



HAL
open science

What does the Canary Say? Low-Dimensional GAN Applied to Birdsong

Silvia Pagliarini, Nathan Trouvain, Arthur Leblois, Xavier Hinaut

► **To cite this version:**

Silvia Pagliarini, Nathan Trouvain, Arthur Leblois, Xavier Hinaut. What does the Canary Say? Low-Dimensional GAN Applied to Birdsong. 2021. hal-03244723v2

HAL Id: hal-03244723

<https://inria.hal.science/hal-03244723v2>

Preprint submitted on 26 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

What does the Canary Say?

Low-Dimensional GAN Applied to Birdsong

Silvia Pagliarini^{1,2,3}, Nathan Trouvain^{1,2,3}, Arthur Leblois^{3*}, and Xavier Hinaut^{1,2,3*}.

1. INRIA Bordeaux Sud-Ouest. 2. LaBRI, Université de Bordeaux, Bordeaux INP, CNRS UMR 5800. 3. IMN, Université de Bordeaux, CNRS UMR 5293, Bordeaux, France.

Abstract—The generation of speech, and more generally complex animal vocalizations, by artificial systems is a difficult problem. Generative Adversarial Networks (GANs) have shown very good abilities for generating images, and more recently sounds. While current GANs have high-dimensional latent spaces, complex vocalizations could in principle be generated through a low-dimensional latent space, easing the visualization and evaluation of latent representations. In this study, we aim to test the ability of a previously developed GAN, called WaveGAN, to reproduce canary syllables while drastically reducing the latent space dimension. We trained WaveGAN on a large dataset of canary syllables (16000 renditions of 16 different syllable types) and varied the latent space dimensions from 1 to 6. The sounds produced by the generator are evaluated using a RNN-based classifier. This quantitative evaluation is paired with a qualitative evaluation of the GAN productions across training epochs and latent dimensions. Altogether, our results show that a 3-dimensional latent space is enough to produce all syllable types in the repertoire with a quality often indistinguishable from real canary vocalizations. Importantly, we show that the 3-dimensional GAN generalizes by interpolating between the various syllable types. We rely on UMAP [1] to qualitatively show the similarities between training and generated data, and between the generated syllables and the interpolations produced. We discuss how our study may provide tools to train simple models of vocal production and/or learning. Indeed, while the RNN-based classifier provides a biologically realistic representation of the auditory network processing vocalizations, the small dimensional GAN may be used for the production of complex vocal repertoires.

Index Terms—Generative Adversarial Networks, Reservoir Computing, sound generation, birdsong, canary, latent space, low-dimensional.

I. INTRODUCTION

The acoustic variety of speech sounds across speakers, rates and contexts makes speech composition difficult to apprehend. Indeed, when different speakers produce the same phonem, the acoustics (i.e. formant frequencies) vary [2], making speech highly variable. Faster speech acoustics differ from slower speech acoustics, and contexts variations can determine a change in the acoustic features [3]. Behavioral and neuroanatomical similarities between songbirds and humans [4][5] make songbirds a particularly tractable model for a *reduced* version of the vocal learning process in humans, enabling to test basic mechanistic hypotheses and to probe the underlying neural substrate [6]. Among the many songbird species, canaries have a large repertoire and their

songs are characterized by a complex syntax with long-time dependencies [7]. These properties make canary songs a reasonable middle ground between human and songbirds with stereotypical songs (like zebra finches) to study vocal learning and production.

The generation and classification of vocalizations are key points to build models of vocal learning (for speech or birdsong). Complex vocalizations (like human words or sentences) are made of sequences of smaller canonical vocal units (such as phonemes). Focusing on the generation of such vocal units allows to concentrate on the faithful reproduction of sounds, while allowing to later combine them to generate complex vocalizations. The generated vocal units must then be classified to be compared to recorded vocal units in order to evaluate the generation process. Once classified, the generated vocal units may undergo (1) a quantitative evaluation, e.g. to compare different generators, and (2) a qualitative evaluation, for instance to identify the structure and continuity of the generated dataset (using dimension reduction methods such as PCA or UMAP [1]). While good classifiers already exist for speech, there is currently no fully automated software for the classification of the vocal units in complex birdsongs like those of canaries.

Generative models produce high-dimensional complex datasets and can realistically reproduce real data such as images or sounds with a diversity of outputs [8], including elements not present in the training dataset [9]. The representation of generated data in an intermediate layer of the generative model called latent space provides information about the structure of the produced dataset. Such a latent space can have various properties (e.g. diversity of outputs, continuous representation) that enables to obtain meaningful representations of the underlying datasets. In the latent space, nearby points have similar properties [9]. By interpolating between two points in the latent space associated with 2 different outputs, a continuous set of intermediate realistic outputs can be generated [10], [11]. Interestingly, Sainburg et al. [12] have shown that birdsong vocalizations can be visualized in low-dimensional representations using UMAP [1]. While current GANs rely on generative networks with high-dimensional latent space (usually 100-dimensional), birdsong vocal units may thus be generated through very low-dimensional latent spaces. A GAN low-dimensional latent space would (1) ease the direct visualization of the space (without the need to use dimension reduction methods), and (2) facilitate the comparison between the structure of the real data space (e.g. spectrograms) and the

* Corresponding authors who co-supervised the study.

structure of the latent space obtained.

An imitative agent (e.g. a child or a bird) aiming at reproducing the sounds it perceives from conspecifics can be modeled as a sensorimotor system. Such an agent has to learn the inverse model that maps perceived sounds to the corresponding motor commands. The imitation process can be modeled in an artificial system through a vocal learning model. Reducing the number of control parameters needed to generate vocalizations is critical in such model to ease the learning: because lower output (control) dimension means less parameters to be learned, and thus less training samples. A generative model producing vocalizations with a low-dimensional latent space would allow to reduce the number of control parameters needed to generate vocalizations. Moreover, low control dimension is useful to analyze the influence of each parameter on the generated vocalization. Even if a greater number of motor dimensions may not be a problem for *motor babbling* exploration strategies (depending on the heterogeneity of the redundancy) [13], learning the inverse kinematics is more difficult when the degree of freedom of the motor space is much higher than the number of controlled dimensions (i.e. the number of possible sensory outputs) [14]. Altogether, a low motor dimension should facilitate the learning when used in sensorimotor models [15]. A generative network producing complex vocalizations with a low-dimensional latent space could therefore also serve as the basis for a future vocal learning model.

Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) are two classes of generative models that enable the construction of a low-dimensional latent space. VAEs are composed of an encoder, that takes a real sample as input and builds a lower-dimensional representation of it (i.e. an element of the latent space), and a decoder (generative model), that takes an element of the latent space as input and generates back the original sample. GANs [8] are an example of generative models which enable to represent high-dimensional distributions in a latent space. The two main components of a GAN are (i) a generator that tries to reproduce a target distribution (e.g. images) given random inputs, and (ii) a discriminator that estimates the probability that the generated data is taken from the training dataset [8]. GANs differentiate from, on the one hand, other generative models like Wavenet [16] because they are not incremental models, and, on the other hand, from VAEs because of the structure of the latent representation. Indeed, the latent representation in VAEs is common between the encoder and the decoder (i.e. it is in the middle of the network). However, in GANs the latent space representation of the generator (encoder) is initially random and independent from the discriminator (decoder) part. Although VAEs are easier to evaluate quantitatively, GANs generate higher quality outputs with respect to VAEs [17]. GANs have already been used to produce complex datasets such as images, music, speech, but on a lesser extent to birdsongs [18]–[20]. Donahue et al. [20] trained their model, called WaveGAN, on audio recordings of speech, drums, piano, and vocalizations of several wild songbird species. However, they trained WaveGAN on all species at once and never on multiple vocalizations from a single bird specie.

Our choice to produce a large set of canary vocalizations with the WaveGAN generator has been determined by (1) the promising but noisy results obtained on a highly variable dataset of birdsongs, and (2) the fact that the training is more stable with respect to other GANs (WaveGAN is a Wasserstein GAN [21]).

In this paper, we trained WaveGAN [20] on an adult canary dataset [22]. The objective of this study is to test the ability of WaveGAN to produce realistic canary syllables for different conditions, varying the latent space dimension and the size of the training dataset. On the one hand, we are interested in finding the minimal latent space dimension to reproduce realistic and diverse vocalizations. On the other hand, we explore the capability of the generator to produce a good output when the network has been trained with datasets of limited size. Indeed, having a lower amount of training data would speed up the computational time to train the generator. Moreover, we also explore the latent space to evaluate (quantitatively and qualitatively) (1) the accuracy and diversity of the produced syllables, (2) the continuity of the latent space (i.e. the ability of the generator to generalize), and (3) the structure of the latent representations.

The paper is organized as follows. Section II-A describes the preprocessing of the data. Section II-B contains a brief introduction to the architecture of WaveGAN. In section II-C, we present the training dataset and detail the different WaveGAN training conditions we used. Section II-E introduces the metrics used to evaluate the performance of the generator. The tools used to evaluate the model are described in Sections II-D and II-F. Section III-A applies quantitative and qualitative analyses on the training dataset. Section III-B shows the analyses we carried out to evaluate the ability of the generator to reproduce the whole repertoire. Section III-C shows the organization of the latent space. Section III-D and Section III-E show the quantitative analyses presented in Section III-B applied to compare the generator performances when using different conditions on the latent space dimension or on the training dataset size. Section IV summarizes the results obtained, and details the advantages and the limitations of our approach. Moreover, we discuss how such a generator model could be used in future work within the framework of vocal learning models.

All the details of the implementation are available at <https://github.com/spagliarini/low-dimensional-canary-GAN>.

II. METHODS

A. Data pre-processing

Canaries sing sequences of syllables organized in phrases: a phrase is a short segment of the song where the same syllable is repeated many times [7]. The initial dataset was composed of a repertoire of 27 classes of syllables organized in labeled phrases (i.e. each phrase was already assigned to a specific class), manually sorted from a set of recordings at 44100Hz from an adult canary [22]. Among the 27 classes, we focused on 16 classes that had the highest number of samples labeled as: *A, B1, B2, C, D, E, H, J1, J2, L, M, N, O, Q, R, V*. For each recorded phrase, we performed the following steps:

- 1) We downsampled the signal to a rate of 16000 Hz.
- 2) We reduced low-frequency noise using a high-pass filter of order 5 with a cutting frequency of 700 Hz.
- 3) We selected three parameters for syllable detection: the amplitude threshold, the minimal duration of the syllable, and the duration of the gap (i.e., the silence between two consecutive syllables).

Figure 15 in the supplementary material shows a summary of the selection procedure for syllable *J2*. The automatic selection based on an amplitude threshold, the duration of the syllable and the duration of the gap, could fail to select correctly the syllable. We performed a visual inspection after selection to reject misclassified and miscut elements. The resulting dataset contains 72, 155 syllables from the 16 classes. In supplementary material, Figures 16 and 17 show 100 examples of samples that have been selected from each class after the pre-processing.

B. WaveGAN

WaveGAN architecture is based on the architecture of DCGAN [10], initially used for image synthesis. DCGAN [10] generator is a CNN where transposed convolution is used to upsample low-resolution feature maps into a high-resolution image. As for DCGAN, the generator and the discriminator of WaveGAN are CNNs. To process audio signals, a larger receptive field is introduced in WaveGAN. This is similar to what Van den Oord et al. [23] did in WaveNet, where dilated convolutions has been used to increase the effective receptive field of the model.

Donahue et al. [20] used WGAN-GP[24] strategy during training. The role of the discriminator has changed: it is not making a direct choice to assess whether a sample is real or fake. Instead, the discriminator acts as a critic and allows it to be trained until it reaches an optimum. An advantage of the critic is that it can't saturate, whereas the classic discriminator could learn too quickly, becoming unreliable [21]. The loss (i.e., the objective function that is optimized during training) is based on the computation of Wasserstein distance, which gives stability to the GAN [19]. Finally, the introduction of a gradient penalty term in the loss function (driven by the hyperparameter λ), enables a faster training and less parameter tuning [24]. Please refer to the supplementary material (see Figure 18) for further details about the architecture of WaveGAN.

C. Experimental setup

We selected a balanced training dataset in order to consider the same number of samples per syllable type. Among all the selected syllables, we used a subset of 16k syllables: 1k syllables per class. From now on, we will refer to this balanced dataset as the training dataset. Each syllable has been padded with silence to obtain recordings having a length of exactly 1s, to create a dataset resembling the speech dataset that was originally used to train WaveGAN [20]. We used the balanced training dataset to train the classifier (described in Section II-D) and the network (described in Section II-B).

We used the original WaveGAN [20] setup to train the network. We kept the original network architecture with gradient penalty option, with $\lambda = 10$ and Adam optimizer in the training phase. We used a batch size of 64 samples and we trained the discriminator 5 times more than the generator.

We tested various conditions for the latent space dimension (indicated as ld), varying it from $ld = 6$ to 1. In the following, we will refer to one of these conditions using, for instance, 6-dimensional WaveGAN to refer to a GAN trained with a 6-dimensional latent space. We also varied the size of the training dataset. We first trained WaveGAN with the complete dataset, then reduced it by a factor of 2 ($8k$ syllables) and finally by a factor of 4 ($4k$ syllables).

