



HAL
open science

Audio-visual speech enhancement with a deep kalman filter generative model

Ali Golmakani, Mostafa Sadeghi, Romain Serizel

► **To cite this version:**

Ali Golmakani, Mostafa Sadeghi, Romain Serizel. Audio-visual speech enhancement with a deep kalman filter generative model. International Conference on Acoustics Speech and Signal Processing (ICASSP), IEEE, Jun 2023, Rhodes island, Greece. hal-03833814v2

HAL Id: hal-03833814

<https://inria.hal.science/hal-03833814v2>

Submitted on 4 Jan 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

AUDIO-VISUAL SPEECH ENHANCEMENT WITH A DEEP KALMAN FILTER GENERATIVE MODEL

Ali Golmakani, Mostafa Sadeghi, and Romain Serizel

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

ABSTRACT

Deep latent variable generative models based on variational autoencoder (VAE) have shown promising performance for audio-visual speech enhancement (AVSE). The underlying idea is to learn a VAE-based audio-visual prior distribution for clean speech data, and then combine it with a statistical noise model to recover a speech signal from a noisy audio recording and video (lip images) of the target speaker. Existing generative models developed for AVSE do not take into account the sequential nature of speech data, which prevents them from fully incorporating the power of visual data. In this paper, we present an audio-visual deep Kalman filter (AV-DKF) generative model which assumes a first-order Markov chain model for the latent variables and effectively fuses audio-visual data. Moreover, we develop an efficient inference methodology to estimate speech signals at test time. We conduct a set of experiments to compare different variants of generative models for speech enhancement. The results demonstrate the superiority of the AV-DKF model compared with both its audio-only version and the non-sequential audio-only and audio-visual VAE-based models.

Index Terms— Audio-visual speech enhancement, generative model, variational autoencoder, deep Kalman filter.

1. INTRODUCTION

Audio-visual speech enhancement (AVSE) is the task of estimating a clean speech signal given a noisy audio recording, as well as visual information (e.g., lip images) of the speaker [1]. Visual data provide complementary information that could be very helpful for speech enhancement, especially when the audio recording is highly noisy [1, 2]. Furthermore, visual data are robust with respect to acoustic noise and could help discriminate between the target speaker and potential concurrent speakers. Over the last decade and with the unprecedented progress made in deep learning, the AVSE problem has been extensively revisited [1, 3–5].

A dominant AVSE approach is to design and train a deep neural architecture that fuses audio and visual features, extracted from video and noisy audio data, respectively, to estimate the clean speech signal directly. This approach is data-driven and, as such, its success and generalization performance depend heavily on the amount of training data and their diversity, e.g., in terms of noise types. In contrast to this *supervised* AVSE framework, a recent alternative paradigm is to combine the classical model-based methods, e.g., maximum a posteriori (MAP) estimation, with the expressive power of deep neural networks (DNNs) to perform *unsupervised* AVSE [6–9]. More precisely, in a pre-training phase, the statistical

characteristics of speech signals in the time-frequency domain are learned via a deep generative model based on variational autoencoders (VAEs) [10], with only *clean* audio-visual (AV) data. The learned speech generative model, serving as a deep speech prior, is then combined with a parametric statistical model for noise, whose parameters along with the clean speech signal are estimated following an expectation-maximization (EM)-based approach. As noise is modeled at test time, unsupervised AVSE can adapt to unseen noise situations and has a potentially better generalization performance than its supervised counterpart [6].

The AV-VAE models developed so far for unsupervised AVSE do not account for the sequential nature of speech data, as they rely on a statistical independence assumption between consecutive speech time frames, and thus ignore their intrinsic correlations. Recently, some dynamical variants of VAEs, called DVAEs [11], have been used for audio-only speech enhancement [12], which effectively model the temporal dynamics of speech data and improve the enhancement performance with respect to VAEs. Nevertheless, this comes at the cost of complicating the EM step at test time, due to the complex temporal dependencies of latent variables in the models.

