

CAN WE TRUST DEEP SPEECH PRIOR?

Ying Shi¹, Haolin Chen², Zhiyuan Tang², Lantian Li^{2}, Dong Wang^{2*}, Jiqing Han^{1*}*

Harbin Institute of Technology, Harbin, China
Center for Speech and Language Technologies, Tsinghua University, China

ABSTRACT

Recently, speech enhancement (SE) based on deep speech prior has attracted much attention, such as the variational auto-encoder with non-negative matrix factorization (VAE-NMF) architecture. Compared to conventional approaches that represent clean speech by shallow models such as Gaussians with a low-rank covariance, the new approach employs deep generative models to represent the clean speech, which often provides a better prior. Despite the clear advantage in theory, we argue that deep priors must be used with much caution, since the likelihood produced by a deep generative model does not always coincide with the speech quality. We designed a comprehensive study on this issue and demonstrated that based on deep speech priors, a reasonable SE performance can be achieved, but the results might be suboptimal. A careful analysis showed that this problem is deeply rooted in the disharmony between the flexibility of deep generative models and the nature of the maximum-likelihood (ML) training.

Index Terms— Speech enhancement, deep generative model, maximum likelihood

1. INTRODUCTION

Single-channel speech enhancement (SE) aims to restore the quality and possibly also the intelligibility of noise corrupted speech [1, 2]. Most modern SE methods are based on the maximum-likelihood (ML) statistical framework. The most popular statistical model is the complex Gaussian on the short-time Fourier transform (STFT) coefficients [3], although some super-Gaussian models have been experimented [4]. The choice for Gaussian is not only theoretically sound (e.g., justified by the central limit theorem [5]), but also computationally attractive. Specifically, by assuming that clean speech and noise spectra are both zero-mean Gaussians and are independent and additive, the ML inference leads to the spectra subtraction approach [6, 7] or the Wiener filter [8, 9], depending on different model assumptions.

The ML-based statistical SE approach relies on the covariance estimation of the Gaussians for clean speech and noise. Non-negative matrix factorization (NMF) [10] is perhaps the most popular variance model, which assumes that the variances at different time-frequency (T-F) bins possess a low-rank structure. Although elegant in theory, the low-rank covariance model may be too simple to represent complex signals such as speech. Recently, several authors have proposed to use deep generative models to represent speech. For example, the VAE-NMF architecture [9, 11, 12, 13] uses a variational auto-encoder (VAE) to model the covariance of the complex Gaussian for speech signal, while the noise model is still based on NMF. Subakan et al. [14] models the power spectrum of speech signals as a Poisson distribution and trains a deep neural net to predict the variance of the distribution. The central idea of these studies is to learn a powerful prior model for clean speech, which we call a deep speech prior. In theory, this deep prior can provide very rich knowledge for clean speech, hence a powerful enhancement. This has been demonstrated in [9, 11].

In this paper, we argue that caution must be given to use the deep speech prior models. The argument is based on a key observation that a high likelihood measured by a deep generative model does not mean that the sample is similar to the training data, which is clean speech in our case. This means that the likelihood produced by deep speech prior models does not always reflect the quality of the speech, leading to suboptimal results with the ML-based statistical SE framework. We will design a series of experiments to investigate this issue. In particular, our analysis is based on normalization flow (NF) [15], a typical deep generative model that can measure the likelihood accurately. Our results show that reasonable SE results can be obtained with deep speech priors, at least when the noise can be well modeled. However, there is a potential risk that the inference converges to unrealistic speech, due to the unreliable likelihood produced by the deep speech prior models. We attribute this problem to the disharmony between the flexibility of deep generative models and the nature of the ML training criterion.

The rest of the paper is organized as follows. Section 2 revisits the ML-based statistical SE framework, and Section 3 reviews the deep generative models. The flow-based deep speech prior model and the inference algorithm based on gra-

This work was supported by the National Natural Science Foundation of China (NSFC) under the project No. 61633013. L.L. (lilt@csl.org), D.W. (wangdong99@mails.tsinghua.edu.cn) and J.H. (jqhan@hit.edu.cn) are the corresponding authors.

dient descent (GD) are presented in Section 4. Experiments and analysis are presented in Section 5, and Section 6 concludes the paper.

