition, is a demanding subject that gave birth to numerous approaches, mostly based on advanced signal processing-based algorithms. However, these techniques are often non-generic, allowing the extraction of definite physical properties of the signal (pitch, octave), but not allowing arbitrary vocabularies or more general annotations. On top of that, these techniques are one-sided, meaning that they can extract symbolic data from an audio signal, but cannot perform the reverse process and make symbol-to-signal generation. In this paper, we propose a bijective approach for signal/symbol translation by turning this problem into a density estimation task over signal and symbolic domains, considered both as related random variables. We estimate this joint distribution with two different variational auto-encoders, one for each domain, whose inner representations are forced to match with an additive constraint, allowing both models to learn and generate separately while allowing signal-to-symbol and symbol-to-signal inference. In this article, we test our models on pitch, octave and dynamics symbols, which comprise a fundamental step towards music transcription and label-constrained audio generation. In addition to its versatility, this system is rather light during training and generation while allowing several interesting creative uses that we outline at the end of the article.

1. INTRODUCTION

Music Information Retrieval (MIR) is a growing domain of audio processing that aims to extract information (labels, symbolic or temporal features) from audio signals [1, 2]. This field embeds both musical and scientific challenges paving the way to a large variety of tasks. Such abundant industrial and creative applications [3] have attracted the interest of a large number of researchers with plentiful results. Among the diverse sub-tasks included in MIR, *music transcription* comprises an active research field [4, 5] which is not only interesting by itself but finds generic applicability as a sub-task for other MIR objectives (cover recognition, key detection, symbolic analysis). Music transcription can be described

Copyright: © 2019 Axel Chemla–Romeu-Santos^{1,2}, Stavros Ntalampiras¹, Philippe Esling², Goffredo Haus¹, Gérard Assayag¹ et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 Unported License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

as associating symbols to audio signals composed of one or more musical instruments. Thus, this field embeds pitch and multi-pitch estimation tasks but also other musical dimensions, such as dynamics. Currently, most pitch estimation techniques are based on fundamental frequency detection [6]. However, such approaches may prove insufficient in multi-pitch contexts, where the need for more sophisticated approaches appears crucial.

In parallel, the recent rise of *generative systems* provided interesting alternatives to supervised machine learning approaches focusing on classification [7]. These unsupervised learning models aim to discover the inner structure of a dataset based on a reconstruction task. Such methods are usually defined as probabilistic density estimation approaches, Bayesian inference and auto-encoding structures. Among those, the *Variational Auto-Encoders* (VAE) provides a powerful framework, which explicitly targets the construction of a *latent space* [8]. Such spaces are high-level representations with the ability to reveal interesting properties about the inner structure of different types of data [8][9], and also more recently in audio [10]. Such learning procedures can be mixed with supervised learning to perform label extraction and conditional generation, showing the flexibility and the efficiency of this approach. Last but not least, latent spaces can also be explicitly shared by several systems acting on different data domains, providing an elegant way of performing domain-to-domain translation or multi-modal learning [11].

In this article, we propose a generative modeling approach to musical transcription by formulating it as a density estimation problem. Our approach allows to directly model pairs (x, y) , where x represents the spectral features and y represents the corresponding musical annotations. Following a multi-modal approach inspired by Higgins & al. [12], we train two different VAEs on these separate domains whose latent representations are progressively shared through explicit distribution matching. In addition to providing a Bayesian formulation of musical transcription compatible with arbitrary vocabularies, our method also naturally handles the reverse audio generation process, and thus allows both *signal-to-symbol* and *symbol-to-signal* inference. Furthermore, direct data/symbol generation is also available by latent space exploration, providing an interesting method for creative audio synthesis. Finally, we bind our transcription approach with a novel source-separation approach, based on explicit source decomposition with disjoint decoders. The idea behind our method is to use the knowledge previously acquired on individual instruments in order to ease their recognition in the mixture signal. A novel form of inference network is trained on the product space of the decoders latent space, with additional latent dimensions that per-

forms Bayesian inference directly over mixture coefficients.

2. STATE OF THE ART

Here, we provide a brief state-of-the-art of the most common approaches for musical transcription. Then, we introduce *variational auto-encoders* and detail their use for cross-modal inference and generation.

2.1. Automatic music transcription

Automatic music transcription (AMT) aims at closing the gap between acoustic music signals and their corresponding musical notation. The main problem in AMT is detecting multiple and possibly overlapping in time pitches. Classical approaches for pitch and multi-pitch extraction are mostly based on spectral or spectral analysis using fundamental harmonics localization [13], such as the Yin algorithm [6]. As these methods were originally conceived for monophonic signals, their extension to multi-pitch estimation contexts often implies recursive processes (multi-fundamental recognition, harmonic subtraction) that reduce their efficiency. In parallel, other methods relying on spectrogram factorization have been proposed. These are based on the decomposition of the spectrogram into a linear combination of non-negative factors. These include Non-negative Matrix Factorisation (NMF) [14] or probabilistic latent component analysis (PLCA) [15]. However, spectrogram factorization methods usually fail to identify a global optima, a limitation which led many researchers to hypothesize the need for supplementary external knowledge to attain more accurate decompositions [16, 17].