We first trained 5 instances per latent space condition (keeping the dataset size fixed) and dataset size condition (keeping the latent space dimension fixed), to observe the performance of the generator. In Figure 10 and Figure 14, we selected the best instance for each parameter set (no or late overtraining). We compared the results to choose the optimal combination of latent space dimension and dataset size. We then trained 10 instances of WaveGAN with the training dataset introduced at the beginning of this Section. For all the instances, we trained the network until epoch ~ 1000 , saving the model every ~ 15 epochs. After training, we used the generator to produce new syllables at every saved epoch. We evaluated the generated data using both a quantitative and a qualitative measure, as described in Sections II-E1 and II-E2.

D. Classifier

To characterize the original dataset as well as the production of our WaveGAN, we initially trained classifiers to recognize the syllable classes present in the original dataset. The two classifiers (*classifier-REAL* and *classifier-EXT*) used during the evaluation phase (described in Section II-E) are *Echo State Networks* (ESNs) [25], a type of artificial recurrent neural network (RNN). ESNs are part of the Reservoir Computing paradigm) [26] that use random RNNs and train only the output layer. ESNs are often considered as temporal Support Vector Machines (SVMs): they work in a similar way by embedding input data into a high-dimensional space using non-linear transformations. However, unlike SVMs, ESNs are designed to manipulate sequential data and are relevant candidates for sound classification [27], [28].

The classifiers were fed with Mel-Frequency Cepstral Coefficients (MFCCs) representations of the syllables, a low dimensional spectral representation of sound. We extracted 20 MFCCs features per time step, one time step being defined as the result of a spectral analysis window of 64ms, using a Hanning window (often called *window width*; param. *win_length* in *librosa* library) applied to the 16kHz audio signal, with a 32 ms jump between each frame (often called *frame stride*; param. *hop_length*). Because the GAN generated syllables tend to have higher amplitude than the real ones, we used only the first and the second derivative of the extracted MFCCs to remove any influence of the signal amplitude in the representations. Otherwise, the amplitude difference would bias the classifiers decisions, as it would be artificially easier

to separate real samples from generated samples only by comparing the average power of the signals. First and second derivatives of the MFCC signal are also known to be good representations of vocal signals, as they give relevant clues on the temporal dynamics of the vocalizations.

The ESNs trained on the classification tasks [27] are built using *ReservoirPy* Python library [29], [30]¹ and are described by the following equations:

$$\begin{cases} x(n) = (1 - \alpha)x(n - 1) \\ \quad + \alpha \tanh(W_{in}u(n) + Wx(n - 1)) \\ y(n) = W_{out}x(n) \end{cases} \quad (1)$$

where $u(n)$, $x(n)$ and $y(n)$ are respectively the input features, the internal state of the network and the output vector at time step n . W_{in} stores the connection weights between the inputs and the units of the network. These weights are sampled from a discrete bi-modal distribution, i.e. are randomly chosen from the set $\{1, -1\}$. The proportion of non-zero connection weights is fixed to 10%. W_{in} is defined in $\mathbb{R}^{N \times I}$, where N is the number of recurrent units and I is the dimension of the input. In our case $I = 41$, with input features being 20 derivatives of MFCCs, 20 second derivatives of MFCCs, and a constant bias equal to 1. Each set of features is scaled by multiplying the corresponding connection weights in W_{in} by a constant. The first derivatives and the bias are scaled by a factor 1, and the second derivatives are scaled by a factor 0.7. W stores the connection weights of the recurrent units. These weights are sampled from a standard normal distribution, with a proportion of non-zero connections between neuronal units fixed to 10%. W is defined in $\mathbb{R}^{N \times N}$, with N equal to 1000. W is scaled using a factor equal to a fixed *spectral radius* of 0.5 divided by the largest absolute eigenvalue of W . The α parameter, called *leaking rate*, is set to 0.05. It controls the time constant of the ESNs and allows information from past internal states to be fed to future internal states. All connection weights defined in W_{in} and W are fixed during the training phase of the ESNs, as opposed to machine learning algorithms using gradient descent. Only the $W_{out} \in \mathbb{R}^{N \times V}$ readout matrix is learned during training, where V is the output dimension corresponding to the vocabulary size used to classify the input features, i.e. $V = 16$ for *classifier-REAL* and $V = 21$ for *classifier-EXT*. The readout weights are learned using a linear regression between all the internal states x generated from the inputs and all the expected values of the output y . A $L2$ regularization coefficient of value 10^{-8} is applied during the linear regression: thus this corresponds to a ridge regression.

The classifiers then output a vector $\hat{y}(n)$ of dimension V representing its activation for each time step n and for each category of syllable in the vocabulary. Then, these output activities are summed up over the whole sequence of MFCC frames. A softmax operation finally provides a probability distribution representing the chance for the syllable to belong to one of the classes of the vocabulary. A further analysis of the robustness of the classifiers can be found in Appendix III-B.

E. Evaluation

We evaluated the performance of the generator across epochs and across training conditions. As explained in Section II-C, after training, we used the generator to produce new syllables every 15 epochs and we used the classifier described in Section II-D to identify the generated syllables as elements of one class of the vocabulary. The vocabulary is the set of classes known to the classifier. The vocabulary includes all syllable types and may include additional types of generations (see below). First, we used the balanced training dataset to train a classifier able to provide the probability of each sample belonging to each class of a vocabulary composed by the 16 classes of the repertoire. We refer to this model as *classifier-REAL*. Then, we trained a classifier able to provide the probability of each syllable belonging to 21 classes: the 16 classes of the repertoire, a white noise (*WN*) class, an overtraining (*OT*) class, and three *EARLY* classes, respectively obtained from epochs 15, 30 and 45. To train the classifier to recognize the alternative unknown classes, in addition to the usual training dataset, we used three additional sets of generated samples. In summary, to train the classifier we used:

- *EARLY 15*: 1k samples of early generations, obtained after ~ 15 epochs using two different instances of a 3-dimensional WaveGAN (500 samples per instance);
- *EARLY 30*: 1k samples of early generations, obtained after ~ 30 epochs using two different instances of a 3-dimensional WaveGAN (500 samples per instance);
- *EARLY 45*: samples of early generations, obtained after ~ 45 epochs using two different instances of a 3-dimensional WaveGAN (500 samples per instance);
- *OT*: 1k samples obtained when two instances of a 3-dimensional WaveGAN reach overtraining (500 samples per instance);
- *WN*: 1k samples of artificial white noise.

The two different instances used to define the classes above are instances *Ex 0* and *Ex 1* (where *Ex* stands for *example*) in Figure 34. We will refer to this classifier as *classifier-EXT* (where *EXT* stands for extended).

For simplicity, we will call the set of unknown classes X . We define a class $x \in X$ as a class representing either white noise, or samples containing a lot of noise and, in general, resembling early productions of the generator (i.e., belonging to *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*).

1) **Quantitative evaluation**: For GANs, quantitative evaluations allow to probe if the model is able to reproduce a wide enough variety of samples [31], and provide a preliminary measure of whether or not the generated data resembles the training data. To monitor the generator’s performance across time we quantified how many different classes it produces and how many syllables are produced in each class. Thereby, we can compare the performance and stability of WaveGAN across different instances of training. We also computed the Inception Score (IS) at several epochs of the training, observed its evolution and compared it with the IS obtained from the training dataset. The IS provides a method of objective evaluation for GANs performance [20][24][32]. IS has been proposed by Salimans et al. [18] as a quantitative measure to

¹<https://github.com/reservoirpy/reservoirpy>

evaluate GANs: a pre-trained deep learning neural network model for image classification provides the probability of each image belonging to each class. This information is then summarized in the IS, which is defined as follows:

$$IS = \exp(\mathbb{E}(KL(p(y|x)||p(y)))), \quad (2)$$

where KL stands for Kullback-Leibler divergence. Given a problem with N classes, $IS \in [1, N]$. IS provides both a measure of the quality of generated samples and of the entropy of the generations, indicating the generator ability to produce a wide set of new data [18].

2) **Qualitative evaluation:** First, we based our qualitative analysis on spectrogram analysis. Then, we explored the latent space to study its structure and the continuity of the generations. On the one hand, we computed the mean spectrogram of the generated data, and we compared it with the mean spectrogram of the dataset and with the repertoire. To obtain the mean spectrogram, for each class, we first aligned the syllables’ envelope, then we computed the mean of the spectrograms of all the syllables belonging to that class. We observed the spatial organization of the data using Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) [1]. We applied UMAP to the spectrograms of the samples. Further details about the algorithm can be found in Section II-F. On the other hand, we compared the spectrogram space with the latent space to observe the organization of the latent vectors with respect to the syllables they produce. We generated syllables for each small variation of the latent vector.

- **One component variation.** We selected a random latent vector $z \sim R^3([-1, 1])$ to generate a baseline syllable using a 3-dimensional WaveGAN after training. Then, we moved one by one the components of the vector by a variation step equal to $v_{step} = 0.05$. To explore critical points (i.e., where a bigger variation arises a sudden non-smooth change between two syllables), we moved one by one the components of the vector by a variation step equal to $v_{step} = 0.01$ and $v_{step} = 0.001$. We compared the generations obtained by each variation (one per component of the latent vector) with a set of $16k$ generated data from the same epoch.
- **Interpolation between two syllables.** We first selected 2 syllables s_1 and s_2 and their correspondent latent vectors z_1 and z_2 , where $z_1, z_2 \in R^3([-1, 1])$. Then, we moved in the latent space from z_1 to z_2 using a variable step depending on the distance between the components of z_1 and z_2 . That is, the step applied to each component i is:

$$step[i] = \frac{|z_1[i] - z_2[i]|}{N_{step}} \quad (3)$$

where N_{steps} is the number of steps. We used $N_{steps} = 1000$. We used *classifier-EXT* to identify the syllables and see which class is assigned to the transition between s_1 and s_2 . We compared the variations with a bigger set of generated data in the spectrogram syllable space obtained using UMAP [1].

Finally, we used human judgment to provide an additional qualitative evaluation. We asked two humans to participate in a syllable recognition test organized as described below.

- **Human training phase.** Recognition of a sample of 100 syllables from the training dataset: for the first 50 the person is authorized to look at the repertoire to classify the syllables. After each guess, the person can also look at the correct answer to continue learning to classify syllables. Then, the last 50 syllables have to be recognized. No help is allowed here, and no correct answer can be seen. In the training phase, the available classes are the 16 classes of the repertoire.
- **Human testing phase.** Recognition of a sample of 200 syllables generated after training, without the possibility to look at the repertoire. In the testing phase, the available classes are the 16 classes of the repertoire plus a general X class which, ideally, represents the alternative unknown classes (three *EARLY* classes, *OT* class and *WN* class) recognized by the classifier.

Then, we computed the proportion of agreement beyond the chance agreement (i.e. percentage of agreement that would have occurred anyway) using Cohen’s kappa coefficient [33], [34]. If $\kappa_{Cohen} = 0$, then the agreement is equal to chance agreement. Alternatively, $\kappa_{Cohen} \in (0, 1]$ represents a positive agreement and $\kappa_{Cohen} = 1$ represents a perfect agreement between two judges [33]. We computed κ_{Cohen} between the participants and the *classifier-EXT* and between participants.

F. Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP)

Similarly to t-SNE, Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) [1] is a dimension reduction technique. It can be used to perform non-linear dimension reduction. With respect to t-SNE, it has a higher computational power (i.e., it is faster). Moreover, it has the advantage of not being local (i.e. it takes into account the distance between points that are far in the space). This enables a synthetic, clear and simple representation of the original manifold. The power of using UMAP [1] to reduce the complexity of songbirds spectrograms has first been shown by Sainburg et al. [12]: an exhaustive set of representations show how UMAP helps in representing not only the repertoire of a single species but also different species at the same time. The result is a two-dimensional representation. The axis are not meaningful to identify a discriminant feature (i.e., they do not represent, for example, the pitch of the syllable, or another syllable-related feature). Instead, they are hyperparameters that well represent the given dataset: for instance, the size of the neighborhood used to estimate the manifold structure of the data and the minimum distance apart that points are allowed to be. The tuning of these hyperparameters allows to obtain a more or less local representation of the data (where a higher size of the neighborhood translates in a more local representation) and to pack or not pack together the points in the clusters (where a higher distance allows a more sparse representation).

III. RESULTS

A. Analysis of the training dataset

This section describes the training dataset and highlights the performance of our evaluation metrics on it. Virtually all the samples group into 16 well segregated and mostly non-overlapping clusters (Figure 1a) in the Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) [1] representation (see Section II-F). Besides, the representation shows similarities between syllables *B1* and *B2*, syllables *J1* and *J2*: the clusters of syllables *B1* (cream cluster) and *B2* (light orange cluster), the clusters of syllables *J1* (dark green cluster) and *J2* (light brown cluster) lie very close. These similarities can be noticed also in the spectrogram representation (Figure 1b). For this reason, in further representations using UMAP syllables *B1* and *B2*, *J1* and *J2* have been grouped, respectively, into syllable *B* (keeping the light orange color to represent the cluster) and *J* (keeping the dark green color to represent the cluster). Each template shown in Figure 1a can be compared with the correspondent mean spectrogram (Figure 1b).