2. REVISIT STATISTICAL SPEECH ENHANCEMENT

2.1. ML-based statistical SE framework

Let \mathbf{x}_t denote a segment of clean speech of length M , and \mathbf{n}_t denote a segment of noise signal in the same length. Suppose the noise is additive, the noisy speech \mathbf{y}_t can be written by:

$$\mathbf{y}_t = \mathbf{x}_t + \mathbf{n}_t. \quad (1)$$

The clean speech can be estimated from the noisy speech following the ML principle¹:

$$\begin{aligned} \hat{\mathbf{x}}_t &= \arg \max_{\mathbf{x}_t} p(\mathbf{x}_t | \mathbf{y}_t) \\ &= \arg \max_{\mathbf{x}_t} p(\mathbf{x}_t) p(\mathbf{y}_t | \mathbf{x}_t) \\ &= \arg \max_{\mathbf{x}_t} p(\mathbf{x}_t) p(\mathbf{n}_t), \end{aligned} \quad (2)$$

where we have assumed \mathbf{x}_t and \mathbf{n}_t are independent.

Let $\mathbf{s}_t^x, \mathbf{s}_t^n, \mathbf{s}_t^y$ denote the complex spectrum of $\mathbf{x}_t, \mathbf{n}_t, \mathbf{y}_t$ respectively. The superposition of Eq.(1) still holds since the Fourier transform is a linear operator.

$$\mathbf{s}_t^y = \mathbf{s}_t^x + \mathbf{s}_t^n. \quad (3)$$

The ML estimation for \mathbf{s}_t^x has a similar form as in Eq.(2):

$$\hat{\mathbf{s}}_t^x = \arg \max_{\mathbf{s}_t^x} p(\mathbf{s}_t^x) p(\mathbf{s}_t^n). \quad (4)$$

If $p(\mathbf{s}_t^x)$ and $p(\mathbf{s}_t^n)$ are all complex Gaussians and the covariances are diagonal, i.e.,

$$\begin{aligned} p(s_{tf}^x) &= \mathcal{N}_C(0, \sigma_{tf}^x) \\ p(s_{tf}^n) &= \mathcal{N}_C(0, \sigma_{tf}^n), \end{aligned}$$

the posterior $p(\mathbf{s}_t^x | \mathbf{s}_t^y)$ can be represented in the following element-wised form:

$$p(s_{tf}^x | s_{tf}^y) = \mathcal{N}_C\left(\frac{\sigma_{tf}^x}{\sigma_{tf}^x + \sigma_{tf}^n} y_{tf}, \frac{\sigma_{tf}^n \sigma_{tf}^x}{\sigma_{tf}^n + \sigma_{tf}^x}\right).$$

Therefore, the ML estimation for \mathbf{s}_t^x can be written by:

$$\hat{s}_{tf}^x = \frac{\sigma_{tf}^x}{\sigma_{tf}^x + \sigma_{tf}^n} y_{tf}. \quad (5)$$

Note that this is the form of Wiener filter [8, 9].

¹This is essentially the *maximum a posteriori* (MAP) principle, though we regard it as a generalized ML, considering that it maximizes the conditional likelihood.

2.2. NMF prior model

The paramount task of the statistical SE approach is to estimate the variance σ_{tf}^x and σ_{tf}^n for each T-F bin. Note that they correspond to the power spectral density (PSD) of clean speech and noise, respectively. Traditional algorithms focus on the noise PSD, either heuristically (e.g., by voice activity detection [16, 7]) or principally (e.g., by speech present probability [17, 18, 19]). Since this estimation heavily relies on the assumption that the noise is time-invariant (at least change slowly), it cannot deal with dynamic noises. A principle solution is to learn the patterns of the speech and noise *in prior*, and then infer the clean speech based on these prior models. For example, hidden Markov models (HMMs) [20] or Gaussian mixture models (GMM) [21] were used as the prior models in early research.