Recently, deep learning approaches have been proposed to address the multi-pitch detection problem. For instance, piano transcription task has been tackled via a variety of neural networks in [18, 19, 20, 21]. Interestingly, the MusicNet dataset [22] includes multi-instrument music conveniently structured to address polyphonic music transcription. Finally, a method based on convolutional neural networks is presented in [23], which aims at learning meaningful representations allowing accurate pitch approximation in polyphonic audio recordings.

2.2. Generative models and variational auto-encoders

2.2.1. Variational inference

Generative models define a class of unsupervised machine learning approaches aiming to recover the probability density $p(\mathbf{x})$ underlying a given dataset. This density is usually conditioned on another set of random variables \mathbf{z} , called *latent variables*. This set acts as a higher-level representation that controls the generation in the data domain. Formally, generative models can be described as modeling the joint probability $p(\mathbf{z}, \mathbf{x}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$, where $p(\mathbf{z})$ acts as a Bayesian *prior* over the latent variables. The *generative process* takes a latent position \mathbf{z} to produce the corresponding probability density $p(\mathbf{x}|\mathbf{z})$ in the data domain. Conversely, we also want to estimate the *posterior* distribution $p(\mathbf{z}|\mathbf{x})$, that gives the latent distribution corresponding to a data sample \mathbf{x} . Retrieving this posterior distribution from a given generative process is called *Bayesian inference*, and is known to be a very robust inference framework. Unfortunately, this inference is generally intractable for complex distributions or requires limiting assumptions on both generative and inference processes. *Variational inference* (VI) is a framework that overcomes this intractability by turning Bayesian

inference to an optimization problem [24]. To do so, variational inference posits a parametric distribution $q(\mathbf{z}|\mathbf{x})$ that can be freely designed, and optimizes this distribution to approximate the real posterior $p(\mathbf{z}|\mathbf{x})$. This optimization is performed thanks to the following bound

$$\log p(\mathbf{x}) \geq \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [p(\mathbf{x}|\mathbf{z})] + D_{KL}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})] = \mathcal{L}_{\text{ELBO}}(q) \quad (1)$$

where D_{KL} denotes the Kullback-Leibler divergence. We can see that maximizing the right term of this inequality inherently optimizes the evidence $p(\mathbf{x})$ of our model. This bound, called the *Evidence Lower-BOund* (ELBO), can be interpreted as the sum of a likelihood term $p(\mathbf{x}|\mathbf{z})$ and of a divergence term that enforces the approximated posterior $q(\mathbf{z}|\mathbf{x})$ to match the prior $p(\mathbf{z})$. This variational formulation is less restrictive than direct Bayesian inference, as it only requires the tractability of these two terms. Thus, we are able to model complex dependencies between \mathbf{x} and \mathbf{z} for both $p_{\theta}(\mathbf{x}|\mathbf{z})$ and $q_{\phi}(\mathbf{z}|\mathbf{x})$ while retaining the benefits of a Bayesian formulation [25].

2.2.2. Variational auto-encoder and cross-modal learning

To define the approximate distribution, we can model both generative and inference models as normal distributions

$$\begin{aligned} q(\mathbf{z}|\mathbf{x}) &= \mathcal{N}(\boldsymbol{\mu}_q(\mathbf{x}), \boldsymbol{\sigma}_q^2(\mathbf{x})) \\ p(\mathbf{x}|\mathbf{z}) &= \mathcal{N}(\boldsymbol{\mu}_p(\mathbf{z}), \boldsymbol{\sigma}_p^2(\mathbf{z})) \end{aligned}$$

such that parameters $(\boldsymbol{\mu}_q, \boldsymbol{\sigma}_q^2)$ and $(\boldsymbol{\mu}_p, \boldsymbol{\sigma}_p^2)$ are respectively obtained by deterministic functions $f_{\theta}(\mathbf{x}; \theta)$ and $g_{\phi}(\mathbf{z}; \phi)$. When these functions are parametrized as neural networks, we obtain the original *Variational Auto-Encoder* (VAE) formulation proposed by Kingma & al. [8]. The prior is usually defined as an isotropic normal distribution $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, which acts as a regularizer to enforce the independence of latent dimensions. Similar to auto-encoding architectures, $f_{\theta}(\mathbf{x}; \theta)$ and $g_{\phi}(\mathbf{z}; \phi)$ are respectively called the *encoder* and the *decoder* of the system. These functions are jointly trained until convergence on parameters $\{\theta, \phi\}$ with a back-propagation algorithm. Despite the apparent simplicity of its formulation, this system allows very expressive encoding and generative processes while providing a highly structured latent space, whose smoothness is provided by the D_{KL} reconstruction term.