The level of confidence of the classifier in making the correct assignment can be represented using the confusion matrix relative to the predictions (see Figure 2). The confusion matrix for *classifier-EXT* (Figure 2) reveals the possible misclassification of syllables *B1/B2* and *J1/J2*. Indeed, as syllables *B1* and *B2* (or syllables *J1* and *J2*) differ the rate at which they are sung more than in frequency content, these syllables may be incorrectly classified when identified individually. The confusion matrix obtained from *classifier-REAL* can be found in the supplementary material (Figure 19).


We compute the Inception Score (IS) of the training dataset based on Eq. 2 to quantify the variety of sounds present in the data. As there are 16 classes of syllables in the dataset, the range of the IS is $IS \in [1, 16]$. The labeling of all single syllables of the training dataset with *classifier-REAL* leads to an inception score of $IS_{train} = 15,92$, reflecting on the clear distinction between classes and their diversity.

B. Evaluation of the model


We trained WaveGAN over the dataset described in Section II-A, to study the ability of the generator to reproduce the various classes of syllables present in the dataset and their variability within a given class. We saved the model every ~ 15 epochs until epoch ~ 1000 . At the beginning of training (epoch 0 in Figure 3), the generator produces for all the classes a noisy output which does not resembles a real syllable. At epoch 15, the generator starts to produce a sound which is coherent in duration but remains noisy and unclear. The resemblance increases over time (at epoch 45, 106 and 514), and finally, at epoch 984 the generator produces syllables showing spectrograms similar with the training data syllables (Figure 1). An example for each class of the repertoire is provided in Figure 23 (supplementary material).

Quantitative evaluation

We used the classifier described in Section II-E1 to obtain a quantitative analysis of the model’s production across



(a) UMAP representation and template.



(b) Mean spectrogram.

Fig. 1. **Repertoire.** (a) Spectrogram syllable space of the training dataset. This representation has been obtained using UMAP [1]. Each cluster represents a class of the repertoire and a template syllable of each class is highlighted with the corresponding spectrogram (an arrow connects each cluster to the corresponding template). Syllables *B1* (cream cluster) and *B2* (light orange cluster), syllables *J1* (dark green cluster) and *J2* (light brown cluster) lie very close, indeed the syllables are very similar and correspond to different speed of repetitions when embedded within a canary phrase. They will be merge in the following experiments. (b) Mean spectrogram computed after envelope alignment of the waveforms. To obtain the spectrogram syllable space and the mean spectrogram we used $1k$ syllables per class (i.e., the training dataset) and their real labels. No classifier has been applied to assign each syllable to the correct class.

learning. Every 15 epochs, we generated $1k$ samples and we used *classifier-EXT* to calculate the probability distribution of the produced syllables to belong to the 16 syllable classes. Figure 4 shows the distribution obtained from the classifier at 4 example epochs: epochs 15, 106, 212, 318, 514 and 984. Each column represents one of the 21 classes of the vocabulary: the 16 classes of the repertoire and 5 alternative unknown classes $x \in X$ (i.e., three classes of *EARLY* generations, the overtraining (*OT*) class, and the artificial white noise (*WN*)). These results have been obtained from an instance of training

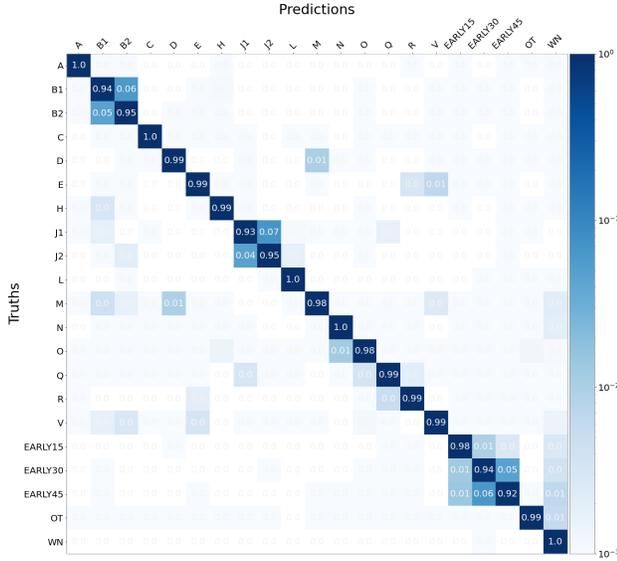


Fig. 2. **Confusion matrix classifier-EXT.** Confusion matrix (i.e. the level of confidence of the classifier in making the correct assignment) relative to *classifier-EXT*. Each square represents the level of confidence of the classifier in making the correct assignment. The level of confidence is expressed on a logarithmic scale using shades of blue.




Fig. 3. **Generations across time: syllable: example of three selected syllable across time.** The syllables has been first generated at epoch 984 and recognized using *classifier-EXT*. Then, the latent vector associated at each syllable has been used to generate the same syllable at epochs 0,15, 45, 106 and 514. From the top to the bottom: syllable *D*, syllable *M*, and syllable *N*. From the left to the right: at epoch 0 the generator produces a noisy sound, not resembling to a real sound; at epoch 15 the generator produces a sample that is coherent in duration but remain noisy and unclear; at epoch 45 the syllables are more distinguishable than at earlier epochs, but remains noisy; as the training goes on, the generation resembles more and more to the real syllables (epochs 106 and 514). The generation obtained at epoch 984 has a clear distinguishable shape, which can be compared with the one shown in Figure 1. Here, the generator obtained from instance *Ex 6* in Figure 5 has been used to generate the syllables across time; latent space dimension $ld = 3$; *Ep.* stands for epoch.

where the latent space dimension was fixed at $ld = 3$. The number of syllables belonging to an alternative class decrease with time, whereas syllables classified as belonging to the training repertoire become more frequent. After ~ 200 epochs

TABLE I
INCEPTION SCORE (IS) FOR ALL INSTANCES. FOR OUR DATASET, $IS \in [1, 16]$.

Dataset	Epoch	IS
Training	-	15,92
Ex 0	800	13,51
Ex 1	211	8,91
Ex 2	515	5,65
Ex 3	512	12,29
Ex 4	605	11,89
Ex 5	408	11,08
Ex 6 (baseline)	984	13,07
Ex 7	988	13,61
Ex 8	802	12,66
Ex 9	605	12,19

of training, the generator is able to cover all the syllable classes of the repertoire (Figure 5a). In parallel, the total variance (considering all classes of the repertoire and the 5 alternative unknown classes) decreases over time (Figure 5c). Indeed, at early stages of training the majority of the samples generated are classified by *classifier-EXT* as elements of an alternative unknown class $x \in X$ (Figures 5d, 5e, and 5f). Then, the number of samples recognized by *classifier-EXT* as belonging to a class of the repertoire increases leading to an increasing in the average number of syllables per class (Figure 5b). Importantly, syllables are produced in all classes in a relatively even distribution across the various classes of the repertoire. After epoch ~ 600 the average number of produced syllables assigned to each class remains stable. At the same time, the number of syllables belonging to one of the alternative classes $x \in X$ (*EARLY15*, *EARLY30* and *EARLY45*) decreases over time (Figures 5d, 5e, and 5f). Classes *EARLY15* and *EARLY30* decrease faster with respect to *EARLY45*, which never reaches a ratio comparable to the one found in the training data. Although these results suggest that the generator produces better and better samples over time, further analysis are needed to understand if the quality of the produced samples keep increasing after epoch 600.

To compare 10 different training instances of 3-dimensional WaveGAN, we first computed the Inception Score (IS) using Eq. 2 at each saved epoch for the 10 instances. Figure 6 shows the IS across time (referred to as IS_{train} and defined in Section III-A) represented here by the black line. Table I summarizes the IS of the training dataset and each of the selected instances: we considered the highest IS within the time range shown in Figure 6. The IS of *Ex 7*, *Ex 8* and *Ex 9* increase across time and, eventually, reach a maximum value between 12 and 14 (for reference, the IS obtained from the training dataset is $IS_{train} = 15,92$); contrarily, the IS of *Ex 2* and *Ex 3* increase at the beginning but remains low during all the training. The IS of *Ex 1*, *Ex 5*, *Ex 4*, *Ex 6* and *Ex 10* suddenly break around epoch 800: this behavior is consistent across multiple iterations of the GAN shown in Figure 34. We refer to this behavior as overtraining. A complete analysis of the 10 instances of training of a 3-dimensional WaveGAN are shown in Figures 34 and 35 (supplementary material).




Fig. 4. **Distribution of syllable classes recognized by the classifier over training epochs.** Distribution of $1k$ generated samples after 15, 106, 212, 318, 514, and 984 epochs of training. Each column represents a class of the vocabulary: a syllable from the repertoire, or an alternative unknown class $x \in X$ (*EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). The number of syllables belonging to a class $x \in X$ decreases over time, whereas the average number of syllables belonging to the classes of the repertoire increases. The latent space dimension is $ld = 3$ and *classifier-EXT* has been used to classify the generated data.

Qualitative evaluation

For qualitative evaluation, we focused on instance *Ex 6* of training because it displays a high maximal Inception Score (see Figure 6 and Table I) and remains stable until epoch ~ 1000 (see Figure 5). *Ex 6* is also the example we used since the beginning of Section III-B as an example.

The generator produces a higher number of good syllables over time and, at the end of the training, all the syllables can be produced by the generator. The mean spectrogram obtained from the generated data at epoch 984 can be compared with the one obtained from the training dataset (Figure 1b). Figure 8 displays the spectrogram syllable space of the generated data and the training data obtained using UMAP [1]. The generated data are initially relatively homogeneous and quite different from the syllables from the training dataset and are therefore grouped together as an additional cluster well separated from the training data clusters (light blue points in Figures 8(a)). Across training of the WaveGAN, the generated data diversify, leading to a spread of the corresponding dots in the spectrogram syllable space (Figures 8(b-d)). Moreover, as the generated syllables resemble more and more the training data, their spectrogram syllable space comes close or within the corresponding cluster of the training data (brown points). Further observations regarding the comparison between the spectrogram syllable space of training and generated data are discussed in Figures 25 and 26 in the supplementary material.

The spectrogram syllable space obtained using UMAP [1] for the generated data shows smoother transitions from one cluster to another (Figure 9). These transitions can be either represented by points belonging to one of the classes of the repertoire (e.g., between *J* and *C*) or to class *X* (e.g.,

between *N* and *Q*). This observation highlights an interesting perspective for which elements classified in class *X* can be seen as intermediate elements between two syllables. We will deal with this concept in Section III-C. For further observations, Appendix IV-E shows how, combining the spectrogram syllable space and the mean spectrogram, it is possible to compare three consecutive epochs of training to check the stability of the generator.

Finally, human judgment confirms the goodness of the classifier and of the generated data. Each judge evaluated the same set of 200 samples generated at epoch 984 using the generator of instance *Ex 6*. The judges had the possibility to classify each of them as an element of one of the 16 classes of the repertoire or as an element of an alternative unknown class $x \in X$. The Cohen’s kappa coefficients κ_{Cohen} obtained across judges and with respect to *classifier-EXT* are comparable: $\kappa_{Cohen}(Judge1, Judge2) = 0,74$, $\kappa_{Cohen}(Judge2, classifier-EXT) = 0,79$ and $\kappa_{Cohen}(Judge1, classifier-EXT) = 0,73$ (these values are collected in Table III in the supplementary material).