In this paper, we extend the unsupervised AVSE framework to audio-visual DVAE models, with a focus on respecting the computational efficiency of original non-sequential models for speech enhancement. To this end, we develop an audio-visual extension of the deep Kalman filter (DKF) model [13], as the simplest DVAE variant in terms of temporal dependencies and architecture, which assumes a first-order Markov model on the latent variables. The proposed AV-DKF model efficiently incorporates visual data for clean speech generative modeling. Furthermore, we propose a dedicated EM-based methodology for parameter learning and speech estimation at test time. Our experimental results demonstrate the superiority of the developed AV-DKF framework for speech enhancement compared with both its audio-only variant considered by Bie *et al.* [12], and the standard AV-VAE model proposed in [6].

The rest of the paper is organized as follows. Section 2 reviews speech generative modeling and enhancement based on standard and dynamical VAEs. The proposed speech generative modeling and enhancement frameworks are detailed in Section 3. Experimental results are then presented in Section 4, followed by the conclusions in Section 5.

2. BACKGROUND

Let us denote the short-time Fourier transform (STFT) representations of clean speech signals as $\mathbf{s}_{1:T} = \{\mathbf{s}_t\}_{t=1}^T$, where $\mathbf{s}_t = [s_{ft}]_{f=1}^F \in \mathbb{C}^F$. A latent variable $\mathbf{z}_t \in \mathbb{R}^L$ ($L \ll F$) is associated to each STFT time frame \mathbf{s}_t . The VAE framework then involves modeling the joint distribution of the observed and latent variables with some parametric Gaussian forms. The standard VAE model assumes the following factorization $p_\theta(\mathbf{s}_{1:T}, \mathbf{z}_{1:T}) =$

Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER, and several Universities as well as other organizations (see <https://www.grid5000.fr>).




Fig. 1: Schematic diagram of the proposed AV-DKF generative model (without explicit architecture of the prior network). **MLP**: multilayer perception, **RNN**: recurrent neural network, **VE**: video encoder, \oplus : addition, **C**: concatenation, **S**: sampling in the latent space.

$\prod_{t=1}^T p_\theta(\mathbf{s}_t, \mathbf{z}_t) = \prod_{t=1}^T p_\theta(\mathbf{s}_t | \mathbf{z}_t) p_\theta(\mathbf{z}_t)$. This modeling framework has been extended to the audio-visual case by conditioning the two distributions on visual features [6]. No temporal modeling is considered here, which is not realistic for speech STFT time frames. To resolve this issue, DVAEs consider the following factorization: $p_\theta(\mathbf{s}_{1:T}, \mathbf{z}_{1:T}) = \prod_{t=1}^T p_\theta(\mathbf{s}_t | \mathbf{s}_{1:t-1}, \mathbf{z}_{1:t}) p_\theta(\mathbf{z}_t | \mathbf{s}_{1:t-1}, \mathbf{z}_{1:t-1})$. The two distributions involved in this factorization are parameterized by some deep neural network (DNN) architectures, known as the *decoder* and *prior* networks, respectively.

Parameter inference, i.e., learning θ , necessitates computation of the posterior distribution $p_\theta(\mathbf{z}_{1:T} | \mathbf{s}_{1:T})$, which is highly intractable due to the non-linear generative model. As a solution, a variational approximation is employed where a Gaussian form parameterized by a DNN, called the *encoder*, is introduced to approximate the intractable posterior [10]. For DVAEs, this writes $q_\psi(\mathbf{z}_{1:T} | \mathbf{s}_{1:T}) = \prod_{t=1}^T q_\psi(\mathbf{z}_t | \mathbf{s}_{1:T}, \mathbf{z}_{1:t-1})$. The inference phase involves joint learning of the model parameters by optimizing the so-called evidence lower bound (ELBO) of the intractable data log-likelihood $\log p_\theta(\mathbf{s}_{1:T})$ with stochastic gradient-based algorithms [10, 12]. Depending on how the observed and latent variables are structured in the decoder and prior, several variants of the DVAE arise. In particular, DKF, as the simplest DVAE variant, assumes the following joint factorization $p_\theta(\mathbf{s}_{1:T}, \mathbf{z}_{1:T}) = \prod_{t=1}^T p_\theta(\mathbf{s}_t | \mathbf{z}_t) p_\theta(\mathbf{z}_t | \mathbf{z}_{t-1})$, i.e., with a first-order Markov model on the latent variables.