3. CROSS-MODAL VAE FOR MUSIC TRANSCRIPTION

3.1. Signal/symbol transfer through shared latent spaces

In this paper, we propose to reformulate the audio transcription problem as the estimation of a joint probability density $p(\mathbf{x}, \mathbf{y})$, where \mathbf{x} represents the spectral information of the analyzed audio signal and \mathbf{y} represents the corresponding set of symbolic information. Previous works showed the efficiency of VAEs for audio processing when used on spectral frames, in terms of both representational and generative abilities [10, 26]. However, we intend here to estimate not only the probability density $p(\mathbf{x})$, but also the joint probability density $p(\mathbf{x}, \mathbf{y})$. Considering \mathbf{y} as label information, some approaches proposed to include an additional discriminator on the latent space, that is jointly trained during the learning process [27]. Here, we take inspiration from the SCAN approach proposed by Higgins & al., that trains a mirrored VAE on symbolic

data whose latent representation is constrained to match the latent space obtained from the signal VAE [12]. Hence, modelling our symbolic information as binary vectors $\mathbf{y} = [y_1, \dots, y_L]$, we can train this VAE over the label space

$$q(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}_q(\mathbf{y}), \boldsymbol{\sigma}_q^2(\mathbf{y}))$$

$$p(\mathbf{y}|\mathbf{z}) = \prod_{i=1}^L \mathcal{B}(\mu_{p,i}(\mathbf{z})) \text{ or } \prod_{i=1}^L \text{Cat}(\boldsymbol{\mu}_{p,i}(\mathbf{z}))$$

where $\mathcal{B}(\mu_{p,i}(\mathbf{z}))$ denotes a Bernoulli distribution of mean $\mu_{p,i}$ for binary symbols, and $\text{Cat}(\boldsymbol{\mu}_{p,i}(\mathbf{z}))$ denotes a Categorical distribution of classwise probabilities $\boldsymbol{\mu}_{p,i}$ in the case of multi-label symbols. We enforce its latent representation to fit the one obtained with the signal VAE by adding a term to the ELBO

$$\mathcal{L}_{\text{scan}}(q) = \mathcal{L}(q) + D_{KL}[q(\mathbf{z}|\mathbf{x})||q(\mathbf{z}|\mathbf{y})] \quad (2)$$

such that the latent distributions provided by the two inference processes match for a given pair (\mathbf{x}, \mathbf{y}) . The ordering of terms in the Kullback-Leibler divergence is chosen such that the distribution $q(\mathbf{z}|\mathbf{y})$ is forced to cover the whole mass of $q(\mathbf{z}|\mathbf{x})$. Hence, the correct label for a given \mathbf{x} is encouraged even for low-probability areas of $q(\mathbf{z}|\mathbf{x})$. Both VAEs are jointly trained, so that the latent representation obtained is a compromise between both auto-encoder performances. It should be noted that, as both signal and symbolic VAEs are independent, we are still able to perform semi-supervised learning for incomplete pairs (\mathbf{x}, \mathbf{y}) by training only one of the two auto-encoders.

3.2. Bidirectional signal-to-symbol mappings

Our approach extends the multi-pitch detection problem on several aspects. First, our model is inherently bi-directional as we can recover symbolic inference with the process

$$p(\mathbf{y}|\mathbf{x}) = p(\mathbf{y}|\mathbf{z})q(\mathbf{z}|\mathbf{x})$$

This can be understood as a Bayesian formulation of audio semantic labeling. Hence, multi-pitch transcription is simply a special case of our formulation, where \mathbf{y} is defined as being solely the pitch information. Furthermore, we can also naturally handle *signal generation from symbolic constraints*, by taking the reverse process

$$p(\mathbf{x}|\mathbf{y}) = p(\mathbf{x}|\mathbf{z})q(\mathbf{z}|\mathbf{y})$$

such that we can recover the appropriate spectral distribution from the symbolic data, as depicted in Fig. 1. Another interesting property of our method is its applicability to arbitrary symbols. In this paper, we model symbolic information \mathbf{y} as a triplet [pitch class, octave, dynamics], where we add dynamics estimation to the pitch estimation task. Thus, we have $p(\mathbf{y}|\mathbf{x}) = p(\mathbf{y}^p|\mathbf{z})p(\mathbf{y}^o|\mathbf{z})p(\mathbf{y}^d|\mathbf{z})$, where each $p(\mathbf{y}^i|\mathbf{z})$ is defined as a categorical distribution. We use this property to extend this method to multi-pitch applications, where \mathbf{x} is a mixture signal with M different sources. Hence, we formulate the symbolic information as a product $p(\mathbf{y}|\mathbf{z}) = p(\mathbf{y}_1|\mathbf{z}) \dots p(\mathbf{y}_M|\mathbf{z})$, where each $p(\mathbf{y}_i|\mathbf{z})$ follows the previous specification. In addition to performing multi-pitch estimation, it also specifies the corresponding instrument if a given symbolic ordering is held during training. Finally, our formulation can be extended to polyphonic instruments in a straightforward manner. In this case, we simply replace the above conditioning by $p(\mathbf{y}|\mathbf{z}) = p(\mathbf{y}^p|\mathbf{x})p(\mathbf{y}^o|\mathbf{z})$, where we define $p(\mathbf{y}^p|\mathbf{z})$ to be a Bernoulli distribution over a one-hot pitch vector.