NMF [10] is perhaps the most popular prior model. Initial application of NMF to SE is largely heuristic, which distributes the PSD of noisy speech to PSDs of clean speech and noise, e.g., $\mathbf{V}_y \approx \mathbf{V}_x + \mathbf{V}_n$, where each PSD is in the form of $\mathbf{V} \approx \mathbf{W}\mathbf{H}$, where $\mathbf{W}_{0,+}^{K \times L}$ represents L bases, and $\mathbf{H} \in R_{0,+}^{L \times T}$ represents the activations on the bases [22, 23]. Usually $L \ll K$, and so NMF represents a low-rank decomposition for the PSD. Févotte et al. [24] demonstrated that if the spectrum follows a complex Gaussian and the variances on all the T-F bins are assumed to be in form of $\mathbf{\Sigma} = \mathbf{W}\mathbf{H}$, the ML estimation with respect to \mathbf{W} and \mathbf{H} is equivalent to the NMF on \mathbf{V} with the Itakura-Saito (IS) divergence. Due to this equivalence, the ML estimation for the variances (equivalently, the PSD) of clean speech and noise can be conducted by IS-based NMF.

3. DEEP SPEECH PRIOR: POTENTIAL AND RISK

3.1. Deep speech prior

A key advantage of the NMF-based SE approach is that the activation matrix \mathbf{H} is estimated on-the-fly during inference, thus compensating for the variation on the magnitude of the signal. Moreover, since the distributions of both clean speech and noise are Gaussian, the inference is simply in the form of Wiener filtering. The disadvantage is that the low-rank form of the variances may prevent it from representing complex signals, e.g., speech.

Recently, deep generative models have been extensively studied, including the auto-regressive model [25, 26], variational auto-encoder (VAE) [27], normalization flow [28, 29, 30], and generative adversarial network (GAN) [31]. A key advantage of the deep generative models is that they can represent complex natural data, by either density estimation or sampling, therefore suitable for modeling speech signals.

Several groups have proposed to use deep generative models to present clean speech in SE, which we call deep speech prior models. Among the proposals, VAE-NMF is in particular attractive [9, 11, 12, 13]. In this architecture, VAE is used

to model the covariance of the complex Gaussian for clean speech, and the noise model is still based on NMF. Since both the models are still Gaussian, the inference remains in the form of Wiener filtering.

3.2. Pitfall of deep speech prior

Although interesting, some recent research observations cast a shadow on the employment of deep speech prior models. For example, Nalisnick et al. [32] and Choi et al. [33] concurrently found that a normalization flow may assign out-of-domain data a high likelihood, even higher than that of the training data. They also found that the same problem occurred in VAE and PixelCNN. This is very weird at the first glance, as all the models are trained with the ML criterion, and therefore should assign the highest likelihood to the training data.

To investigate if the same problems exists with the deep speech prior, we trained a VAE-NMF model using the code released² by Leglavie et al. [11], and plotted the histogram of the likelihood with different types of speech data. Since the likelihood cannot be exactly obtained with VAE, we instead use the evidence lower bound (ELBO). Note that the VAE training maximizes ELBO, so the ELBO values on the training data (clean speech) are expected to be the highest. The results are shown in Fig.1. It can be seen that the clean speech and noisy speech are well separated, which is nice; however, the white noise seems obtain higher ELBO than the noisy speech, which is not consistent with the intuition. More seriously, the whispering babble obtains even higher ELBO than the clean speech.

The disharmony between likelihood (or ELBO) and speech quality (more precisely, the resemblance to the training data) is a disaster for the ML-based statistical SE, as the ML principle does not optimize the speech quality, and may result in unrealistic speech.

4. DEEP SPEECH PRIOR WITH NORMALIZATION FLOW

In this paper, we will conduct a comprehensive study for the deep speech prior. To this purpose, we made several decisions as shown below. All these decisions help us isolate the issue with the deep speech prior model.