The speech enhancement phase consists in combining the pre-trained speech generative model (the learned speech prior) with a parametric Gaussian model for noise, usually based on a non-negative matrix factorization (NMF) variance model [14]. The NMF parameters are then learned from the observed noisy STFT time frames, followed by speech signal estimation based on Wiener filtering. Here, one would also need to compute the posterior distribution of latent variables, which is intractable. A variational EM approach is proposed in [12] that fine-tunes the pre-trained clean encoder on the noisy observations to approximate the intractable posterior.

3. PROPOSED FRAMEWORK

3.1. Audio-visual DKF Generative model

We follow the DKF generative model [11], and propose to extend it to the audio-visual case. Given some clean AV training data $\mathbf{u}_{1:T} = \{\mathbf{s}_t, \mathbf{v}_t\}_{t=1}^T$, with \mathbf{v}_t being the visual feature vector at time frame t extracted using a video encoder, the generative model is defined as

follows:

$$\begin{cases} p_\theta(\mathbf{s}_t | \mathbf{z}_t, \mathbf{v}_t) = \mathcal{N}_c(\mathbf{0}, \text{diag}(\sigma_{\theta_s}^2(\mathbf{z}_t, \mathbf{v}_t))), \\ p_\theta(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{v}_t) = \mathcal{N}(\boldsymbol{\mu}_{\theta_z}(\mathbf{z}_{t-1}, \mathbf{v}_t), \text{diag}(\sigma_{\theta_z}^2(\mathbf{z}_{t-1}, \mathbf{v}_t))), \end{cases}$$

where $\mathcal{N}_c(\mathbf{0}, \boldsymbol{\Sigma})$ denotes a circularly symmetric complex Gaussian distribution, σ_{θ_s} , $\boldsymbol{\mu}_{\theta_z}$, $\sigma_{\theta_z}^2$ are parametric non-linear functions realized by some DNN architectures, and $\theta = \{\theta_s, \theta_z\}$. The approximate posterior takes the following form:

$$\begin{aligned} q_\psi(\mathbf{z}_{1:T} | \mathbf{u}_{1:T}) &= \prod_{t=1}^T q_\psi(\mathbf{z}_t | \mathbf{r}_t) \\ &= \prod_{t=1}^T \mathcal{N}(\boldsymbol{\mu}_\psi(\mathbf{r}_t), \text{diag}(\sigma_\psi^2(\mathbf{r}_t))) \end{aligned} \quad (1)$$

where $\mathbf{r}_t = \{\mathbf{z}_{t-1}, \mathbf{u}_{t:T}\}$ collects all the conditioning variables, and $\boldsymbol{\mu}_\psi$, σ_ψ^2 are DNN-parameterized non-linear functions, i.e., the encoder. As with the previous works, the encoder takes the modulus square of STFT data as input. Learning the set of parameters, i.e., $\Phi = \{\theta, \psi\}$, amounts to optimizing the ELBO:

$$\begin{aligned} \mathcal{L}(\Phi; \mathbf{u}_{1:T}) &= \sum_{t=1}^T \mathbb{E}_{q_\psi(\mathbf{z}_t | \mathbf{u}_{1:T})} \{\log p_\theta(\mathbf{s}_t | \mathbf{z}_t, \mathbf{v}_t)\} - \\ &\quad \sum_{t=1}^T \mathbb{E}_{q_\psi(\mathbf{z}_{t-1} | \mathbf{u}_{1:T})} \{\mathcal{D}_{\text{KL}}(q_\psi(\mathbf{z}_t | \mathbf{r}_t) \| p_\theta(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{v}_t))\}, \end{aligned} \quad (2)$$

where $\mathcal{D}_{\text{KL}}(q \| p)$ denotes the Kullback–Leibler (KL) divergence between q and p . The two expectations can be computed recursively, as detailed in [11]. A single-sample Monte-Carlo approximation of the two expectations is computed followed by the *reparametrization trick* [10] before optimizing the parameters. The proposed AV-DKF architecture is shown in Fig. 1.