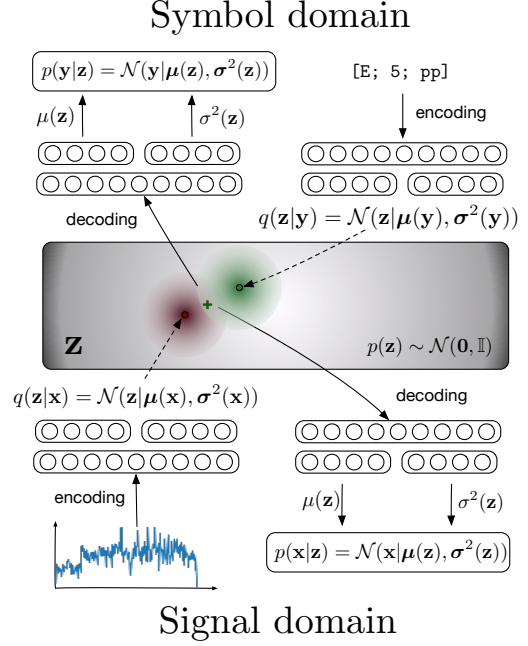


Figure 1: Multi-modal variational auto-encoding process. Our model is based on two separate variational auto-encoder, one in the signal domain and one in the symbolic domain, sharing a common latent space.

4. EXPERIMENTS

4.1. Datasets

To evaluate our approach we use the Studio One Line (SOL) [28], a database that contains solo instrument recordings for every note across their tessitura. Each note is recording over a range of different dynamics (*ff, mf, pp*). Here, we selected five instruments: violin, alto-sax, flute, C-trumpet and piano, for a total amount of 800 files. First, audio files are all resampled to a sample rate of 22050Hz. Then, we transform the raw audio data to the spectral domain by using a *Non-Stationary Gabor Transform* (NSGT) [29]. Interestingly, this multi-resolution spectral transform allows to define custom frequency scales, while remaining invertible. Here, we use a constant-Q scale with 48 bins per octave. For each model training, we split our dataset with 80% as training and 20% as test sets. As our dataset is composed of monophonic signals, we randomly create instrument signal mixtures during training such that every combination is seen during the training.

4.2. Models

To show the efficiency of our proposal, we rely on VAEs with very simple architectures. Nevertheless, depending on the complexity of the input data, we adjust the dimensionality of both the latent space and hidden layers. For single-instrument models we use 32 dimensions for the latent space, and define both encoding and decoding functions for the signal VAE as 2-layers multi-layer perceptrons (MLP) with 2000 hidden units. For the symbolic auto-encoder, encoding and decoding MLPs have 2 layers and 800 hid-

Table 1: Signal reconstruction and transfer performances

	$-\log p(\mathbf{x} \mathbf{z})$	ISD	$-\log p(\mathbf{x} \mathbf{y})$	ISD
Alto-Sax (Sax)	-694.1	0.093	-416.6	0.177
Violin (Vn)	-671.4	0.104	-551.1	0.151
Trumpet-C (TpC)	-706.9	0.073	276.71	0.35
Flute (Fl)	-706.2	0.076	-379.2	0.147
Piano (Pn)	-813.5	0.044	-361.13	0.112
Sax + Vn	-358.71	0.364	-27.37	0.852
Sax + Vn + Fl	-268.7	0.624	692.4	3.813

den units. For mixtures of two different instruments, the number of hidden units for the signal encoders/decoders are set to 5000. For the mixture of three instruments, hidden layers have 5000 and 1500 units for the signal and the symbolic encoders / decoders respectively. All models are trained using the ADAM optimizer, and we use the warm-up procedure that slowly brings the regularization from 0 to 1 during the first 100 epochs. As recommended by Higgins & al., the additional term presented in (2) is scaled up to a factor 10. The learning rate is first set to 1e-3, and is increasingly reduced as the derivative of the error decreases.

4.3. Evaluation

In addition to performing a standard evaluation on the test set, we also evaluate our model on a separate dataset containing recordings of flute arpeggios, scales and melodies [30] with source audio files and aligned MIDI files. Unfortunately, this dataset does not provide information about symbolic dynamics, so we do not evaluate the dynamics inference on this set. We compare the efficiency of our model with results obtained from a baseline approach. To this end, we rely on an architecture similar to our model, but designed in a supervised way to emphasize the gain provided by our model. This baseline classifier first performs a Principal Component Analysis (PCA) from the signal data to perform dimensionality reduction, mocking the compression between the input data and the latent space. Then, we use a 2-layer MLP with the same amount of hidden units than the corresponding symbolical decoder, to output the desired labels. The whole system is trained on a standard cross-entropy loss. The classifier is trained until convergence with the same optimization strategy.