- We use normalization flow (NF) rather than VAE as the speech prior model. The NF model can measure the likelihood exactly, rather than an approximation (i.e., ELBO) as with VAE. Specifically, we choose the WaveGlow [15] model, which has shown great potential in speech synthesis and speech conversion.
- We build the SE model in the time domain, and model the speech prior $p(\mathbf{x}_t)$ directly, rather than assume

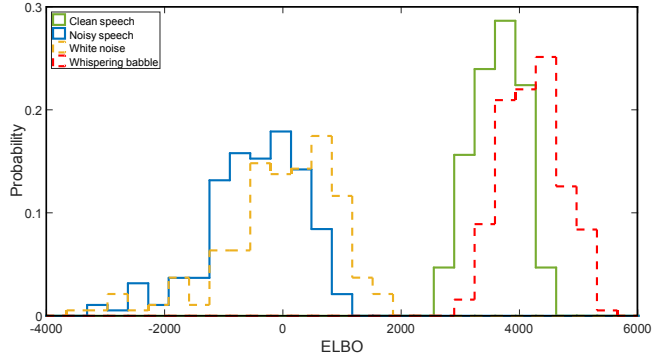


Fig. 1. Histogram of ELBO for clean speech (green solid stairs), noisy speech (blue solid stairs), white noise (orange dashed stairs), whispering babble (red dashed stairs), computed from a VAE trained with clean speech.

a Gaussian and then model the covariance (as in the VAE-NMF approach). This eliminates all the assumptions/approximations made when transferring to the frequency-domain and conducting modeling there.

- We constrain our study with white noise. By white noise, the distribution in the time domain is a Gaussian and can be modeled precisely. This will exclude potential problems caused by inaccurate noise modeling, and so fully expose the problems caused by the deep speech prior model.
- The inference is conducted based on the original ML principle shown in Eq.(2), based on the gradient descent (GD) algorithm. Therefore, it does not rely on any specific assumption/approximation, and so will neither conceal any problem of the model nor introduce non-existing problems.

We will start from reviewing the normalization flow model, and then describe the SE-approach with a flow-based deep speech prior model.

4.1. Revisit normalization flow and WaveGlow

Normalization flow (NF) [28, 29, 30, 34] is a typical deep generative model. It employs an invertible transformation $f^{-1}(\mathbf{x})$ that transforms a complex distribution $p(\mathbf{x})$ to a simple Gaussian. Formally, $\mathbf{z} = f^{-1}(\mathbf{x})$ where $p(\mathbf{z})$ is a Gaussian, i.e., $\mathbf{z} \sim \mathcal{N}(\mathbf{z}; 0, I)$. According to the principle of distribution transformation, we have:

$$p(\mathbf{x}) = p(\mathbf{z}) \left| \det \frac{\partial f^{-1}}{\partial \mathbf{x}} \right|, \quad (6)$$

where $\det \frac{\partial f^{-1}}{\partial \mathbf{x}}$ is the determinant of the Jacobian of f^{-1} at \mathbf{x} . The model is trained with the ML criterion, where the objective is as follows:

²<https://github.com/sleglavie/MLSP-2018>

$$\sum_i \log p(\mathbf{x}_i) = \sum_i \log p(\mathbf{z}_i) + \sum_i \log \left| \det \frac{\partial f^{-1}}{\partial \mathbf{x}_i} \right|. \quad (7)$$

After the training converges, the likelihood $p(\mathbf{x})$ can be computed according to Eq.(6). Papamakarios’s review paper and the references therein are good resource for the NF model [34].

WaveGlow is a special NF model [15]. The structure is inherited from Glow [30], with modification to suit for processing speech signals. In our experiment, we reused the WaveGlow architecture presented in [15]. The input is a window of 80 speech samples in the time domain, and the output is the latent variable with the same dimension as the input. The entire flow involves 12 blocks, and each consists of two components: an invertible 1×1 convolution layer and an affine coupling layer. The output of the latent variable \mathbf{z} is arranged in a multi-scale way, i.e., 20 dimensions are output every 4 flow blocks. Different from the original WaveGlow model, we do not need the conditional input as our focus is in the marginal distribution of $p(\mathbf{x})$. Readers can refer to Ryan’s paper [15] for more details of the WaveGlow structure.

4.2. Speech enhancement with flow-based prior

A nice property of the NF model is that the exact likelihood $p(\mathbf{x})$ can be computed via Eq.(6), without any assumption on the distribution (as in NMF) or any approximation on the computation (as in VAE). This permits us pursuing the ML inference for the clean speech by optimizing Eq.(4) directly, by using any numerical optimization algorithm, such as gradient descent (GD).