C. Latent space exploration

To explore the structure of the latent space, we used *Ex 6* at epoch 984. We investigate $16k$ generated sounds that can be represented using UMAP [1] (Figure 9) or in the latent space, each generated sound being associated to a given vector in the latent space (Figure 10). The latent space is dense (Figure 10a), i.e. a high number of latent vectors give rise to a syllable that *classifier-EXT* recognizes as an element of the repertoire, and only a minority of elements is classified as belonging to class *X*. The topology of the latent space is

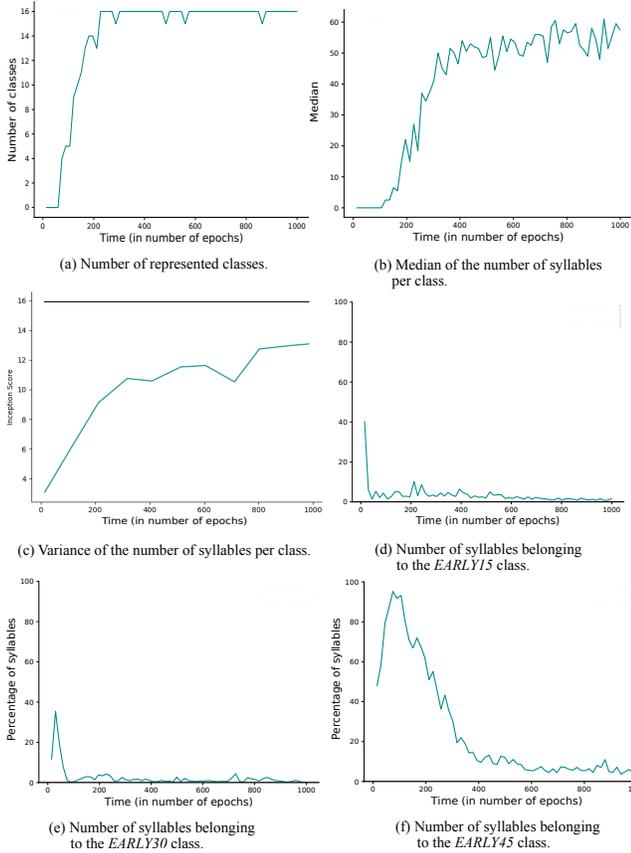


Fig. 5. **Analysis of a 3-dimensional generator using classifier-EXT.** Statistical analysis performed on the classifier distribution of one instance of training. Panel (a) shows the number of classes represented by the generated data: that is, at each epoch, how many syllables of the repertoire are covered by the generator. The number of syllable reaches the maximum value 16 relatively quickly and then stays high during all the training. Panel (b) shows the median of the number of elements per class: across time, with some instability, we can observe an increase of the median across time. To compute the quantities in panels (a) and (b), alternative unknown classes $x \in X$ have not been taken into account. Panel (c) shows the evolution of the variance of how many syllable per class have been produced. Here, $x \in X$ classes are included. The variance starts at a high value when the majority of the samples produced are not classified as syllables of the repertoire, then it decreases when the generator becomes better at producing syllables. Panels (d-f) show the percentage of syllables that are classified as belonging to one of the alternative classes $x \in X$: from the left to the right, classes *EARLY15*, *EARLY30* and *EARLY45*.

consistent with the acoustic structure of the produced syllables: clusters that are close in the latent space represent syllable classes with strong similarities in their spectrograms, and the distance between clusters thus reflects on acoustic differences between syllables. Clusters that are close in the data (Figure 9) are close in the latent space (Figure 10): for instance, the clusters relative to syllables *B* (cream cluster), *H* (light green cluster), *L* (light blue cluster), *R* (light pink cluster) and *J* (turquoise cluster). Moreover, some clusters (e.g., syllable *O*) show an higher sparseness than others (dark blue cluster in Figures 10e and 10f).

To explore the continuity of the latent space we used two different strategy to (1) explore the latent space in all the direction starting from one point (one component variation described in Section II-E2), and (2) observe the transition




Fig. 6. **Inception Score (IS) across time.** IS relative to the generator of 10 instances of a 3-dimensional WaveGAN, with respect to IS_{train} (black line) In particular, *Ex 6* (light blue line) is the instance used as baseline. the IS for some instances (e.g., *Ex 6* in light blue and *Ex 7* in gray) increases over time, whereas the IS remains low over time in the case of instance *Ex 2* (orange line). Alternatively, it increases over time until it drops (*Ex 0* in purple and *Ex 9* in red). Here, for each instance we generated 16k samples every ~ 100 epochs and we used *classifier-REAL* to classify them.




Fig. 7. **Mean spectrogram.** Mean spectrogram of 1k syllables generated at epoch 984. All the syllables are clear and distinguishable. They can be compared with the repertoire in Figure 1. Here, the training has been done using 3 latent dimensions ($ld = 3$), the generator obtained from instance *Ex 6* in Figure 5 has been used to generate the syllables across time and *classifier-EXT* has been used to identify the generated syllables.

from one syllable to another (three components variation described in Section II-E2). First, we selected a random latent vector $z \sim R^3([-1, 1])$ to create a baseline syllable, then continuously move from one representation to another by changing one dimension of the latent space by a fix variation step. As shown in Figure 11, most gradual changes in the latent representation of a syllable leads to a small incremental change in the acoustics of the produced syllable as revealed by the spectrogram representation. However, for a step of 0.05 in the latent space, some transitions are still steep, as highlighted by the red squares in Figure 11. These non-smooth transitions usually represent the transitions between syllable classes (e.g., first *C* and *J2*, then *V* and *V* on Figure 11). For these particular transitions, we considered the two consecutive syllables obtained from the first variation step (i.e., two consecutive syllables contained in two consecutive red squares), and we applied a smaller variation step. The bottom-left panel of Figure 11 shows the latter operation applied to the first non-smooth transition in the first component and highlights a new non-smooth transition that needs to be investigated. Using the same procedure, we applied a smaller variation step and reached a point at which we cannot see any more clear evidence of non-smooth transitions (bottom-right panel of Figure 11). The investigation of the second non-




Fig. 8. **Spectrogram syllable space across time.** Spectrogram syllable space obtained from the training dataset ($16k$ syllables) and $1k$ syllables generated at epochs (a) 15, (b) 106, (c) 514 and (d) 984 using UMAP [1]. Brown points represent the training data, while light blue points the generated data. These four figures are different because the analyzed dataset differs for the $1k$ generations specific of each epoch. Here, the training has been done using $ld = 3$, the generator of instance *Ex 6* has been used to generate the syllables.




Fig. 9. **Spectrogram syllable space of the generated data.** Spectrogram syllable space of $16k$ generated syllables obtained using UMAP [1]. Each cluster/color correspond to one class of the repertoire and class X (in white) represents the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Here, the training has been done using $ld = 3$, the generator obtained from instance *Ex 6* in Figure 5 has been used to generate the syllables at epoch 984 and *classifier-EXT* has been used to identify the generated syllables.

smooth transition highlighted in the upper panel of Figure 11, and additional examples from the second and the third components can be found in the supplementary material.

A better visualization of smooth transitions between syllable classes can be observed in the spectrogram syllable space obtained using UMAP [1] shown in Figure 12a. The exploration obtained varying, one by one, the first component of a random latent vector $z \sim R^3([-1, 1])$ by a step equal to $v_{step} = 0.001$ shows a smooth transition between one syllable to another. It is not surprising that a difficulty in the classification arises between syllable *V* (magenta cluster) and syllable *C* (orange cluster): indeed, syllable *V* is a syllable that shares its content with other syllables of the repertoire, causing uncertainty and errors for the classifier. The spectrogram syllable space of the training dataset (Figure 1a) and then the of the generated




Fig. 10. **Structure of the latent space.** We selected $16k$ random vectors $z \sim R^3([-1, 1])$ from a 3-dimensional latent space, and we generated $16k$ syllables. Then, we used *classifier-EXT* to classify each new generation. Each point in panels (a-f) corresponds to a 3-dimensional latent vector. Each color corresponds to a class of the vocabulary (i.e., a class of the repertoire or an alternative class). As for the spectrogram syllable space obtained using UMAP [1] in Figure 9, the alternative classes have been represented as a cumulative class X. Panel (a) shows a global representation containing all the $16k$ latent vectors, giving an idea of the density of the latent space. Panels (b-f) represent five slices from the whole space. Points close in the latent space are also close in the spectrogram syllable space (Figure 9) obtained using UMAP [1]. The difference between the structure of the clusters can be observed: some of them are more sparse than others. For instance on panels (d) and (e), the dark blue cluster representing syllable *O* has several white points – i.e., elements belonging to class X – inside or at the border of the cluster. The syllables have been produced by the 3-dimensional generator obtained from instance *Ex 6*.

data (Figure 9) already show multiple cluster locations for syllable V . Additional figures and comments can be found in the supplementary material (Figures 21, 22 and 32).




Fig. 11. Exploration of the latent space: one component variation. First component. We selected a random latent vector $z \sim R^3([-1, 1])$ to create a baseline syllable. Then, we moved one by one the components of the vector by a variation step equal to $v_{step} = 0.05$. We observed all the syllables produced to look at how they evolve and if there are non-smooth transitions. The upper panel shows the exploration of the first component of the latent vector obtained with $v_{step} = 0.05$. The syllables highlighted by the red squares (i.e., first C and $J2$, then V and V) are an example of non-smooth transitions. For these particular transitions, we considered the two consecutive syllables obtained from the first variation step (i.e., two consecutive syllables contained in two consecutive red squares) and we applied a variation step of $v_{step} = 0.01$ to the first component of the latent vector to generate intermediate syllables. The bottom-left panel shows the latter operation applied to the first non-smooth transition in the first component and highlights a new non-smooth transition that needs to be investigated. To do so, as shown in the bottom-right panel, we used a variation step equal to $v_{step} = 0.001$ to generate intermediate syllables. We have reached a point at which we cannot see any more clear evidence of non-smooth transitions. The syllables have been produced by the 3-dimensional generator obtained from instance *Ex 6* and the name of the syllable for this analysis has been provided by *classifier-EXT*.

We computed the three components transition from syllable M to syllable D . These two syllables are not adjacent in the spectrogram syllable space of the generated data (see Figure 9). The variational data across the generated data (left panel of Figure 12) give rise to a smooth change of syllable class (right panel of Figure 12). Interestingly, the transition between syllable M and syllable D (Figure 12, right panel) shows that the intermediate syllables between two classes are sometimes recognized as class X (white points between syllable M (turquoise cluster) and syllable N (blue cluster)). More examples of transition between syllables can be found in Figure 33 (supplementary material).

The latent space exploration, combined with the obtained for the generated data and shown in Figure 9, highlights the fact that the generator is not only producing samples from the training dataset but also other samples. That is, the generator is able to change from one syllable to another without abrupt changes. Instead, a small change in the latent space results in a small change in the sound space. This allows to move realistically between two syllables.

D. Latent space dimension

Until now, we have shown results obtained from a 3-dimensional WaveGAN. We now show a comparison between different N -dimensional WaveGAN generators, where $N =$




Fig. 12. Latent vector exploration. Panel (a) represents the first component variation. A latent vector has been randomly selected and a step $v_{step} = 0.001$ has been applied to its component, one by one, in order to generate $1k$ syllables. Panel (b) represents the transition between M (turquoise cluster) and syllable D (yellow cluster). The transition between one syllable to the other has been done generating intermediate syllables as described in Equation 3 (Section II). The left panel shows the variational data (generated at each step) (light blue points) and $16k$ generated data from the same model at the same epoch (brown points). Each cluster/color in the right panels corresponds to one class of the repertoire and class X (in white) represents the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Here, the training has been done using $ld = 3$. The generator obtained from instance *Ex 6* of the training has been used *classifier-EXT* has been used to identify the generated syllables from the generator of instance *Ex 6*.

1, 2, 3, 4, 5, 6. Although all conditions allows the generator to reach the ability of producing all types of syllables present in the training dataset, 1-dimensional and 2-dimensional WaveGANs converge later with respect to higher dimensional WaveGANs (Figure 13a). At the same time, 1-dimensional and 2-dimensional WaveGANs show a slower increase in the average number of syllables belonging to each class of the repertoire (Figure 13b) and a slower decrease in the variance of the number of syllables per class. As already mentioned in Section III-B, a generalized instability could appear towards the end of the training (what we refer to as overtraining). This overtraining phenomenon appears across various latent space dimensions (e.g., overtraining happens for $ld = 4, 5, 6$ - respectively, the blue, the green and the pink line). This can be seen in a decrease in the median evolution (Figure 13b) and in the sudden drop in the inception score (Figure 13c). A qualitative measure of the generated syllables obtained from WaveGANs of different dimensions (e.g., for different conditions of the latent space dimension) can be found in Figure 37 (supplementary material).

In conclusion, a 3-dimensional WaveGAN works nicely and as good as higher dimension WaveGANs and, at the same time, lower dimensional WaveGANs are longer to be trained and/or




Fig. 13. **Comparison between different latent space dimensions.** Each line represents one instance of training at a particular latent space dimension. Panel (a) shows how many syllables of the repertoire are covered by the generator across time. Panel (b) shows how many syllable per class have been generated in average. The dark lines show the evolution of the mean, whereas the light lines shows the evolution of the median. To build these two panels (a) and (b), $1k$ generations have been generated every 15 epochs, *classifier-EXT* has been used to provide the classification, and only the repertoire’s classes have been taken into account when plotting. Panel (c) shows the evolution of the inception score over time. To build this panel $16k$ syllables have been generated every 200, and *classifier-REAL* has been used to provide the classification. A 3-dimensional WaveGAN (orange line) reaches convergence as good as higher-dimensional WaveGANs (blue, green and magenta lines) and better than lower dimensional WaveGANs (purple and yellow lines). We varied the latent space dimension as $ld = 1, 2, 3, 4, 5, 6$ and we kept the training dataset fixed. A complete version of this figure can be found in the supplementary material (Figure 36).

do not show a similar performance.

E. Training dataset dimension

We compared three datasets of different size: 4k syllables, 8k syllables, and 16k syllables (the latter is the same used for all our previous analysis, and the example shown in this section is the same as the one shown in Section III-B). For all the datasets, the generator is only able to produce a limited number of syllable types at the beginning of the training, whereas it becomes able to reproduce all of them after some epochs of training (Figure 14a). The smaller dataset (green line) shows a slower convergence compared to the 16k dataset training. Likewise, the average number of syllables generated per class across time (Figure 14b) increases more quickly when WaveGAN is trained with a dataset of 16k syllables (blue line) and 8k syllables (red line), than when WaveGAN is trained with a dataset of 4k syllables (green line). Similarly, the inception score (Figure 14c) shows an increase with the size of the training dataset: the bigger the dataset, the higher the inception score. Such quantitative evaluation is coherent with the qualitative evaluation obtained comparing the mean spectrogram at the end of the training for the different conditions (see Figure 39 in the supplementary material). In conclusion, a larger dataset allows WaveGAN to converge faster and better (Figure 14).