5. RESULTS

In this section, we present the results of our methods. The source code, audio examples and additional figures and results are available on our support page <https://domkirke.github.io/latent-transcription/>.

5.1. Signal reconstruction and transfer performances

First, we analyze the results obtained on the SOL examples. Signal reconstruction and transfer scores are provided in Table 1, relying on two evaluation metrics. The first metric is the log-likelihood of the original spectrum with respect to the distribution decoded by the model. The second is the Itakura-Saito Divergence (ISD), a metric that reflects the perceptual dissimilarity between the original and reconstructed spectrum [31]. Both scores are presented for signal-to-signal reconstruction (left) and symbol-to-signal inference (right). In addition to these scores, reconstruction examples

Table 2: Symbolic inference reconstruction and classification results (successively pitch, octave and dynamics). Scores without parenthesis are reconstruction scores obtained within the symbolic domain, while scores in parenthesis are obtained when performing transfer from the signal domain

		$-\log p(\mathbf{y} \mathbf{z})$	Success Ratio (%)	loose (%)	Baseline (loose)
Sax	p	-1.0 (-1.0)	100% (100%)	-	94%
	o	-1.0 (-1.0)	100% (100%)	-	97%
	d	-1.0 (-1.0)	100% (100%)	-	46%
Vn	p	-1.0 (-1.0)	100% (100%)	-	89.8%
	o	-1.0 (-1.0)	100% (100%)	-	99.0%
	d	-1.0 (-1.0)	100% (100%)	-	35.3%
TpC	p	-1.0 (-1.0)	99.9% (100%)	-	76.1%
	o	-1.0 (-1.0)	100% (100%)	-	99.8%
	d	-0.998 (-1.0)	99.7% (100%)	-	47.8%
Fl	p	-1.0 (-1.0)	100% (100%)	-	52.4%
	o	-1.0 (-1.0)	100% (100%)	-	81.8%
	d	-1.0 (-1.0)	100% (100%)	-	41.4%
Pn	p	-1.0 (-1.0)	100% (100%)	-	51.6%
	o	-1.0 (-1.0)	100% (100%)	-	63.9%
	d	-1.0 (-0.999)	99.9% (100.0%)	-	40.0%
Sax + Vn	p	-0.534 (-0.871)	54.0% (87.9%)	62.6% (81.6%)	65.3%
	o	-0.782 (-0.980)	84.6% (99.2%)	94.9% (88.7%)	79.1%
	d	-0.712 (-0.939)	74.4% (95.9%)	82.4% (66.3%)	52.0%
Sax + Vn + TpC	p	-0.381 (-0.725)	38.6% (75.0%)	62.6% (84.5%)	56.6%
	o	-0.377 (-0.641)	42.4% (67.8%)	79.3% (88.7%)	62.3%
	d	-0.347 (-0.616)	34.6% (62.4%)	66.9% (69.5%)	41.2%

are depicted in Fig. 2. We can see that performances in both signal reconstruction and transfer decrease with the number of instruments, as the complexity of the incoming signal increases. Both reconstruction and signal-to-transfer scores are almost perfect in the case of solo instruments, providing convincing and high-quality sound samples generation. In the case of mixtures of two or more instruments, reconstruction scores maintain an acceptable performance, but symbol-to-signal transfer scores clearly decrease. This observation correlates with the decrease of performance observed in the symbolic domain, as discussed in the following sub-section.

5.2. Symbolic inference performances

Here, we evaluate the performances of our model in the symbolic domain. We provide in Table 2 four different classification scores, separately for each family of labels: octave, pitch class and dynamics. In the case of multi-instrument mixtures, these losses are averaged over every instrument of the mixture. Every column (except for the baseline) show two scores : the first are the scores obtained symbol-to symbol (reconstruction), and the second within parenthesis are the ones obtained signal-to-symbol (transfer).

The first loss, written $-\log p(\mathbf{y}|\mathbf{z})$, denotes the *likelihood* of the true labels with respect to the distributions decoded by the symbolic part of the VAE. The percent scores located at the right of the likelihood correspond to classification scores, obtained by taking the highest probability of the categorical distribution and obtaining the corresponding ratio of well-classified symbols. The first column, called *success ratio*, denotes the classification score obtained by the symbolical VAE. The second column, called *loose ratio*, is specific in the case of mixed instruments, considering a label to be correct regardless of the instrument (we will come back to the motivation behind this score). Finally, the last column display the scores obtained by our baseline classifier, that does not have symbol-to-symbol scores.