In theory, the flow model can be used to model the prior for both the clean speech and the noise signal, and the inference based on Eq.(4) is theoretically optimal if the prior models produce accurate likelihood for clean speech and noise. However, as we have discussed, the likelihood produced by deep generative models, including the WaveGlow model we used, is suspicious. To simplify the investigation, we use the flow model to represent the clean speech only, and constrain the noise to be Gaussian. Therefore, the ML-based inference can be conducted by optimizing the following objective with respect to \mathbf{x}_t :

$$\begin{aligned} \log p(\mathbf{x}_t | \mathbf{y}_t) &\propto \log p(\mathbf{x}_t) + \log p(\mathbf{n}_t) \\ &= \log N(\mathbf{z}_t; \mathbf{0}, \mathbf{I}) + \log \left| \det \frac{\partial f^{-1}}{\partial \mathbf{x}_t} \right| \\ &\quad + \log N(\mathbf{y}_t - \mathbf{x}_t; \mathbf{0}, \sigma \mathbf{I}) \end{aligned} \quad (8)$$

where $\mathbf{z}_t = f^{-1}(\mathbf{x}_t)$ and f^{-1} is the flow model. σ is the variance of the noise. Since all the terms in the above objective can be computed easily, the optimization can be easily conducted by GD. Note that the Gaussian noise is just the white

noise which is ubiquitous in ambient and most destructive for speech signal.

It should be highlighted that both the noise model and the inference are accurate in our SE architecture, which means that any unexpected behaviors should be attributed to the speech prior model $p(\mathbf{x}_t)$. We therefore expose the speech prior model by this specific design.

5. EXPERIMENTS AND ANALYSIS

5.1. Data

We use the TIMIT dataset [35] in our experiments. For the training set, 4620 utterances were selected from 462 speakers (136 females and 326 males), and the total amount of speech was 4 hours. It was used to train the prior models (NMF, VAE and WaveGlow). For the test set, 192 utterances were selected from 25 speakers (8 females and 16 males), and then corrupted by white noise at 4 different signal-to-noise ratios (SNRs) (-6, 0, 6 and 9 dB). Perceptual evaluation of speech quality (PESQ) [36], short-time objective intelligibility (STOI) [37], and frequency-weighted segmental SNR (fwSNRseg) [38] were used as metrics to evaluate the SE performance.

5.2. Models

We test three SE models: NMF, VAE-NMF and Flow-Gauss. The features and settings of these models are described as follows:

- **NMF**: A supervised frequency-domain model where both clean speech and noise are modeled by a complex Gaussian whose variance is modeled by NMF. The model is based on 257-dimensional spectrum. The number of bases is empirically set to 40 for both speech and noise NMFs. The speech bases are trained with the TIMIT training set, and the noise bases are trained with 2 hours of white noise sampled with different levels of power. The inference is based on the IS-based NMF.
- **VAE-NMF**: A supervised frequency-domain model, built on 257-dimensional spectrum. The prior model for the clean speech is a complex Gaussian whose covariance is modeled by a VAE, and the noise model is a complex Gaussian whose covariance is modeled by an NMF. The VAE model is constructed by using the source code provided by Leglaive [11], and trained with the TIMIT training set. The base matrix W of the NMF is trained with 2 hours of white noise with different levels of variances, and the activation H is derived on-the-fly during the inference. The inference is based on Bayesian inference [11].

- Flow-Gauss:** The time-domain approach proposed in Section 4.2. The speech prior model is a WaveGlow and the noise model is a Gaussian $N(\mathbf{0}, \sigma\mathbf{I})$. The WaveGlow is implemented using the source code provided by [15], where the input is a window of 80 samples. The Adam optimizer is used to train the model, with the batch size set to 4 and the learning rate set to 0.0001. The inference is based on GD, using the Adam optimizer with the learning rate of 0.0005. The optimization will be stopped if the likelihood is as high as the clean speech, for which we set a threshold -7.0 in our experiments.

We test two settings for the variance σ of the Gaussian noise: (1) **Flow-Gauss (Est)**, where σ is manually set to be 100, 200, 2000, 4000 for the SNR conditions -6, 0, 6 and 9 dB respectively. Note that for each SNR condition, σ is different for different utterances due to their different speech magnitudes, so the setting for σ here is an estimated value for the strength of the noise. (2) **Flow-Gauss (Opt)**, where the true σ is set for each utterance. Although this is not a practical scenario, the oracle setting for the noise model is useful for us to investigate the problem of the speech prior model.