3.2. Speech Enhancement

The observed noisy speech data are modeled as $\mathbf{x}_t = \sqrt{g_t} \mathbf{s}_t + \mathbf{b}_t$, $t = 1, \dots, \hat{T}$, where \mathbf{b}_t corresponds to noise. The parameters $\mathbf{g}_{1:\hat{T}} = \{g_t\}_{t=1}^{\hat{T}}$ are non-negative scalars to take into account the potentially different loudness between training and test speech data [15]. As the statistical model of \mathbf{s}_t , i.e., the prior distribution, the pre-trained AV-DKF generative model in (1) is used. Following the NMF

approach, two matrices \mathbf{W} , \mathbf{H} of dimensions $F \times K$ and $K \times \tilde{T}$, respectively, with non-negative entries are considered for the variance of \mathbf{b}_t as follows:

$$\mathbf{b}_t \sim \mathcal{N}(\mathbf{0}, \text{diag}(\mathbf{W}\mathbf{h}_t)), \quad (3)$$

where \mathbf{h}_t is the t -th column of \mathbf{H} . As opposed to the previous works [12, 15] that treat $\mathbf{g}_{1:\tilde{T}}$ as model parameters and estimate them using multiplicative update rules, here we propose a probabilistic modeling framework by assuming a gamma prior distribution for each g_t as follows:

$$p(g_t) = \frac{\beta^\alpha}{\Gamma(\alpha)} g_t^{\alpha-1} \exp(-\beta g_t), \quad (4)$$

where $\Gamma(\cdot)$ is the gamma function, and $\alpha, \beta > 0$ (set to some predefined values), are the shape and scale parameters, respectively. As will be shown in Section 4, this new approach results in significantly more stable and improved performance, especially for DKF-based models. Given the observed data $\mathbf{o}_{1:\tilde{T}} = \{\mathbf{x}_t, \mathbf{v}_t\}_{t=1}^{\tilde{T}}$, we follow an EM approach to estimate the set of model parameters $\phi = \{\mathbf{W}, \mathbf{H}\}$, which involves optimizing the expectation of the complete data log-likelihood with respect to the following posterior:

$$p_\phi(\mathbf{z}_{1:\tilde{T}}, \mathbf{g}_{1:\tilde{T}} | \mathbf{o}_{1:\tilde{T}}) \propto p_\phi(\mathbf{x}_{1:\tilde{T}} | \mathbf{z}_{1:\tilde{T}}, \mathbf{g}_{1:\tilde{T}}, \mathbf{v}_{1:\tilde{T}}) p_\theta(\mathbf{z}_{1:\tilde{T}} | \mathbf{v}_{1:\tilde{T}}) p(\mathbf{g}_{1:\tilde{T}}), \quad (5)$$

where the likelihood writes:

$$p_\phi(\mathbf{x}_{1:\tilde{T}} | \mathbf{z}_{1:\tilde{T}}, \mathbf{g}_{1:\tilde{T}}, \mathbf{v}_{1:\tilde{T}}) = \prod_{t=1}^{\tilde{T}} p_\phi(\mathbf{x}_t | \mathbf{z}_t, g_t, \mathbf{v}_t) = \prod_{t=1}^{\tilde{T}} \mathcal{N}(\mathbf{0}, \text{diag}(g_t \sigma_{\theta_s}^2(\mathbf{z}_t, \mathbf{v}_t) + \mathbf{W}\mathbf{h}_t)). \quad (6)$$