We note that symbolic reconstruction and signal-to-symbol scores are almost perfect in the case of single-source signals, outperforming the equivalent baseline system. We argue that is due to two main aspects of the proposed approach. First, thanks to the re-




Figure 2: Signal & symbolic reconstruction samples in 1-instrument, 2-instruments and 3-instruments settings. The first column shows the original spectral & symbolic contents : for spectra, the blue line represents the final mixture, and the thinner lines the different components of the mixture. For symbols, labels are grouped by family and are symbolized by a peak at the correct label. The second column represents reconstruction results. For spectra, the orange line represents the reconstructed spectra (the original spectra in blue is left for comparison). Regarding the symbolic reconstructions, the corresponding categorical distributions are displayed right to the original one-hot vectors. The third column finally shows the transfer results, where we decode a latent position given by the encoder of the other domain.

construction task, the construction of the latent space is organized to reflect the inner structure of both signal and symbol domains. The latent space can be thus understood as a feature space, carrying higher-level information that allow signal/symbolic coupling to be more efficient. Second, the Bayesian approach matching the latent spaces allows a smoother and more efficient mapping than a deterministic approach, that would just provide pairwise mappings between incoming examples.

In the case of instrument mixtures, scores are decreasing as the complexity of both the spectrum and symbolic distributions increase. While the system still performs convincingly with two instruments mixtures, it struggles with mixtures of more than three instruments. We argue that this is partly due to a combinatorial problem, as can be seen when analyzing the *loose classification ratio* : indeed, it becomes harder for the model to accurately affect a label to the corresponding instrument as the number of different possible sources increase. This can be seen with the *loose ratio* in table 2 : where the classification ratio significantly increase if the label is considered correct regardless the affected instrument. This effect can also be seen figure 2, where some peaks are correct but unfortunately distributed to the wrong instrument. A more subtle strategy to tackle this effect has to be considered ; we leave this to a future work.

5.3. Monophonic flute transcription

Here, we analyze the results obtained with an external dataset of flute recordings, as depicted in Table 3. Performances in symbolic inference is still convincing, showing that our model does not suffer from strong over-fitting. Compared to the results obtained with the reference dataset 1, the reconstruction results ob-

Table 3: Results obtained on an external flute dataset

Likelihood	ISD	Class. Ratio	Baseline
2648 (1057)	1.065 (0.632)	65.4%	63.8%
		81.9%	76.8%

tained here have decreased. This is due to several points : first, we have trained the model solely on the stationary part of each instrument signals, such that the attack and release of the signal are not understood by our model. This anomaly is clearly perceptible when listening to the reconstructions. Second, a more subtle comparison between the reference dataset and this dataset showed important differences in terms of harmonic content. Specifically, a 1-octave lower harmonic is globally present in this dataset, and not in the reference one. This may explain the important decrease in octave classification, and may indicate that an increased amount of various instruments of the same type may be required to enforce the generalization of the model.

6. DISCUSSION AND FUTURE WORKS

6.1. Performance aspects

We think that the efficiency of the proposed approach mainly relies on the hybridization of its learning process, that combines both unsupervised and supervised learning. Indeed, while each encoder learns to extract domain-dependant features in an unsupervised manner, latent spaces are matched by enforcing a supervised coupling of signal/symbol pairs. This process thus intends to learn transferable features, that are then used by each decoder to project them back into their respective data domains. Furthermore, this

process allows the model to train on incomplete data, such that each domain’s encoding / decoding functions can still be trained individually even if some signal / symbol couplings are missing. This means that the training method is scalable to bigger datasets where some symbolic information may be absent, such that incomplete data can yet be used to reinforce the reconstruction abilities of the system.

However, in spite of the strengths of the proposed approach, the actual state of the model suffers some issues, that we aim to tackle in the future. The first main issue is that, while the system performs well in the single-instrument case, its performance weakens with two instruments and clearly fails when applied on more. We think that this falls to several reasons. First, we think this is due to the capacity of the model, as we still use very simple systems even for complex signals like the 3-instruments case. Secondly, the complexity of the problem is such that, as we showed when comparing the loose and non-loose version of the classification ratio, the system struggles to correctly allocate the good label to the good instrument. We think that the incoming signal representation may be not precise enough to alleviate some ambiguities, as for examples in the case of octaves or fifths where instrument identification may be hard to disentangle. Furthermore, the model does not prevent instrument-wise symbolic outputs to focus on the same spectral components, and thus to perform redundant symbol predictions, and thus may also lead to a permutation problem. Finally, another issue with the proposed model is that the temporal evolution is not considered by the system. Including temporal features could bring decisive enhancements : in addition to allow full-sound generation and increased pitch and dynamics inference, it may even be mandatory for applying our model to custom symbolic dictionaries (playing modes, temporal symbols such as trills...) and provide a substantial advantage over more casual pitch-detection methods.