IV. DISCUSSION

A. Summary of the main results

In this paper we tested the ability of WaveGAN [20] to produce canary syllables when the latent space dimension is low. Originally, Donahue et al. [20] trained WaveGAN on a set of wild recordings from different bird species: due to the variability of the dataset and to the lack of enough recordings for each specie, the generated samples were noisy. By introducing a clean dataset of recordings from a single adult canary, we simplified the complexity of the training

dataset and obtained qualitatively better results. The fact that we deal with single syllables simplifies the analysis of the performance of the generator and the exploration of the latent space.

WaveGAN has been originally trained with a 100-dimensional latent space. We showed that WaveGAN can produce good syllables even if the latent space dimension is reduced to $ld = 3$ (see Figure 3). The generator can produce sounds with similar acoustic properties as the canary syllables used to train it, with a diversity of produced sounds that covers the training dataset and interpolate in between training data. Moreover, we showed that introducing a dataset of a lower size results in a slower learning (see Figure 14). Finally, we explored the structure of the latent space: we showed the continuity by studying (1) how small changes in the latent space influence the generations (see Figure 11); and (2) the interpolation between well-defined syllables (see Figure 12a).

B. Limitations

There is still a poor understanding of the mechanisms involved in the training of GANs. A limitation of WaveGAN [20], and GANs in general, is the absence of a stopping criteria for the training: this introduces an incertitude about how to evaluate the learning and its stability (i.e., in number of epochs rather than depending on the loss function value). Moreover, the hyperparameter sensibility and the limitations related to overfitting may introduce some complications in the exploration of GANs performance.

C. Evaluation of the performances

There is currently no consensus on how to evaluate generative models, and especially GANs [17], [31]. However, generative models are generally evaluated using both a quantitative and a qualitative measure of their output. A quantitative evaluation is useful to determine whether or not the generator is producing only examples on which the classifier was trained.




Fig. 14. **Comparison between datasets of a different size.** We varied the dataset size as $d = 23456, 3600, 1600$ and we kept fix the latent space dimension at $ld = 3$. Each line in this figure represents one instance of training at a particular dataset size condition. Panel (a) shows how many syllables of the repertoire are covered by the generator across time. Panel (b) shows how many syllable per class have been generated in average. The dark lines show the evolution of the mean, whereas the light lines shows the evolution of the median. To build these two panels (a) and (b), $1k$ generations have been generated every 15 epochs, *classifier-EXT* has been used to provide the classification, and only the repertoire’s classes have been taken into account when plotting. Panel (c) shows the evolution of the inception score over time. To build this panel $16k$ syllables have been generated every 200, and *classifier-REAL* has been used to provide the classification. A dataset of bigger size (blue line) allows better and faster convergence than having a dataset of lower sizes (red and green lines). A complete version of this figure can be found in the supplementary material (Figure 38).

A qualitative measure is necessary to truly understand the quality of the produced samples and their comparability to real recordings in term of whether or not they are comprehensible to an external judge. The drawbacks of a qualitative measure is that it presupposes the interpretability of the data by humans, which is not always possible [31]. Moreover, it could be biased, expensive, not efficient in detecting overfitting.

1) *Quantitative evaluation:* Inception Score (IS) [18] is a broadly used quantitative measure which gives an idea about whether or not the generator model is able to produce a wide enough diversity in the samples. Nevertheless, IS is not able to detect if the generator is producing only examples that belong to the training dataset or if the generator enter ‘collapse mode’ (i.e., it produces always a small set of classes) or overfitting [35], [36]. Thus, IS should not be used as a holistic method to evaluate the performance of a model and must be combined with another evaluation method [35]. For instance, Donahue et al. [20] measured the Euclidian distance in the space of log-Mel spectrograms (1) within the training and the generated data (to measure the intra-dataset diversity with respect to the training data), and (2) between the training and the generated data (to show the ability of WaveGAN of producing sounds not belonging to the training dataset).

During the training, the median of the number of syllables produced per class continuously increase over time (see Figure 34). A similar behavior can be seen for the IS (see Figure 6). Eventually, one can observe a drop if overtraining begins: we call *overtraining* the fact that, at a certain epoch, there is a drop in the IS (see Figure 6) or in the statistical properties of the classifier distribution (see Figure 34b). However, the concept of overtraining remains unclear. We tried to characterize it using a specific class (*OT*) which arises at advanced stages of training, but never becomes the most represented class (see Figure 35). Instead, after the training drops, there is an exponential increase in the number of element assigned to early classes (see Figure 34(d-f)).

2) *Qualitative evaluation:* A qualitative evaluation based on human judgment has been proposed by Salimans et al. [18] and Denton et al. [37] to evaluate GANs trained on MNIST

or CIFAR-10 (image datasets). Similarly, Donahue et al. [20] used human judgment to evaluate WaveGAN performance when training on SC09 (speech dataset). First, we rely on the mean spectrogram obtained from the generated data over time (see Figure 7 and Figure 24) to observe at once how the average quality of syllables increases over time, becoming more and more comparable with the repertoire (see Figure 1). Then, we rely on the UMAP [1] representation of the spectrograms to compare the generated data with the training data. On the one hand, the generated data belong to the same cluster as the training data (see Figure 8) if represented together. On the other hand, a continuity across clusters arises when observing the generated data alone (see Figure 9). In the case of data not familiar to humans (such as canary syllables), one could design a behavioral protocol to test the accuracy by measuring how canaries perceive the generated data. Although such an experiment could be very interesting, it is complicated to settle and require an expertise outside of the machine learning domain.

D. Structure of the latent space

The exploration of the latent space allows to study whether or not the generator is able to produce a continuous space, and interpolate between its element [10]. Moreover, it allows to compare the structure of the sound output space (spectrograms) and the latent space: a desirable property is that clusters that are close by in the sound space are close by in the latent space [9]. The exploration of the latent space by interpolating between different latent space elements has been proposed by Drysdale et al. [38] after training WaveGAN on a drums dataset, and by Blaauw and Bonada [39] after training a Variational Autoencoder (VAE) on a speech corpus. Alternatively, Hsu et al. [11] train a VAE on a speech corpus and decompose the latent space into orthogonal latent attribute representations and explore how speech attribution changes (e.g., they move from a female speaker to a male speaker).

In our study, we analyzed (1) how a small change in the latent space affects the generated syllable, (2) the transition

between two different generated syllables, and (3) whether or not clusters that are close by in the spectrogram syllable space (obtained using UMAP) are close by in the latent space. The training dataset generates well separated clusters in the spectrogram syllable space, (see Figure 1a), whereas the clusters obtained from generated data appears more continuous with transition “paths” between clusters (see Figure 9). By interpolating between two syllables in the latent space using small steps, it is possible to move continuously in the spectrogram syllable space (see Figure 12a). Alternatively, if the applied step in the latent space is too big, there are non-smooth transitions between syllables (see Figure 11). To our knowledge, there is no other GAN performing exploration of one latent dimension at a time (see Figures 11 and 12a). This is possible because we have a low-dimensional latent space while previous works have high-dimensional ones.

E. Perspectives

As future work to this study, several perspectives could be considered. In the next paragraphs, we summarize how the WaveGAN generator could help the understanding of syllable generation in songbirds when using a different training dataset, or when integrating it in a more complex model (e.g., a vocal learning model).

1) *Different training dataset*: One could explore the possibility of enlarging the training dataset by including (1) more classes of syllables, (2) more birds (in terms of number), (3) recordings from different species or (4) recordings from juveniles. In principle, the addition of more syllables or recordings from different canaries would introduce complexity in the training dataset and would increase the computational time (i.e., the number of epochs needed to obtain a generator able to output good syllables). On the contrary, the attempt of using the model on other species is not trivial. For instance, although zebra finches represent an important model for vocal learning, they produce syllables with non-harmonic components. Thus, it is more difficult to identify/evaluate for a human observer compared to canary syllables. Training WaveGAN on both juveniles and adult data is an interesting and rather straightforward extension of this work. For instance, it can be useful to design a motor control function that could produce any possible canary sounds, from juveniles to adults. This would be useful in computational experiments in order to model the sound productions at different stages of learning. It could also be useful for behavioral experiments with birds.

2) *Vocal learning model*: Our results support the use of generative models as motor function in vocal learning models, as an alternative to other sound synthesis. Indeed, motor control is most often modeled using Ordinary Differential Equations (ODEs) representing the sound generation organ in vocal learning models [40]–[43]. Such models are usually based on the anatomical structure of the vocal tract and the respiratory apparatus, and, especially for humans, can include a large number of parameters. However, learning to control a high-dimensional motor apparatus in current theoretical frameworks for sensorimotor learning remains challenging [14], [44], [45] and often leads to unrealistic convergence time [15]. A lower

dimensional motor control reduces the number of parameters that need to be learned, thus simplifying the learning (w.r.t an equivalent topology of latent space). Reducing the number of dimensions controlled by the learning algorithm therefore appears as a necessary first step to design vocal learning models with more biologically plausible learning rules. Consequently, we believe in the possibility that generative models, and in particular the generator of WaveGAN, can be used as motor control in vocal learning models. Biologically, they would represent a premotor layer rather than the control parameters of a vocal organ. Such premotor “module” simplifies the dimensionality of the learning problem and can be seen as an equivalent of embodiment: appropriate body features can facilitate the motor control algorithms [46]. Moreover, our approach using a continuous generator is validated these experimental data that show the ability of syrinx to produce a variety of sounds is present early during song ontogeny [47].

3) *Generation of canary songs*: A generator model able to produce sequences of syllables would be difficult to evaluate: indeed, it would be difficult to perform the qualitative analysis we propose here (e.g., the mean spectrograms or the transition from one syllable to another in the latent space). Alternatively, a model able to generate single syllables, as the one proposed in this work, may be an adequate tool to generate full canary songs. Such model would require (1) an estimation of the distribution of the delay within syllables of the same class, and (2) a probabilistic model estimating (i) how many time a syllable is repeated and (ii) the upcoming syllable in the phrase. Indeed, canary songs are composed of sequences of phrases, which are themselves repetitions of the same syllable with smooth transitions. The number of syllable repetitions in one phrase is variable. Thus, a generator trained to produce bouts of a few seconds would not be able to reproduce the variability of canary song with much flexibility. On the contrary, a generator like the one proposed, which is able to produce smooth controllable transitions between syllables, is likely to be a good tool to generate full canary songs. Finally, unrelated to birdsong domain, such low-dimensional GAN and the diverse explorations it enables, can help in a general fashion for future works on generative models.

ACKNOWLEDGMENT

We would like to thank Catherine Del Negro, Aurore Cazala and Juliette Giraudon for the recording and transcription of the canary data. We also thank Inria for the CORDI-S PhD fellowship grant and the LabEx BRAIN for the PhD extension grant. Experiments presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d’Aquitaine (see <https://www.plafrim.fr>).

REFERENCES

- [1] L. McInnes *et al.*, “Umap: Uniform manifold approximation and projection for dimension reduction,” *arXiv preprint arXiv:1802.03426*, 2018.
- [2] J. Hillenbrand *et al.*, “Acoustic characteristics of american english vowels,” *The Journal of the Acoustical society of America*, vol. 97, no. 5, pp. 3099–3111, 1995.
- [3] J. L. Miller and A. M. Liberman, “Some effects of later-occurring information on the perception of stop consonant and semivowel,” *Perception & Psychophysics*, vol. 25, no. 6, pp. 457–465, 1979.
- [4] P. Kuhl, “Early language acquisition: cracking the speech code,” *Nature reviews neuroscience*, vol. 5, no. 11, p. 831, 2004.