5.3. Basic results

The basic SE results with different models are shown in Table 1. It can be seen that all the SE methods obtained noticeable performance gains in terms of the three metrics. The VAE-NMF approach generally outperforms the NMF approach, demonstrating that the deep speech prior model is more powerful. Moreover, the flow-based approach obtained comparable performance as VAE-NMF, even better in terms of fwSNRseg. In high-SNR conditions, the flow-based models seem more superior.

5.4. Analysis 1: Maximizing likelihood does not maximize quality

Although promising, we found that the Flow-Gauss model does not result in an optimal SE performance, even with the oracle setting for noise. To see this, we randomly selected 30 utterances from the training set and corrupted them by noise. With the Flow-Gauss (Opt) setting, the optimization is conducted for 500 iterations, and the log likelihood, PESQ, STOI and fwSNRseg are shown in Fig. 2.

It can be found that the likelihoods in different SNR conditions are constantly increased with more iterations, even higher than the likelihood of clean speech in the training set (read dot line). However, the three SE metrics show a different pattern: they are firstly increased (indicating a better speech quality), and then are decreased with more iterations. This clearly demonstrated that a higher likelihood does not mean a better speech quality.




Fig. 2. The change of log likelihood, PESQ, STOI and fwSNRseg during the inference, based on the Flow-Gauss (Opt) model.




Fig. 3. Footprint of the inference process with the Flow-Gauss (Opt) model, plot by t-SNE [39]. The red star is the target clean speech, and the green dots represent the enhanced speech at different steps of the optimization. [#] is the log likelihood.

Fig. 3 shows the optimization process for one noisy speech, where the enhanced speech after a particular iteration is presented by a green dot. It can be observed that the optimization process attempted to approach the target clean speech, however it just passed by the target and continued its journey to pursue a higher likelihood.

5.5. Analysis 2: Unreliable likelihood is caused by the deep prior model

Since the noise model is ideal and the ML inference is theoretically solid, the failure of the ML-oriented optimization must be attributed to the deep speech prior model, i.e., WaveGlow in our setting. Specifically, the model has assigned a high likelihood to an unreal speech and so the ML optimization was attracted to a weird state, as shown in Fig. 3. This argument can be easily verified by maximizing $p(\mathbf{x})$ provided by the prior model and see where the optimization goes. The results are shown in Fig. 4.

It can be observed that the optimization indeed tries to

Table 1. Basic results with different SE models.

Models	-6dB			0dB			6dB			9dB		
	PESQ	STOI	fwSNRseg	PESQ	STOI	fwSNRseg	PESQ	STOI	fwSNRseg	PESQ	STOI	fwSNRseg
Noisy	1.032	0.582	4.068	1.046	0.722	4.816	1.098	0.849	6.323	1.161	0.898	7.446
NMF	1.144	0.592	4.202	1.307	0.733	5.799	1.523	0.830	7.685	1.647	0.865	8.742
VAE-NMF	1.211	0.668	3.307	1.319	0.800	5.590	1.521	0.887	8.776	1.738	0.917	10.203
Flow-Gauss (Est)	1.096	0.637	5.160	1.301	0.791	7.627	1.693	0.894	10.585	1.920	0.925	12.044
Flow-Gauss (Opt)	1.094	0.647	5.215	1.304	0.799	7.772	1.718	0.898	10.796	1.977	0.929	12.388

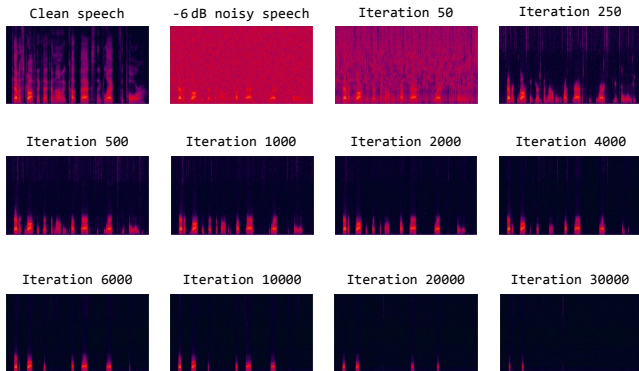


Fig. 4. The enhanced speech from the optimization process that maximizes the prior likelihood $p(\mathbf{x})$.

remove the noise at the beginning, however it does not stop when a reasonable enhancement has been obtained. After thousands of iterations, the remaining spectrum tends to be zero. It indicates that the speech prior assigns a higher likelihood to silence. Notably, this observation coincides with the results in Fig. 1 where the speech prior model is a VAE.