Unfortunately, there is no closed form expression for (5). However, as an efficient approximate approach inspired by [16, 17], we try to find the mode of the posterior distribution of $\mathbf{z}_{1:\tilde{T}}, \mathbf{g}_{1:\tilde{T}}$:

$$\mathbf{z}_{1:\tilde{T}}^*, \mathbf{g}_{1:\tilde{T}}^* = \underset{\mathbf{z}_{1:\tilde{T}}, \mathbf{g}_{1:\tilde{T}}}{\text{argmax}} \sum_{t=1}^{\tilde{T}} \log p_\phi(\mathbf{x}_t | \mathbf{z}_t, g_t, \mathbf{v}_t) + \log p_\theta(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{v}_t) + \log p(g_t), \quad (7)$$

which, after substituting from (1), (4), and (6), can be optimized by a few iterations of a gradient-based solver, e.g., Adam [18]. In the maximization (M) step, the NMF parameters are updated according to the following approximate problem:

$$\mathbf{W}, \mathbf{H} \leftarrow \underset{\mathbf{W}, \mathbf{H}}{\text{argmax}} \sum_{t=1}^{\tilde{T}} \log p_\phi(\mathbf{x}_t | \mathbf{z}_t^*, g_t^*, \mathbf{v}_t), \quad (8)$$

which can be solved with multiplicative update rules as similarly done in [15]. The overall inference algorithm iterates between (7) and (8). Once $\phi^* = \{\mathbf{W}^*, \mathbf{H}^*\}$ is learned, the speech signal is estimated as the posterior mean (element-wise division):

$$\begin{aligned} \hat{\mathbf{s}}_{1:\tilde{T}} &= \mathbb{E}_{p_{\phi^*}(\mathbf{s}_{1:\tilde{T}} | \mathbf{o}_{1:\tilde{T}})} \{\mathbf{s}_{1:\tilde{T}}\} \\ &= \mathbb{E}_{p_{\phi^*}(\mathbf{z}_{1:\tilde{T}}^*, \mathbf{g}_{1:\tilde{T}}^* | \mathbf{o}_{1:\tilde{T}})} \left\{ \mathbb{E}_{p_{\phi^*}(\mathbf{s}_{1:\tilde{T}} | \mathbf{z}_{1:\tilde{T}}^*, \mathbf{g}_{1:\tilde{T}}^*, \mathbf{o}_{1:\tilde{T}})} \{\mathbf{s}_{1:\tilde{T}}\} \right\} \\ &\approx \left\{ \frac{g_t^* \sigma_{\theta}^2(\mathbf{z}_t^*, \mathbf{v}_t)}{g_t^* \sigma_{\theta}^2(\mathbf{z}_t^*, \mathbf{v}_t) + \mathbf{W}^* \mathbf{h}_t^*} \odot \mathbf{x}_t \right\}_{t=1}^{\tilde{T}}. \end{aligned} \quad (9)$$

4. EXPERIMENTS

In this section, we provide a performance evaluation of our proposed AV-DKF speech enhancement framework against some baseline methods, including A-VAE [15], AV-VAE [6], and A-DKF [12]. To measure the quality of the enhanced speech signals, we use standard metrics, including the scale-invariant signal-to-distortion ratio (SI-SDR) in dB [19], the short-term objective intelligibility (STOI) measure [20], ranging in $[0, 1]$, and the perceptual evaluation of speech quality (PESQ) score [21], ranging in $[-0.5, 4.5]$. For all the metrics, the higher, the better.

Datasets. For training all the VAE variants, we used the TCD-TIMIT corpus [22]. This dataset contains AV speech data from 56 English speakers (39 for training, 8 for validation, and 9 for testing) with an Irish accent, uttering 98 different sentences, each ~ 5 -second long, and sampled at 16 kHz. This amounts to ~ 8 hours of data. There is a corresponding video file for each utterance that captures a frontal view of the speaker at a rate of 30 frames per second. For all the videos, the lip region of interest (ROI) are already extracted as 67×67 images (a sample is shown in Fig. 1). The STFT of the speech data is computed with a 1024 samples-long (64 ms) sine window, 75% overlap, without zero-padding, yielding STFT frames of length $F = 513$. The ROI images for each video were upsampled across time to make visual and audio frame rates equal. Moreover, we augmented each video into 75 different samples by applying natural image transformations such as random translation up to 10%, random scaling and crop up to 10%, and random brightness and contrast jitter up to 40%.