- [5] M. Chakraborty and E. Jarvis, "Brain evolution by brain pathway duplication," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 370, no. 1684, p. 20150056, 2015.
- [6] A. Doupe and P. Kuhl, "Birdsong and human speech: common themes and mechanisms," *Annual review of neuroscience*, vol. 22, no. 1, pp. 567–631, 1999.
- [7] J. E. Markowitz *et al.*, "Long-range order in canary song," *PLoS Comput Biol*, vol. 9, no. 5, p. e1003052, 2013.
- [8] I. Goodfellow *et al.*, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [9] A. Roberts *et al.*, "A hierarchical latent vector model for learning long-term structure in music," *arXiv preprint arXiv:1803.05428*, 2018.
- [10] A. Radford *et al.*, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [11] W.-N. Hsu *et al.*, "Learning latent representations for speech generation and transformation," in *Proc. Interspeech 2017*, 2017, pp. 1273–1277. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-349>
- [12] T. Sainburg *et al.*, "Finding, visualizing, and quantifying latent structure across diverse animal vocal repertoires," *PLoS computational biology*, vol. 16, no. 10, p. e1008228, 2020.
- [13] F. Benureau, "Self exploration of sensorimotor spaces in robots," pp. English. NNT : 2015BORD0072 . tel-01 251 324v2, 2015.
- [14] M. Rolf *et al.*, "Goal babbling permits direct learning of inverse kinematics," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 216–229, 2010.
- [15] S. Pagliarini *et al.*, "A bio-inspired model towards vocal gesture learning in songbird," in *ICDL-EpiRob*. IEEE, 2018, pp. 269–274.
- [16] A. v. d. Oord *et al.*, "Pixel recurrent neural networks," *arXiv preprint arXiv:1601.06759*, 2016.
- [17] I. Goodfellow, "Nips 2016 tutorial: Generative adversarial networks," *arXiv preprint arXiv:1701.00160*, 2016.
- [18] T. Salimans *et al.*, "Improved techniques for training gans," in *Advances in neural information processing systems*, 2016, pp. 2234–2242.
- [19] H.-W. Dong *et al.*, "Musegan: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment," in *AAAI*, 2018.
- [20] C. Donahue *et al.*, "Adversarial audio synthesis," *arXiv preprint arXiv:1802.04208*, 2018.
- [21] M. Arjovsky *et al.*, "Wasserstein gan," *arXiv preprint arXiv:1701.07875*, 2017.
- [22] G. Giraudon *et al.*, "Labeled songs of domestic canary m1-2016-spring (serinus canaria) (version 0.0.1) [data set]." *Zenodo*, vol. <http://doi.org/10.5281/zenodo.4736597>.
- [23] A. v. d. Oord *et al.*, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [24] I. Gulrajani *et al.*, "Improved training of wasserstein gans," in *Advances in neural information processing systems*, 2017, pp. 5767–5777.
- [25] H. Jaeger, "The "echo state" approach to analysing and training recurrent neural networks-with an erratum note," *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, vol. 148, no. 34, p. 13, 2001.
- [26] M. Lukoševičius and H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Computer Science Review*, vol. 3, no. 3, pp. 127–149, 2009.
- [27] N. Trouvain and X. Hinaut, "Canary song decoder: Transduction and implicit segmentation with ESNs and LTSMs," in *ICANN*. Springer, 2021.
- [28] L. Pedrelli and X. Hinaut, "Hierarchical-task reservoir for online semantic analysis from continuous speech," *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [29] N. Trouvain *et al.*, "Reservoirpy: an efficient and user-friendly library to design echo state networks," in *ICANN*. Springer, 2020, pp. 494–505.
- [30] X. Hinaut and N. Trouvain, "Which hype for my new task? hints and random search for Echo State Networks hyperparameters," in *ICANN*. Springer, 2021, pp. 83–97.
- [31] A. Borji, "Pros and cons of gan evaluation measures," *Computer Vision and Image Understanding*, vol. 179, pp. 41–65, 2019.
- [32] P. Isola *et al.*, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [33] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and psychological measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [34] R. Artstein and M. Poesio, "Inter-coder agreement for computational linguistics," *Computational Linguistics*, vol. 34, no. 4, pp. 555–596, 2008.
- [35] S. Barratt and R. Sharma, "A note on the inception score," *arXiv preprint arXiv:1801.01973*, 2018.
- [36] J. Gui *et al.*, "A review on generative adversarial networks: Algorithms, theory, and applications," *arXiv preprint arXiv:2001.06937*, 2020.
- [37] E. L. Denton *et al.*, "Deep generative image models using a laplacian pyramid of adversarial networks," in *Advances in neural information processing systems*, 2015, pp. 1486–1494.
- [38] J. Drysdale *et al.*, "Adversarial synthesis of drum sounds," in *Proceedings of the International Conference on DAFx*, 2020.
- [39] M. Blaauw and J. Bonada, "Modeling and transforming speech using variational autoencoders," *Morgan N, editor. Interspeech 2016; 2016 Sep 8-12; San Francisco, CA: ISCA; 2016. p. 1770-4.*, 2016.
- [40] A. Amador *et al.*, "Elemental gesture dynamics are encoded by song premotor cortical neurons," *Nature*, vol. 495, no. 7439, p. 59, 2013.
- [41] T. Gardner *et al.*, "Simple motor gestures for birdsongs," *Physical review letters*, vol. 87, no. 20, p. 208101, 2001.
- [42] G. Westerman and E. R. Miranda, "Modelling the development of mirror neurons for auditory-motor integration," *Journal of new music research*, vol. 31, no. 4, pp. 367–375, 2002.
- [43] S. Maeda, "Compensatory articulation in speech: analysis of x-ray data with an articulatory model," in *Eurospeech*, 1989.
- [44] R. Reinhart, *Reservoir computing with output feedback*. PhD Thesis. Bielefeld University, Germany, 2011.
- [45] R. Parr and S. Russell, "Reinforcement learning with hierarchies of machines," in *Advances in neural information processing systems*, 1998, pp. 1043–1049.
- [46] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.
- [47] A. Maxwell *et al.*, "Syringeal vocal folds do not have a voice in zebra finch vocal development," *Scientific reports*, vol. 11, no. 1, pp. 1–18, 2021.

SUPPLEMENTARY MATERIAL

I. AUTOMATIC SELECTION

A. Pre-processing

As introduced in Section II-A, to select the syllables we used three parameters: amplitude threshold, minimal duration of the syllable and duration of the gap between two consecutive syllables. Figure 15 shows a summary of the selection procedure for syllable *J2*. The upper panel shows the downsampled phrase. The lower panel highlights each syllable of the phrase: for each syllable, the green dashed line represents the onset (i.e., the beginning of the syllable), and the blue dashed line represents the offset (i.e., the end of the syllable).




Fig. 15. **Selection of syllable *J2***. The upper panel shows the spectrogram of an example phrase of syllable *J2*. The lower panel shows the selected syllables: the green vertical lines represent the onset of each syllable, the blue vertical lines represent the offset of each syllable. We used onset and offset to determine where each syllable begins and ends.

B. Semi-automatic error detection

We manually check the syllables after the automatic selection. In this section we explain which errors could arise from the automatic selection and how we solved them.

First, our algorithm could have failed in cutting a recording because of the presence of a too short gap between syllables. In this case we obtained samples containing more than one syllables. Alternatively, some samples could be too short, which means they contain only a part of the syllable. Finally, a bad initial classification (e.g. a phrase of class *A* was wrongly classified as class *M*) could lead either to well-selected syllables belonging to the wrong class, or, again, to a not effective cut. Applying a semi-automatic procedure for syllable selection, including errors, we were able to select 78827 syllables from the 16 classes. Some syllables (such as syllables belonging to class *O*, *N*, or *C*) are more difficult to select. Often the automatic selection based on amplitude threshold, duration of the syllable and of the gap, fails to select completely the syllable. That is, the beginning or the end of the syllable is systematically not recognized. In this case, we tried to correct manually the selection by adding a certain amount of samples after (or before) the automatic selection: in this way, more syllables can be correctly isolated.

As a preliminary solution for these errors, we applied a filter on the duration of the samples to eliminate samples that are too short (usually, we do not consider syllables shorter than $10ms$) or too long (usually, we do not consider syllables longer than $300ms$). Of course, as for the selection parameters, these threshold values could change depending on the syllable type. This procedure allows to eliminate error due to a cutting failure and resulting in samples containing more than one syllable, samples containing a very short sound (which is not always a syllable), or, occasionally, the wrong, misclassified, syllable. Using a threshold based on the duration, we could remove 6672 errors. That is, we had an error equal to $8,46\%$ after the semi-automatic selection. The clean dataset contains 72155 syllables from the 16 classes. A random set of 100 single syllables (selected from the phrases) belonging to each class of the repertoire are collected in Figures 16 and 17.

Errors due to an a priori misclassified phrase can't always be solved applying a threshold based on the syllable duration. On the one hand, if the difference between two syllables is evident, the error can be corrected simply using a threshold based on the syllable duration. This is the case of a phrase *M* wrongly classified as phrase *A*. The mean duration of syllable *A* is much smaller than the mean duration of syllable *M*: this means that we can fairly assume that a $100ms$ sample can't be syllable *A* (which has a shorter average duration). Vice versa, we can assume that a $30ms$ sample does not belong to class *M*. On the other hand, if two syllables share their duration distribution, it becomes more difficult to get rid of samples coming from a wrongly classified phrases. This is the case of a phrase *M* wrongly classified as phrase *D*: the two syllables have not only a similar duration distribution, but also a similar structure (Figure 16, panel *D*). To understand if this type of error can affect our work, we need to think about the use we want to make of the training dataset (e.g. which network we want to train and its characteristics). On the one side, we aim to use the dataset to train the generator of the GAN to produce realistic samples: from the structure of the network, we know that the generator does not care about the class of each samples (indeed, it does not know the distribution of the training data). Since the generator does not have access to the classification, it is not a problem to have well selected samples in the wrong class. On the other side, we want to use the dataset to train a classifier able to determine for each sample the class it belongs to. We need to have a clean dataset to be able to teach the classifier. Nevertheless, after a visual inspection of the samples and after a validation of the classifier using the training dataset, we can assume that this type of error represent a low percentage of the total amount of error.




Fig. 16. **Example of single syllables.** For each class of the repertoire, 100 random selected syllables from the totality of the selected syllables. To select the syllables we used three parameters: amplitude threshold, minimal duration of the syllable and duration of the gap between two consecutive syllables. From the top: A, B1, B2, C, D, E, H, J1.




Fig. 17. **Example of single syllables** For each class of the repertoire, 100 random selected syllables from the totality of the selected syllables. To select the syllables we used three parameters: amplitude threshold, minimal duration of the syllable and duration of the gap between two consecutive syllables. From the top: *J2*, *L*, *M*, *N*, *O*, *Q*, *R*, *V*.

II. APPENDIX II: WAVEGAN ARCHITECTURE

Figure 18 shows the structure of the generator G and the discriminator D , highlighting the architecture, the input and the output of the models, and the value function definition (i.e., V). The training data are pre-processed and stored in a tuple of $np.float32$ tensors representing audio waveforms (x in Figure 18). The generator model G takes as input the latent vector $z \sim U(-1, 1)$. In the original paper, z is an 100-dimensional vector, but we will use several lower dimensional inputs in this study. The upper part of Figure 18 shows the architecture of G : it has been taken from DCGAN [10] generator and it has been modified with an additional layer to obtain as output 16384 samples (i.e., 1 s of sound). Similarly, the discriminator model D takes as input vectors of length 16384 and gives as output an object of shape (16, 1024). D can take as input the training data x or the output of the generator $G(z)$, respectively resulting, after deconvolution, in $D(x)$ and $D(G(z))$. These two quantities are the variables of the value function V . To define V , Donahue et al. [20] removed the batch normalization from the generator and discriminator, and they used WGAN-GP [24] strategy during training. This strategy consists in the introduction of a gradient penalty term in the loss function, driven by the regularization hyperparameter λ .




Fig. 18. **WaveGAN architecture.** The architecture of the generator model G is the same architecture used in DCGAN [10] with an additional convolutional layer to obtain 16384 samples (i.e., 1 s of sound). The generator takes a latent vector $z \sim U(-1, 1)$ as input. The discriminator model D takes alternatively the training data x and the output of the generator $G(z)$ as input. After deconvolution, a representation of shape (16, 1024) is obtained. The outputs of the discriminator, $D(x)$ and $D(G(z))$ are used as variable of the value function V [24]. Image adapted from [10]

III. CLASSIFIER

A. Analysis of the training dataset using classifier-REAL

As for *classifier-EXT* (Figure 2a), the average number of syllables recognized for each of the 16 classes of the repertoire is close to $1k$ for. This is coherent with the fact that the training dataset contains $1k$ samples per class. Moreover, a confusion between syllable *B1* and syllable *B2*, and between syllable *J1* and syllable *J2* can be seen on the diagonal in correspondence of these elements (sub-diagonal elements are darker than the others for these 4 syllables).




Fig. 19. **Training data analysis with *classifier-REAL*.** Panel (a) shows the distribution obtained using *classifier-REAL* on the training dataset. The average number of syllables per class is $1k$ for each of the 16 classes of the repertoire. Panel (b) shows the confusion matrix (i.e. the level of confidence of the classifier in making the correct assignment) relative to *classifier-REAL*. Each square represent the level of confidence of the classifier in making the correct assignment. The level of confidence is expressed on a logarithmic scale using shades of blue.

B. Robustness of the classifier

The evaluation of the robustness of *classifier-REAL* and *classifier-EXT* has been performed using a 10 folds cross-validation over the three corresponding training datasets. For each fold, 5 different instances of each classifier were randomly initialized, trained and tested. The models were evaluated using an accuracy measure.

The accuracy has been defined as:

$$\text{accuracy}(y, \hat{y}) = \frac{1}{n_{\text{timesteps}}} \sum_{i=0}^{n_{\text{timesteps}}} \mathbb{1}(y(i) - \hat{y}(i)) \quad (4)$$

where $\mathbb{1}(x)$ is the indicator function, and considering a sequence of data x of length $n_{\text{timesteps}}$, a sequence of labels y and a sequence of classifier outputs \hat{y} , both also of length $n_{\text{timesteps}}$. The accuracy represents the capability of the classifier of making correct assignments for each time steps of MFCCs encoding the audio signal. An accuracy score close to 1 therefore indicates that the classifier is able to recognize the syllable category of each sample and to correctly assign this category to each time steps representing this sample.

Both *classifier-REAL* and *classifier-EXT* show an high level of accuracy (Figure 20). As summarized in Table II the mean accuracy for the validation set is 0.9815 ± 0.0024 for *classifier-REAL* and 0.9756 ± 0.0025 for *classifier-EXT*.