We argue that the unreliable likelihood associated with the deep speech prior model is caused by the ML criterion used for model training. More specifically, the goal of the ML training is to increase the probability density at the locations of the training data. To achieve this goal, the optimizer will try to twist the density function by using a complex transformation function. However, this twisted density function is very risky and may assign unexpected high likelihood for unreal data.

A simulation experiment is conducted to verify this argument, where we use a flow model based on the MAF structure [40]. MAF is much simpler than WaveGlow, and can help us identify the problem by using a small set of simulation data. We sample 500 2-dimensional data from two Gaussians, and train the MAF (5 blocks) for 2500 iterations. After the training, the transformation and the density function are shown in Fig. 5. It can be seen that with such a simple flow, the transformation is highly complex and the density function is noticeably twisted. Most importantly, there is a large region that is far from the training data but the density is high.




Fig. 5. The density function produced by a 2-dimensional MAF model. The pictures are: (a) the training data; (b) the training data transformed into the latent space; (c) data sampled from the model; (d) a color-grid in the data space; (e) the color-grid transformed to the latent space; (f) density function in the data space. The two squares in (f) mark the two Gaussians of the training data shown in (a).

6. CONCLUSION

We conducted a comprehensive study for the behavior of deep speech prior models in speech enhancement. To ensure a focused investigation, we choose normalization flow as the deep prior model, build the model in the time domain, constrain the noise to be Gaussian, and use gradient descent to perform inference. Our experiments show that our flow-based SE approach obtains comparable or even better performance than the state-of-the-art VAE-NMF approach. However, further analyses show that the present result is far from optimal, due to the unreliable likelihood produced by the flow model. We argue that this is a general problem for all deep generative models based on ML training, and demonstrated that it may be caused by the inharmonious combination of the complex distribution transformation and the ML-based training criterion. Therefore, caution must be given when employing such ML-based deep prior models in speech enhancement. More training data may alleviate this problem, but a principled solution requires more theoretical study.