To test the speech enhancement performance, we used the noisy speech material of the NTCD-TIMIT corpus [23] which has been created by adding six noise types, including *Living Room (LR)*, *White, Cafe, Car, Babble*, and *Street*, with different signal-to-noise (SNR) ratios to the test speech data of the TCD-TIMIT corpus. We randomly selected 5 utterances per noise level and noise type from each test speaker, which resulted in 1350 test utterances.

Models architectures. The A-VAE and AV-VAE models share the same architectures as the baseline VAE model experimented in [12, 17], where the encoder and decoder comprise a single fully connected (FC) hidden layer with 128 nodes and tanh activation functions. For the DKF models, we followed a similar architecture as the one proposed in [12], which consists of a backward long short-term memory (LSTM) network, a combiner function in the encoder, and a gated transition function in the prior network [13]. The decoder comprises a multilayer perceptron (MLP) with four hidden layers of dimensions 32, 64, 128, and 256, with the tanh activation functions. We took the pre-trained A-VAE and A-DKF models of [12], trained on the (audio-only) Wall Street Journal (WSJ0) corpus [24], and fine-tuned them on our training data.¹ Similarly, we fine-tuned AV-VAE and AV-DKF from their audio-only pre-trained counterparts [12].

For the AV models, the raw visual data are processed by a feature extraction network, called VE (video encoder) in Fig. 1, before feeding to the model. For that, we incorporated a pre-trained model that is part of the lipreading network proposed in [25]. We took the initial blocks of the visual network, which include a 3D convolutional module and a ResNet architecture module. We used the computed features as the visual input of our model, without fine-tuning the VE network during the training process. The introduced ‘‘skip connection’’ in Fig. 1 helps stabilize the contribution of visual information, since parts of the model are fine-tuned from A-DKF.

¹https://github.com/XiaoyuBIE1994/DVAE_SE

Table 1: Average values of the SI-SDR, PESQ, and STOI metrics for the input (unprocessed) and output (enhanced) test speech signals. For each method, **top row:** g_t updated by multiplicative rules [12, 15], **bottom row:** g_t updated according to (7).

Metric	SI-SDR (dB)					PESQ					STOI				
	-5	0	5	10	15	-5	0	5	10	15	-5	0	5	10	15
Input	-12.80	-7.72	-2.91	2.04	7.25	1.51	1.76	2.05	2.37	2.85	0.20	0.30	0.43	0.56	0.69
A-VAE	-7.37	-1.92	3.78	8.65	13.07	1.63	1.91	2.20	2.50	2.85	0.21	0.32	0.45	0.59	0.72
	-8.46	-2.60	3.02	8.11	13.01	1.67	1.95	2.25	2.58	2.90	0.22	0.32	0.47	0.60	0.73
AV-VAE	-6.86	-0.83	4.70	9.38	13.90	1.74	2.00	2.31	2.61	2.90	0.20	0.31	0.45	0.59	0.72
	-6.65	-0.86	4.47	9.26	13.77	1.75	2.03	2.34	2.65	2.93	0.22	0.33	0.47	0.61	0.73
A-DKF	-6.50	-1.41	1.99	4.36	5.55	1.48	1.67	1.87	2.02	2.13	0.22	0.33	0.45	0.55	0.64
	-7.02	-0.92	4.76	10.39	14.96	1.78	2.08	2.41	2.75	3.03	0.22	0.35	0.50	0.65	0.77
AV-DKF	-5.04	-0.21	2.93	4.92	5.48	1.39	1.61	1.82	1.97	2.07	0.22	0.33	0.44	0.55	0.63
	-3.78	1.78	7.19	11.66	15.81	1.94	2.24	2.54	2.80	3.05	0.25	0.38	0.52	0.66	0.77




Fig. 2: Effect of g_t update on AV-DKF speech enhancement. From top to bottom, left to right: noisy, clean, output of multiplicative update rule [12, 15], output of optimization-based update rule (7).