Classifier	Mean accuracy	
	Train	Validation
<i>classifier-REAL</i>	0.9832 ± 0.0005	0.9815 ± 0.0024
<i>classifier-EXT</i>	0.9777 ± 0.0006	0.9756 ± 0.0025

TABLE II

MEAN ACCURACY. COMPARISON BETWEEN *classifier-REAL* AND *classifier-EXT* IN TERMS OF THE MEDIAN OF THE ACCURACY. FOR EACH MODEL, THE ACCURACY HAS BEEN COMPUTED FOR THE TRAINING SET AND FOR THE VALIDATION SET. THE HIGHEST MEAN VALUE HAS BEEN REACHED WITH *classifier-REAL* WHERE NO ALTERNATIVE UNKNOWN CLASSES HAVE BEEN INTRODUCED.

C. Certainty of the classifier

For each syllable, the classifier described in Section II-D provides a distribution which determines to which class the classifier assign that particular syllable. As mentioned, a soft-max operation is then applied in order to obtain a distribution bounded between 0 and 1. That is, each syllable is assigned to a particular class with a probability given by the max value of the resulting N -dimensional vector, where N is the number of classes present in the vocabulary. For instance, $N = 21$ for *classifier-EXT*. Such a classifier gives high scores for the training data, for which the majority of the syllables is assigned to a class with $p_s > 0.9$ (brown points in Figure 21b). Only a few syllables are assigned to a class with a probability $p_s \leq 0.9$.

Differently, *classifier-EXT* shows a higher uncertainty while classifying the generated data (Figure 22b). Although the majority of the syllables is assigned to a class with $p_s > 0.9$ (brown points), the number of syllables assigned to a particular class with a lower probability increases. Moreover, a higher uncertainty is often connected with the presence of syllables assigned to class *X* (left panel of Figure 22b).




Fig. 20. **Accuracy of the classifiers.** Comparison between *classifier-REAL* (in the left) and *classifier-EXT* (on the right) in terms of the accuracy. For each model, the accuracy has been computed for the training set (gray rectangles) and for the validation set (white rectangles). The red lines represent the median accuracy relative to each set for each model. The white point visible for *classifier-EXT* represents an outlier, determined by the fact that the accuracy distribution is sharp. The highest accuracy has been reached with *classifier-REAL* where no alternative unknown classes have been introduced.




Fig. 21. **Certainty of the classifier: training dataset.** Panel (a) shows the spectrogram syllable space of the training data as in Figure 1. Each point corresponds to a spectrogram, and each color represents a class of the repertoire. Panel (b) shows the probability of each syllable (each point) to belong to a particular class of the repertoire. Each point corresponds to a spectrogram, and each color corresponds to an interval of probability starting at the value indicated in the legend and ending at the next color value. For example, brown points belong to a certain class of the vocabulary with a probability $0.9 < p_s \leq 1$. A higher uncertainty of *classifier-EXT* can be observed when it is applied to the generated data: indeed, an higher number of syllables are assigned to a certain class with a probability $p_s \leq 0.9$ (all the points but the brown points).




Fig. 22. **Certainty of the classifier: generated data.** Panel (a) shows the spectrogram syllable space of the generated data as in Figure 9. Each point corresponds to a spectrogram, and each color represents a class of the vocabulary (repertoire and class X). Panel (b) probability of each syllable (each point) to belong to a particular class of the repertoire (right panel), with class X highlighted in white (left panel). Each point corresponds to a spectrogram, and each color corresponds to an interval starting at the value indicated in the legend and ending at the next color value. For example, brown points belong to a certain class of the vocabulary with a probability $0.9 < p_s \leq 1$. An higher uncertainty of *classifier-EXT* can be observed when it is applied to the generated data: indeed, an higher number of syllables are assigned to a certain class with a probability $p_s \leq 0.9$ (all the points but the brown points).

IV. EXTENSION OF THE QUALITATIVE ANALYSIS

A. Generations across time

Complete version of Figure 3 generations across time for each class of the repertoire at epochs 0,15, 45, 106, 514 and 984.




Fig. 23. **Generations across time.** Example of one selected syllable per class across time. Each syllable has been first generated at epoch 984 and recognized using *classifier-EXT*. Then, the latent vector associated at each syllable has been used to generate the same syllable at epochs 0,15, 45, 106 and 514. At epoch 0 the generations do not vary from one class to another. At epoch 15 the generations start to be coherent in duration but remain noisy and unclear. At epoch 45 some syllables are more distinguishable than at earlier epochs, but the majority remains noisy and indistinguishable. As the training goes on, the generations resemble more and more the real syllables (epochs 106 and 514). The generations obtained at epoch 984 have a clear distinguishable shape. Here, the generator obtained from the instance *Ex 6* in Figure 5 has been used to generate the syllables across time; latent space dimension $ld = 3$; *Ep.* stands for epoch.

B. Mean spectrogram across time

At early epochs of training, when not all the classes are covered by the generator (empty boxes in Figure 7(a-b)), the mean spectrograms are blurry and show syllables difficult to recognize as belonging to a class of the repertoire (see Figure ??). The spectrograms look often as a mix of syllables coming from several classes (e.g, syllable *A* or syllable *R* in Figures 7(a-b)). At epoch 514 (Figure 7c) all the syllables can be produced by the generator and only a few

syllables remain difficult to be recognized as belonging to a class of the repertoire (e.g., syllable *B2* and syllable *L*). Nevertheless, a lot of syllables are clearly recognizable (e.g., syllable *N* and syllable *Q*). Finally, at advanced stages of training, all the spectrograms show a recognizable syllable (epoch 984).




Fig. 24. **Mean spectrogram across time.** Mean spectrogram of 1k the syllables generated at epoch 106 (a), 212 (b) 514 (c) and 984 (d). Empty boxes in panels (a) and (b) mean that at epoch 106 and 212 not all the repertoire can be covered by the generator and no syllables have been recognized by *classifier-EXT* as elements of the not represented classes (e.g., class *B2*). At epochs 106 (panel (a)) and 212 (panel (a)) the correct duration and, eventually, the content of the syllables can be grasped. At epoch 514 (panel (c)) almost all the syllables can be distinguished and only a few remain noisy and unclear (*A*, *E*, *L*, *R*). At epoch 984 (panel (d)) all the syllables are clear and distinguishable. They can be compared with the repertoire in Figure ?? . Here, the training has been done using $ld = 3$, the generator obtained from the instance *Ex 6* in Figure 5 has been used to generate the syllables across time and *classifier-EXT* has been used to identify the generated syllables.

C. Spectrogram syllable space obtained using UMAP

Complete version of Figure 8. The spectrogram syllable space of the generated data and the training data obtained using [1] show that (1) the generated data are grouped together as they were an additional cluster with respect to the ones obtained from the training data (upper panels of Figure 25) and (2) the generated data belong to the same cluster of the training data (bottom panels of Figure 25). On the one hand, the generated data spread over time, moving from being a cluster in itself (light blue points in the upper panels of Figure 25) to taking a conformation compatible with the training data (brown points in the upper panels of Figure 25). On the other hand, the generated data are mostly constituted by syllables belonging to class *X* (the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*) at early stages of training (bottom panels of Figure 25). Later, the majority of the generated data belongs to the same class as the closer cluster of training data in the spectrogram syllable space obtained using [1].

A similar spectrogram syllable space configuration to the one shown in Figure 25 can be obtained considering a balanced subset of ~ 2100 generated syllables (100 per class) instead of 1k random generations at epochs 15, 106, 514 and 984.

The fact that the generated syllables seem to be located at a bigger distance from the training data (Figure 8(a-c, g)) when considering a balanced dataset of 2100 samples (with respect to the same representation obtained from 1k random generated syllables shown in Figure 8(a-c, g)) can be related to the way




Fig. 25. **Spectrogram syllable space across time.** Spectrogram syllable space obtained from the training dataset ($16k$ syllables) and $1k$ syllables generated at epochs 15 (a), 106 (b), 514 (c) and 984 (d) using UMAP [1]. Upper panels show the training data (brown points) and the generated data (blue points). These four figures are different because the analyzed dataset differs for the $1k$ generations specific of each epoch. Bottom panels show the same representation with the classes of syllables highlighted with different colors. Each cluster/color corresponds to one class of the repertoire and class X (in white) represents the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Here, the training has been done using $ld = 3$, the generator of instance *Ex 6* has been used to generate the syllables and *classifier-EXT* has been used to identify both the generated data and the training data.

TABLE III

HUMAN JUDGMENT. COHEN'S KAPPA COEFFICIENT COMPUTED PER EACH COUPLE ACROSS HUMAN JUDGES AND VERSUS *classifier-EXT*. EACH JUDGE EVALUATED 200 SYLLABLES PRODUCED USING THE GENERATOR OF INSTANCE *Ex 6*. THE KAPPA COEFFICIENTS κ_{Cohen} OBTAINED ACROSS JUDGES AND WITH RESPECT TO *classifier-EXT* ARE COMPARABLE.

	Cohen's kappa
Judge 1 vs Judge 2	0,74
Judge 2 vs <i>classifier-EXT</i>	0,79
Judge 1 vs <i>classifier-EXT</i>	0,73

UMAP determines the clusters. Indeed, as much generated data are given to UMAP as much the representation obtained takes into account the difference between the real and the generated data, showing a less compact representation.

D. Human judgement

We provide a comparison between Cohen's kappa computed by the classifier and human judges. The human judges are expert in recognizing the canary syllables by relying on the spectrogram and, at a lesser extent, on the sound. Table III shows that Cohen's kappa is comparable across human judges and the classifier.

E. Stability of the training

The spectrogram syllable space obtained using [1] (Figure 27) and the mean spectrograms (Figure 28) show a similar representation for consecutive epochs (i.e., epochs 969, 984 and 999) which show the stability of instance *Ex 6* around the convergence of the training. The differences in the spectrogram syllable space obtained using [1] (Figure 27) are given by the fact that the set of generated data is different at each epoch.




Fig. 26. **Spectrogram syllable space across time: balanced representation.** Spectrogram syllable space obtained from the training dataset (16k syllables) and 2100 syllables (100 per class, when the class is present, for a total of 21 classes - the 16 classes of the repertoire and 5 alternative unknown classes, here grouped as class X) generated at epochs 15 (a), 106 (b), 514 (c) and 984 (d) using UMAP [1]. Upper panels show the training data (brown points) and the generated data (blue points). These four figures are different because the analyzed dataset differs for the 2100 generations specific of each epoch. Bottom panels show the same representation of panels with the classes of syllables highlighted with different colors. Each cluster/color corresponds to one class of the repertoire and class X (in white) represents the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Here, the training has been done using $ld = 3$, the generator of instance *Ex 6* has been used to generate the syllables and *classifier-EXT* has been used to identify both the generated data and the training data.

F. Latent space exploration

Figures 30 and 31 shows a similar analysis as the one in Figure 11. Here, the second and the third component of the latent space have been varied using a variational step equal to $v_{step} = 0.001$ and a syllable has been generated at each step. The spectrogram syllable space of the newly generated syllables and the generated dataset shown in Figure 9, obtained using [1], shows a continuity in the productions of the generator (light blue lines in the left panels of Figure 32(a-b)).

Moreover, the continuity of the latent space can be observed when observing the transition from a syllable to another, as shown in Figure 12b. Figure 33 shows two more example of transition are shown: (i) from syllable M to syllable V and (ii) from syllable H to syllable N . The transition from syllable H to syllable N shows (1) an interesting stretch of syllable B (cream cluster in the right panels of Figure 33) connecting syllable B and syllable Q and (2) the uncertainty of the classifier in differentiating between syllable H (light green cluster in the right panels of Figure 33) and syllable A (gray cluster in the right panels of Figure 33).




Fig. 27. **Stability of the spectrogram syllable space.** Spectrogram syllable space of the training dataset ($16k$ syllables) and $1k$ generated data at epochs 969 (a), 984 (b) and 999 (d). This representation has been obtained using [1]. The upper panels show the representation of the training data (brown points) and the generated data (light blue points) over time. The bottom panels show the same representation highlighting the classes of the vocabulary. Each cluster/color correspond to one class of the repertoire and class X (in white) represents the cumulative class of alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Although the representation is slightly different (given the fact that the generated syllables vary across time, it is possible to observe how the generated syllable mix well with the clusters obtained from the training data at all epochs), it remains stable across consecutive epochs. Here, the training has been done using $ld = 3$, the generator of instance *Ex 6* has been used to generate the syllables and *classifier-EXT* has been used to identify both the generated data and the training data.




Fig. 28. **Stability of the mean spectrogram.** Mean spectrogram of $1k$ syllables generated at epochs 969 (a), 984 (b) and 999 (c). The three epochs share a similar representation of the syllables. Here, the training has been done using $ld = 3$, the generator of instance *Ex 6* has been used to generate the syllables and *classifier-EXT* has been used to identify the generated data.