7. REFERENCES

- [1] Philipos C Loizou, *Speech enhancement: theory and practice*, CRC press, 2013.
- [2] Robert Rehr, *Robust speech enhancement using statistical signal processing and machine-learning*, Ph.D. thesis, 2019.
- [3] David R Brillinger, *Time series: data analysis and theory*, SIAM, 2001.
- [4] Rainer Martin, “Speech enhancement based on minimum mean-square error estimation and supergaussian priors,” *IEEE transactions on speech and audio processing*, vol. 13, no. 5, pp. 845–856, 2005.
- [5] W Pearlman and R Gray, “Source coding of the discrete Fourier transform,” *IEEE Transactions on information theory*, vol. 24, no. 6, pp. 683–692, 1978.
- [6] Steven Boll, “Suppression of acoustic noise in speech using spectral subtraction,” *IEEE Transactions on acoustics, speech, and signal processing*, vol. 27, no. 2, pp. 113–120, 1979.
- [7] Robert McAulay and Marilyn Malpass, “Speech enhancement using a soft-decision noise suppression filter,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 28, no. 2, pp. 137–145, 1980.
- [8] Peter Vary and Rainer Martin, *Digital speech transmission: Enhancement, coding and error concealment*, John Wiley & Sons, 2006.
- [9] 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 *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 716–720.
- [10] Daniel D Lee and H Sebastian Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [11] Simon Leglaive, Laurent Girin, and Radu Horaud, “A variance modeling framework based on variational autoencoders for speech enhancement,” in *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2018, pp. 1–6.
- [12] S. Leglaive, X. Alameda-Pineda, L. Girin, and R. Horaud, “A recurrent variational autoencoder for speech enhancement,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 371–375.
- [13] Li Li, Hirokazu Kameoka, and Shoji Makino, “Fast mvae: Joint separation and classification of mixed sources based on multichannel variational autoencoder with auxiliary classifier,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 546–550.
- [14] Y Cem Subakan and Paris Smaragdis, “Generative adversarial source separation,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 26–30.
- [15] Ryan Prenger, Rafael Valle, and Bryan Catanzaro, “Waveglow: A flow-based generative network for speech synthesis,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [16] Jongseo Sohn and Wonyong Sung, “A voice activity detector employing soft decision based noise spectrum adaptation,” in *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP’98 (Cat. No. 98CH36181)*. IEEE, 1998, vol. 1, pp. 365–368.
- [17] Hans-Günter Hirsch and Christoph Ehrlicher, “Noise estimation techniques for robust speech recognition,” in *1995 International conference on acoustics, speech, and signal processing*. IEEE, 1995, vol. 1, pp. 153–156.
- [18] Timo Gerkmann and Richard C Hendriks, “Noise power estimation based on the probability of speech presence,” in *2011 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2011, pp. 145–148.
- [19] Aleksej Chinaev, Jahn Heymann, Lukas Drude, and Reinhold Haeb-Umbach, “Noise-presence-probability-based noise PSD estimation by using DNNs,” in *Speech Communication; 12. ITG Symposium*. VDE, 2016, pp. 1–5.
- [20] Yariv Ephraim, “A Bayesian estimation approach for speech enhancement using hidden Markov models,” *IEEE Transactions on Signal Processing*, vol. 40, no. 4, pp. 725–735, 1992.
- [21] Jiucang Hao, Te-Won Lee, and Terrence J Sejnowski, “Speech enhancement using Gaussian scale mixture models,” *IEEE transactions on audio, speech, and language processing*, vol. 18, no. 6, pp. 1127–1136, 2009.
- [22] Paris Smaragdis, “Convolutional speech bases and their application to supervised speech separation,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 1, pp. 1–12, 2006.

- [23] Dong Wang, Ravichander Vippera, Nicholas Evans, and Thomas Fang Zheng, "Online non-negative convolutive pattern learning for speech signals," *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 44–56, 2012.
- [24] Cédric Févotte, Nancy Bertin, and Jean-Louis Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence: With application to music analysis," *Neural computation*, vol. 21, no. 3, pp. 793–830, 2009.
- [25] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al., "Conditional image generation with pixelcnn decoders," in *Advances in neural information processing systems*, 2016, pp. 4790–4798.
- [26] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [27] Diederik P Kingma and Max Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [28] Laurent Dinh, David Krueger, and Yoshua Bengio, "Nice: Non-linear independent components estimation," *arXiv preprint arXiv:1410.8516*, 2014.
- [29] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio, "Density estimation using real nvp," *arXiv preprint arXiv:1605.08803*, 2016.
- [30] Durk P Kingma and Prafulla Dhariwal, "Glow: Generative flow with invertible 1x1 convolutions," in *Advances in neural information processing systems*, 2018, pp. 10215–10224.
- [31] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [32] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan, "Do deep generative models know what they don't know?," *arXiv preprint arXiv:1810.09136*, 2018.
- [33] Hyunsun Choi and Eric Jang, "Generative ensembles for robust anomaly detection," 2018.
- [34] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan, "Normalizing flows for probabilistic modeling and inference," *arXiv preprint arXiv:1912.02762*, 2019.
- [35] John S Garofolo, "TIMIT acoustic phonetic continuous speech corpus," *Linguistic Data Consortium*, 1993, 1993.
- [36] 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 *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*. IEEE, 2001, vol. 2, pp. 749–752.
- [37] 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, 2011.
- [38] Yi Hu and Philipos C Loizou, "Evaluation of objective quality measures for speech enhancement," *IEEE Transactions on audio, speech, and language processing*, vol. 16, no. 1, pp. 229–238, 2007.
- [39] Laurens van der Maaten and Geoffrey Hinton, "Visualizing data using t-SNE," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [40] George Papamakarios, Theo Pavlakou, and Iain Murray, "Masked autoregressive flow for density estimation," in *Advances in Neural Information Processing Systems*, 2017, pp. 2338–2347.