Parameters settings. For all the models, the latent dimension is set to $L = 16$. Moreover, the NMF parameters are initialized with non-negative random entries (the same values for all the methods), with $K = 8$. For the DKF models, we used a sequence length of $T = 50$, as in [12]. Also, we set the gamma parameters in (4) to $\alpha = \beta = 1$, and initialized $\mathbf{g}_{1:T}$ with an all-one vector. Furthermore, $\mathbf{z}_{1:T}$ is initialized by feeding $\mathbf{x}_{1:T}$ (and $\mathbf{v}_{1:T}$, for AV models) to the encoder and taking the mean of q_ψ . All the models are trained with the Adam optimizer, with a learning rate of 0.0001 and a batch size of 128. We used early stopping on the validation set with a patience of 50 epochs, i.e., the training stops if the validation loss does not improve after 50 consecutive epochs. For all the models, the number of EM iterations for speech enhancement was set to 100, where at each iteration, the E-step (7) was performed using the Adam optimizer, for 20 iterations and with a learning rate of 0.001.

Results. The speech enhancement results are reported in Table 1, where for each competing method, we report two sets of results corresponding to the two different approaches for updating the scaling parameter g_t in the observation model. In each cell, the top and bottom rows correspond, respectively, to the multiplicative update rule proposed in [12, 15] and the probabilistic framework proposed in this work, i.e., (4) and (7). By inspecting Table 1 we can draw sev-

eral conclusions. First, as can be clearly seen, the proposed update framework for g_t yields more stable and improved performance than the multiplicative updates, especially for the DKF models. Specifically, for the non-sequential VAEs, i.e., A-VAE and AV-VAE, the PESQ and STOI values consistently improved but for SI-SDR there is a degradation, which is less significant for AV-VAE. The effect of the proposed update rule is more noticeable for the DKF models, without which, the models do not work well, especially for higher noise levels. In fact, we noticed in our experiments that without constraining g_t with a proper regularization, e.g., the gamma prior, the enhancement method for the DKF models would lead to improper noise removal. Fig 2 shows an illustrative example with AV-DKF, where we can see that some speech time frames are wrongly estimated either as zero or noisy, which could be due to too small and too large values for g_t , respectively (cf. Equation (9)).

We also see a clear and consistent performance improvement for the DKF models compared with their non-sequential versions. Specifically, AV-DKF exhibits an average performance gain of about **2.5 dB** in SI-SDR, **0.18** in PESQ, and **0.05** (5%) in STOI over AV-VAE. This signifies the importance of temporal modeling. Comparing the results of AV-DKF and its audio-only version, A-DKF, demonstrates average gains of about **2 dB** in SI-SDR, **0.10** in PESQ, and **0.02** (2%) in STOI. This proves the usefulness of visual information for speech enhancement. Moreover, the amount of improvement is higher for larger amounts of noise, i.e., situations wherein the role of visual modality is more highlighted, meaning that AV-DKF is able to efficiently incorporate the useful information of visual data for speech enhancement. Supplementary material, including audio-visual examples, will be available online.²

5. CONCLUSION

We presented the audio-visual deep Kalman filter (AV-DKF) model to learn the prior distribution of clean speech data for speech enhancement. In contrast to the non-sequential models used in the prior work, the AV-DKF model incorporates visual data more efficiently. Besides, we developed an inference algorithm for speech enhancement based on the learned speech prior. Our experiments confirmed the superiority of the AV-DKF model compared with its audio-only version and non-sequential models. Future work includes extending the proposed framework to other dynamical models [11].