Fig. 29. **Exploration of the latent space: one component variation. First component.** We selected a random latent vector $z \sim R^3([-1, 1])$ to create a baseline syllable. Then, we moved one by one the components of the vector by a variation step equal to $v_{step} = 0.05$. We observed all the syllables produced to look at how they evolve and if there are non-smooth transitions. The upper panel of Figure 11 shows the exploration of the first component of the latent vector obtained with $v_{step} = 0.05$. The syllable V contained in the red square on the right represent a point where a non-smooth transition has been detected. For these particular transitions, we considered the two consecutive syllables obtained from the first variation step (i.e., two consecutive syllables contained in two consecutive red squares) and we applied a variation step of $v_{step} = 0.01$ to the first component of the latent vector to generate intermediate syllables. The bottom panel show the exploration of the non-smooth transition highlighted above. The syllables have been obtained the 3-dimensional generator obtained from instance *Ex 6* and the name of the syllable for this analysis has been provided by *classifier-EXT*.




Fig. 30. **Exploration of the latent space: one component variation. Second component.** We selected a random latent vector $z \sim R^3([-1, 1])$ to create a baseline syllable. Then, we moved one by one the components of the vector by a variation step equal to $v_{step} = 0.05$. We observed all the syllables produced to look at how they evolve and if there are non-smooth transitions. The upper panel shows the exploration of the second component of the latent vector obtained with $v_{step} = 0.05$. The syllable V contained in the red square on the right represent a point where a non-smooth transition has been detected. For these particular transitions, we considered the two consecutive syllables obtained from the first variation step (i.e., two consecutive syllables contained in two consecutive red squares) and we applied a variation step of $v_{step} = 0.01$ to the first component of the latent vector to generate intermediate syllables. The bottom panel show the exploration of the non-smooth transition highlighted above. The syllables have been obtained the 3-dimensional generator obtained from instance *Ex 6* and the name of the syllable for this analysis has been provided by *classifier-EXT*.





Fig. 31. **Exploration of the latent space: one component variation. Third component.** We selected a random latent vector $z \sim R^3([-1, 1])$ to create a baseline syllable. Then, we moved one by one the components of the vector by a variation step equal to $v_{step} = 0.05$. We observed all the syllables produced to look at how they evolve and if there are non-smooth transitions. The upper panel shows the exploration of the third component of the latent vector obtained with $v_{step} = 0.05$. The syllable V contained in the red square on the right represent a point where a non-smooth transition has been detected. For these particular transitions, we considered the two consecutive syllables obtained from the first variation step (i.e., two consecutive syllables contained in two consecutive red squares) and we applied a variation step of $v_{step} = 0.01$ to the first component of the latent vector to generate intermediate syllables. The bottom panel show the exploration of the non-smooth transition highlighted above. The syllables have been obtained the 3-dimensional generator obtained from instance *Ex 6* and the name of the syllable for this analysis has been provided by *classifier-EXT*.



(a) Variation in the second component



(c) Variation in the third component

Fig. 32. **Latent vector components variation.** Panel (a) represents the first component variation, panel (b) represents the second component variation and panel (c) represents the third component variation. The left panels show the variational data (generated at each step) (light blue points) and $16k$ generated data from the same model at the same epoch (brown points). Each cluster/color in the right panels corresponds to one class of the repertoire and class X (in white) represents the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Here, the training has been done using $ld = 3$. A latent vector has been randomly selected and a step $v_{step} = 0.001$ has been applied to its component, one by one. The generator obtained from instance *Ex 6* of the training has been used *classifier-EXT* has been used to identify the generated syllables from the generator of instance *Ex 6*.




Fig. 33. **Transition between two generated syllables.** Panel (a) represents the transition between M (turquoise cluster) and syllable D (yellow cluster), panel (b) the transition between M (turquoise cluster) and syllable V (magenta cluster) and panel (c) the transition between H (light green cluster) and syllable N (blue cluster). The left panels show the variational data (generated at each step) (light blue points) and $16k$ generated data from the same model at the same epoch (brown points). Each cluster/color in the right panels corresponds to one class of the repertoire and class X (in white) represents the cumulative class of the alternative unknown classes (in this case, *EARLY15*, *EARLY30*, *EARLY45*, *OT* and *WN*). Here, the training has been done using $ld = 3$ and *classifier-EXT* has been used to identify the generated syllables from the generator of instance *Ex 6*.

V. EXTENSION OF THE QUANTITATIVE ANALYSIS

A. Different instances of training of a 3-dimensional WaveGAN

The analysis of 10 instances of training of a 3-dimensional WaveGAN shows that not all the instances have the capability of providing a generator able of producing syllables belonging to all the classes of the repertoire (Figure 34a). Instances *Ex 2* (orange line), *Ex 1* (yellow line) and *Ex 5* (magenta line) show an early drop and, eventually, never reach to cover the repertoire. Nevertheless, the other instances show a drop at an advanced stage of the training (i.e., after epoch 600) or never drop. Similarly, the instances showing instability in Figure 34a show instability in the average number of syllables recognized per class (Figure 34b) and in the variance of the number of syllables recognized per class (Figure 34c). Moreover, alternative unknown classes, and in particular class *EARLY45*, are more represented even in advanced epochs of the training for those instances showing instability, as shown in Figures 34(d-f). Such an instability might characterize the beginning of overtraining: the generator starts to produce samples that are recognized as elements of an *EARLY* class or of class *OT*. The latter is more represented in instances showing overtraining (Figure 35). After, the generator is not able to recover. Further analysis are needed to understand how to carefully describe overtraining.




Fig. 34. **Different instances of a 3-dimensional WaveGAN.** Statistical analysis performed on the classifier distribution of 10 instances of training. Panel (a) shows the number of classes represented by the generated data: that is, at each epoch, how many syllables of the repertoire are covered by the generator. An early drop (i.e., before epoch 600) in the number of represented classes can be seen for three instances (i.e., *Ex 2* – orange line, *Ex 1* – yellow line and *Ex 5* – magenta line). Panel (b) shows the average number (dark colored lines) of elements per class and the median (light colored lines): depending on the instance, the mean and the median could increase over time and remain stable (*Ex 6* – light blue line, *Ex 7* – gray line, *Ex. 8* – burgundy line), or increase until a certain epoch and then drop (*Ex 3*–blue line, *Ex 4* – green line, *Ex 5* – magenta lines, *Ex 9*–red line), or remain low for all the duration of the training (*Ex 1* – yellow line, *Ex 2* – orange line). To compute the quantities in panels (a) and (b), alternative unknown classes $x \in X$ have not been taken into account. Panel (c) shows the evolution of the variance of how many syllable per class have been produced. Here, $x \in X$ classes are included. For successful instances of training (*Ex 6* – light blue line, *Ex 7* – gray line, *Ex. 8* – burgundy line), the variance starts at a high value when the majority of the samples produced are not classified as syllables of the repertoire, then it decreases when the generator becomes better at producing syllables. Eventually, it increases again later (*Ex 3* – blue line, *Ex 4* – green line, *Ex 5* – magenta line, *Ex 9* – red line) or never decrease enough (*Ex 1* – yellow line, *Ex 2* – orange line). Panels (d-f) show the percentage of syllables that are classified as belonging to one of the alternative classes $x \in X$: from the left to the right, classes *EARLY15*, *EARLY30* and *EARLY45*. Alternative unknown classes, and in particular class *EARLY45*, are more present even in advanced epochs of the training for those instances where we identify overtraining (all but *Ex. 6* – light blue line, *Ex 7* – gray line and *Ex. 8* – burgundy line).

B. Analysis of the parameters: complete version of the figures

We used the training dataset described in Section II-A to train WaveGAN. Then, we used the experimental setup described in Section II-C. We used the *classifier-EXT* and *classifier-REAL* to identify the syllables as elements of the repertoire and 5 alternative unknown classes $x \in X$ (*EARLY15*, *EARLY30*, *EARLY45*, *OT*, and *WN*).

1) *Latent space dimension*: Complete version of Figures 13 (Figure 36). Accordingly to what observed in the main paper for Figure 36(a-c) and similarly to what shown in Figure 34, the percentage of alternative unknown syllables (here, *EARLY15*, *EARLY30* and *EARLY45*) decreases over time (Figure 36(d-f)).




Fig. 35. **Overtraining.** Comparison between 10 instances of training: percentage of syllables classified by *classifier-EXT* as elements of class *OT* (i.e., overtraining). Instances *Ex 0* (purple line), *Ex 1* (yellow line), *Ex 2* (orange line), *Ex 3* (blue line), *Ex 4* (green line), *Ex 5* (magenta line), *Ex 9* (red line) show an increasing number of syllables classified as belonging to class *OT*. These instances show instability also in the average and variance of the syllables belonging to the classes of the repertoire and in the number of syllables belonging to an *EARLY* class (see Figure 34). Latent space dimension: $ld = 3$.




Fig. 36. **Comparison between different latent space dimension: quantitative measure.** Each line represents one instance of training at a particular latent space dimension. Panel (a) shows how many syllables of the repertoire are covered by the generator across time. Panel (b) shows how many syllable per class have been generated in average. The dark lines show the evolution of the mean, whereas the light lines show the evolution of the median. To build these two panels (a) and (b), $1k$ generations have been generated every 15 epochs, *classifier-EXT* has been used to provide the classification, and only the repertoire's classes have been taken into account when plotting. Panel (c) shows the evolution of the inception score over time. To build this panel $16k$ syllables have been generated every 200, and *classifier-REAL* has been used to provide the classification. Panels (d-f) show the percentage of generated syllables belonging to classes *EARLY15*, *EARLY30* and *EARLY45* across time in comparison with the percentage of syllables belonging to the same class in the training data. A 3-dimensional WaveGAN (orange line) reaches convergence as good as higher-dimensional WaveGANs (blue, green and magenta lines) and better than lower dimensional WaveGANs (purple and yellow lines). We varied the latent space dimension as $ld = 1, 2, 3, 4, 5, 6$ and we kept fix the training dataset.

The fact that the performance obtained for $ld = 3$ is comparable with the performance obtained for $3 > ld \leq 6$ and better than the performance obtained for $1 \leq ld \leq 2$ (Figure 36) is also confirmed by a better mean spectrogram representation 37. A good epoch of training is determined by looking at the classifier distribution (shown on the top of each spectrogram in Figure 37). The mean spectrograms obtained for $ld = 1$ (Figure 37a) and $ld = 2$ (Figure 37b) show noisy representations for several syllables. For instance, but not restricted to, syllables *C*, *H*, *O*. Nevertheless, some syllable representation are influenced by the fact that the trainings used to generated the samples have been done using the preliminary dataset described in Appendix I.




Fig. 37. **Comparison between different latent space dimension: qualitative measure.** Mean spectrogram of $1k$ syllables generated at a good epoch of training determined by looking at the classifier distribution (on top of each spectrogram) for $ld = 1, 2, 3, 4, 5, 6$. Whereas the representations obtained for $ld = 1$ (a) and $ld = 2$ (b) show a noisy representation for several syllables (e.g., syllable M in (b)), $ld = 3$ (c) shows a representation comparable to the one obtained for higher conditions (d-f).

2) *Dataset size*: Complete version of Figure 14 (Figure 38). Accordingly to what observed in the main paper for Figure 36(a-c) and similarly to what shown in Figure 34 and Figure 36, the percentage of alternative unknown syllables (here, *EARLY15*, *EARLY30* and *EARLY45*) decreases over time, and eventually increases when overtraining begins (e.g., the blue line around epoch 900 in Figure 36e).




Fig. 38. **Comparison between datasets of a different size.** We varied the dataset size as $d = 16000, 8000, 4000$ and we kept fix the latent space dimension at $ld = 3$. Each line in this figure represents one instance of training at a particular dataset size condition. Panel (a) shows how many syllables of the repertoire are covered by the generator across time. Panel (b) shows how many syllable per class have been generated in average. The dark lines show the evolution of the mean, whereas the light lines shows the evolution of the median. To build these two panels (a) and (b), $1k$ generations have been generated every 15 epochs, *classifier-EXT* has been used to provide the classification, and only the repertoire's classes have been taken into account when plotting. Panel (c) shows the evolution of the inception score over time. To build this panel $16k$ syllables have been generated every 200, and *classifier-REAL* has been used to provide the classification. Panels (d-f) show the percentage of generated syllables belonging to the alternative classes *EARLY15*, *EARLY30* and *EARLY45* across time in comparison with the percentage of syllables belonging to each garbage class in the training data. A dataset of bigger size (blue line) allows better and faster convergence than having a dataset of lower sizes (red and green lines).

Finally, Figure 39 show the comparison between the mean spectrogram obtained after training WaveGAN with the whole dataset of 16k syllables (a), and the same representation obtained when WaveGAN was trained using half the amount of data (b), and a quarter the amount of data (c). The representations obtained when using a smaller dataset show good results for the majority of the syllables, but introduces noise for several syllables (*A*, *E*, *L*, and *R*).




Fig. 39. **Comparison between different dataset size: qualitative measure.** Mean spectrogram of $1k$ syllables generated at a good epoch of training determined by looking at the classifier distribution (on top of each spectrogram) for dataset of dimension 16k (a), 8k (b) and 4k (c). With respect to the representations obtained by training the model with the whole dataset (a), the representations obtained for (b) and (c) show a noisy representation for several syllables (e.g., syllables A, E, L, and R).