6.2. Creative aspects

Finally, an important aspect of the proposed model is the diversity of creative applications it provides (see figure 3). As generation of both symbolic and signal content is both based on the latent space, one may use it as a continuous control space and meaningfully explore it in either an unsupervised or semi-supervised fashion. Indeed, this space can be explored in a fully unsupervised manner by direct interaction: both signal and symbol information are then generated, such that the user can have a direct symbolic feedback on the data he is generating. Alternately, it can also be used in a semi-supervised fashion, constraining the navigation to the distribution inferred by a given symbol or a given sound. For example, in our case, we can directly generate a note with given pitch, octave and dynamics by inferring a distribution with the symbolic encoder, and then navigate inside it to access the diversity of signals retained under the corresponding label information. This allows us, translating first MIDI information in pitch/octave/dynamics pairs, and then transferring this symbolic information in the signal domain, to generate audio content from a MIDI file. We list below various use cases that can be carried out by our model:

- **sequence generation:** we can use a sequence of labels to recover the corresponding distribution in the latent space that we can freely sample and/or navigate,

- **spectral morphing:** we can take two latent target distributions, and draw a trajectory that we can sample regularly to obtain a smooth transformation between the two target sounds,
- **free trajectory:** take a totally free trajectory in the latent space,
- **symbol extraction:** we can infer symbolic information from an incoming signal, and still train the corresponding signal encoder/decoder with the incoming data. This could be a particularly interesting feature especially in real-time contexts.

Corresponding examples for each of the above navigation strategies are given at [support webpage](#). Finally, also note that, in our example, the vocabulary is easy to learn, such that retrieving the underlying distribution $p(\mathbf{y})$ of the symbolic data itself is not really useful. Indeed, the different labels are all independent, and are approximately equally distributed in their own domain. However, our system can also learn on much more complex vocabularies where learning the underlying distribution in the symbolic domain itself has an interest, and thus open additional perspectives for its use in creative and/or MIR applications.




Figure 3: Diagram showing in-domain and cross-domain inference. If we give a spectrum to the signal encoder and draw from the corresponding distribution a latent position \mathbf{z} , we can decode it with the signal generation network (in-domain) and obtain the reconstruction. Alternatively, we can decode it with the symbol generation network (cross-domain) to perform *signal-to-symbol* generation, that here is equivalent to transcription. Reversely, if we draw a latent position from the symbol inference network with a given set of labels, we can decode it with the signal decoder to perform *symbol-to-signal* generation, that allows us to "play" a MIDI file with our model.

7. CONCLUSION

In this paper, we proposed a novel formulation for bijective signal/symbol translation, based on the latent space matching of domain-wise variational architectures. We studied the benefits and drawbacks of the proposed system, and concluded that while improvable this model performed well and proposed a very interesting alternative to signal-symbol algorithms, and furthermore provided

additional applications that were not possible in previous models. Indeed, our method is bi-directional, and performs well in both audio-to-symbol and symbol-to-audio prediction. Furthermore, our method is compatible with any kind of arbitrary symbolic information, and is then opened to user-defined vocabularies. Besides, as our model is based on a latent space that can be considered as a continuous control space, it is also opened to diverse creative uses as sequence generation, sound interpolation, or free navigation, whether in a supervised manner or semi-supervised manner, guided with symbolic information. For future work, we plan to solve the symbolic ambiguities that raise in the case of numerous instruments, to incorporate temporal features to allow dynamical features extraction, and to design user interfaces to make our model compatible with artistic practises.