²<https://team.inria.fr/multispeech/demos/av-dkf/>

6. REFERENCES

- [1] Daniel Michelsanti, Zheng-Hua Tan, Shi-Xiong Zhang, Yong Xu, Meng Yu, Dong Yu, and Jesper Jensen, “An overview of deep-learning-based audio-visual speech enhancement and separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1368–1396, 2021.
- [2] Zhiqi Kang, Mostafa Sadeghi, Radu Horaud, and Xavier Alameda-Pineda, “Expression-preserving face frontalization improves visually assisted speech processing,” *arXiv preprint arXiv:2204.02810*, 2022.
- [3] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein, “Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–11, 2018.
- [4] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman, “The conversation: Deep audio-visual speech enhancement,” *Proc. Interspeech 2018*, 2018.
- [5] Aviv Gabbay, Asaph Shamir, and Shmuel Peleg, “Visual speech enhancement,” *Proc. Interspeech 2018*, 2018.
- [6] Mostafa Sadeghi, Simon Leglaive, Xavier Alameda-Pineda, Laurent Girin, and Radu Horaud, “Audio-visual speech enhancement using conditional variational auto-encoders,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1788–1800, 2020.
- [7] Mostafa Sadeghi and Xavier Alameda-Pineda, “Mixture of inference networks for vae-based audio-visual speech enhancement,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 1899–1909, 2021.
- [8] Mostafa Sadeghi and Xavier Alameda-Pineda, “Robust unsupervised audio-visual speech enhancement using a mixture of variational autoencoders,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [9] Mostafa Sadeghi and Xavier Alameda-Pineda, “Switching variational auto-encoders for noise-agnostic audio-visual speech enhancement,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [10] Diederik P. Kingma and Max Welling, “Auto-encoding variational bayes,” in *Proc. International Conference on Learning Representations (ICLR)*, April 2014.
- [11] Laurent Girin, Simon Leglaive, Xiaoyu Bie, Julien Diard, Thomas Hueber, and Xavier Alameda-Pineda, “Dynamical variational autoencoders: A comprehensive review,” *Foundations and Trends in Machine Learning*, vol. 15, no. 1-2, pp. 1–175, 2021.
- [12] Xiaoyu Bie, Simon Leglaive, Xavier Alameda-Pineda, and Laurent Girin, “Unsupervised speech enhancement using dynamical variational autoencoders,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2993–3007, 2022.
- [13] Rahul Krishnan, Uri Shalit, and David Sontag, “Structured inference networks for nonlinear state space models,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017, vol. 31.
- [14] Yoshiaki Bando, Masato Mimura, Katsutoshi Itoyama, Kazuyoshi Yoshii, and Tatsuya Kawahara, “Statistical speech enhancement based on probabilistic integration of variational autoencoder and non-negative matrix factorization,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [15] Simon Leglaive, Laurent Girin, and Radu Horaud, “A variance modeling framework based on variational autoencoders for speech enhancement,” in *Proc. IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, September 2018.
- [16] Hirokazu Kameoka, Li Li, Shota Inoue, and Shoji Makino, “Supervised determined source separation with multichannel variational autoencoder,” *Neural computation*, vol. 31, no. 9, pp. 1891–1914, 2019.
- [17] Simon Leglaive, Xavier Alameda-Pineda, Laurent Girin, and Radu Horaud, “A recurrent variational autoencoder for speech enhancement,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [18] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *Proc. International Conference on Learning Representations ICLR*, May 2015.
- [19] Jonathan Le Roux, Scott Wisdom, Hakan Erdogan, and John R Hershey, “SDR—half-baked or well done?,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 626–630.
- [20] Cees H. Taal, Richard C. Hendriks, Richard Heusdens, and Jesper Jensen, “An algorithm for intelligibility prediction of time–frequency weighted noisy speech,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2125–2136, February 2011.
- [21] Antony W. Rix, John G. Beerends, Michael P. Hollier, and Andries P. Hekstra, “Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs,” in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, May 2001.
- [22] Naomi Harte and Eoin Gillen, “TCD-TIMIT: An audio-visual corpus of continuous speech,” *IEEE Transactions on Multimedia*, vol. 17, no. 5, pp. 603–615, 2015.
- [23] Ahmed Hussen Abdelaziz et al., “NTCD-TIMIT: A new database and baseline for noise-robust audio-visual speech recognition,” in *Interspeech*, 2017, pp. 3752–3756.
- [24] John Garofolo, David Graff, Doug Paul, and David Pallett, “CSR-I (WSJ0) sennheiser ldc93s6b,” *Philadelphia: Linguistic Data Consortium*, 1993.
- [25] Brais Martinez, Pingchuan Ma, Stavros Petridis, and Maja Pan-tic, “Lipreading using temporal convolutional networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6319–6323.