8. REFERENCES

- [1] J Stephen Downie, “Music information retrieval,” *Annual review of information science and technology*, vol. 37, no. 1, pp. 295–340, 2003.
- [2] F. Simonetta, S. Ntalampiras, and F. Avanzini, “Multimodal music information processing and retrieval: Survey and future challenges,” in *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*, Jan 2019, pp. 10–18.
- [3] Michael A Casey, Remco Veltkamp, Masataka Goto, Marc Leman, Christophe Rhodes, and Malcolm Slaney, “Content-based music information retrieval: Current directions and future challenges,” *Proceedings of the IEEE*, vol. 96, no. 4, pp. 668–696, 2008.
- [4] Anssi Klapuri and Manuel Davy, *Signal processing methods for music transcription*, Springer Science & Business Media, 2007.
- [5] Emmanouil Benetos, Simon Dixon, Dimitrios Giannoulis, Holger Kirchhoff, and Anssi Klapuri, “Automatic music transcription: challenges and future directions,” *Journal of Intelligent Information Systems*, vol. 41, no. 3, pp. 407–434, 2013.
- [6] Alain De Cheveigné and Hideki Kawahara, “Yin, a fundamental frequency estimator for speech and music,” *The Journal of the Acoustical Society of America*, vol. 111, no. 4, pp. 1917–1930, 2002.
- [7] Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent, “Generalized denoising auto-encoders as generative models,” in *Advances in Neural Information Processing Systems*, 2013, pp. 899–907.
- [8] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [9] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra, “Stochastic backpropagation and approximate inference in deep generative models,” *arXiv preprint arXiv:1401.4082*, 2014.
- [10] Philippe Esling, Axel Chemla-Romeu-Santos, and Adrien Bitton, “Generative timbre spaces with variational audio synthesis,” *CoRR*, vol. abs/1805.08501, 2018.
- [11] Ming-Yu Liu, Thomas Breuel, and Jan Kautz, “Unsupervised image-to-image translation networks,” in *Advances in Neural Information Processing Systems*, 2017, pp. 700–708.
- [12] Irina Higgins, Nicolas Sonnerat, Loic Matthey, Arka Pal, Christopher P Burgess, Matthew Botvinick, Demis Hassabis, and Alexander Lerchner, “Scan: Learning abstract hierarchical compositional visual concepts,” *arXiv preprint arXiv:1707.03389*, 2017.
- [13] Thomas Drugman and Abeer Alwan, “Joint robust voicing detection and pitch estimation based on residual harmonics,” in *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [14] Daniel D. Lee and H. Sebastian Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, pp. 788–791, oct 1999.
- [15] Madhusudana Shashanka, Bhiksha Raj, and Paris Smaragdis, “Probabilistic latent variable models as nonnegative factorizations,” *Computational Intelligence and Neuroscience*, vol. 2008, pp. 1–8, 2008.
- [16] G. Grindlay and D. P. W. Ellis, “Transcribing multi-instrument polyphonic music with hierarchical eigeninstruments,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 6, pp. 1159–1169, Oct 2011.
- [17] G. J. Mysore and P. Smaragdis, “Relative pitch estimation of multiple instruments,” in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, April 2009, pp. 313–316.
- [18] Rainer Kelz, Matthias Dorfer, Filip Korzeniowski, Sebastian Böck, Andreas Arzt, and Gerhard Widmer, “On the potential of simple framewise approaches to piano transcription.,” *CoRR*, vol. abs/1612.05153, 2016.
- [19] Siddharth Sigtia, Emmanouil Benetos, and Simon Dixon, “An end-to-end neural network for polyphonic piano music transcription,” *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 24, no. 5, pp. 927–939, May 2016.
- [20] Curtis Hawthorne, Erich Elsen, Jialin Song, Adam Roberts, Ian Simon, Colin Raffel, Jesse Engel, Sageev Oore, and Douglas Eck, “Onsets and frames: Dual-objective piano transcription,” *arXiv preprint arXiv:1710.11153*, 2017.
- [21] Curtis Hawthorne, Andriy Stasyuk, Adam Roberts, Ian Simon, Cheng-Zhi Anna Huang, Sander Dieleman, Erich Elsen, Jesse Engel, and Douglas Eck, “Enabling factorized piano music modeling and generation with the maestro dataset,” *arXiv preprint arXiv:1810.12247*, 2018.
- [22] John Thickstun, Zaid Harchaoui, and Sham M. Kakade, “Learning features of music from scratch,” in *International Conference on Learning Representations (ICLR)*, 2017.
- [23] Rachel M. Bittner, Brian McFee, Justin Salamon, Peter Li, and Juan Pablo Bello, “Deep salience representations for f0 estimation in polyphonic music,” in *ISMIR*, 2017.
- [24] Tommi S Jaakkola and Michael I Jordan, “Bayesian parameter estimation via variational methods,” *Statistics and Computing*, vol. 10, no. 1, pp. 25–37, 2000.
- [25] Christopher M Bishop, *Pattern Recognition and Machine Learning*, Springer edition, 2006.
- [26] Adrien Bitton, Philippe Esling, and Axel Chemla-Romeu-Santos, “Modulated variational auto-encoders for many-to-many musical timbre transfer,” *arXiv preprint arXiv:1810.00222*, 2018.

- [27] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling, “Semi-supervised learning with deep generative models,” in *Advances in neural information processing systems*, 2014, pp. 3581–3589.
- [28] Guillaume Ballet, Riccardo Borghesi, Peter Hoffmann, and Fabien Lévy, “Studio online 3.0: An internet" killer application" for remote access to ircam sounds and processing tools,” *Journée d’Informatique Musicale (JIM)*, 1999.
- [29] Gino Angelo Velasco, Nicki Holighaus, Monika Dörfler, and Thomas Grill, “Constructing an invertible constant-q transform with non-stationary gabor frames,” *Proceedings of DAFX11, Paris*, pp. 93–99, 2011.
- [30] Elena Agullà Cantos, “Flute audio labelled database for Automatic Music Transcription,” Sept. 2018.
- [31] C. Févotte, N. Bertin, and J. Durrieu, “Nonnegative matrix factorization with the itakura-saito divergence: With application to music analysis,” *Neural Computation*, vol. 21, no. 3, pp. 793–830, March 2009.
- [32] Emmanouil Benetos, Anssi Klapuri, and Simon Dixon, “Score-informed transcription for automatic piano tutoring,” in *Signal processing conference (eusipco), 2012 proceedings of the 20th european. IEEE*, 2012, pp. 2153–2157.
- [33] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol, “Extracting and composing robust features with denoising autoencoders,” in *Proceedings of the 25th international conference on Machine learning. ACM*, 2008, pp. 1096–1103.
- [34] Estefanía Cano, Derry FitzGerald, Antoine Liutkus, Mark D Plumbley, and Fabian-Robert Stöter, “Musical source separation: An introduction,” *IEEE Signal Processing Magazine*, 2018.
- [35] Aditya Arie Nugraha, Antoine Liutkus, and Emmanuel Vincent, “Multichannel audio source separation with deep neural networks.,” .
- [36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.