



Score-Based Training for Energy-Based TTS Models

Wanli Sun, Anton Ragni

¹School of Computer Science, University of Sheffield, Sheffield, UK

wsun20@sheffield.ac.uk, a.ragni@sheffield.ac.uk

Abstract

Noise contrastive estimation (NCE) is a popular method for training energy-based models (EBM) with intractable normalisation terms. The key idea of NCE is to learn by comparing unnormalised log-likelihoods of the reference and noisy samples, thus avoiding explicitly computing normalisation terms. However, NCE critically relies on the quality of noisy samples. Recently, sliced score matching (SSM) has been popularised by closely related diffusion models (DM). Unlike NCE, SSM learns a gradient of log-likelihood, or *score*, by learning distribution of its projections on randomly chosen directions. However, both NCE and SSM disregard the form of log-likelihood function, which is problematic given that EBMs and DMs make use of first-order optimisation during inference. This paper proposes a new criterion that learns scores more suitable for first-order schemes. Experiments contrast these approaches for training EBMs.

Index Terms: speech synthesis, sliced score matching, energy-based models

1. Introduction

Recently, energy-based models (EBMs) [1, 2] have drawn some attention in the text-to-speech (TTS) area [3]. These probabilistic models define the log-likelihood of speech given text as the difference between negative (unnormalised) energy function and the logarithm of the corresponding normalisation term. Due to the difficulty of computing normalisation terms, noise contrastive estimation (NCE) has been previously utilised to train EBMs in applications such as speech recognition [4], machine translation [5], text generation [6], and TTS [3]. The NCE compares unnormalised log-likelihoods of ‘positive’ (reference) examples to unnormalised log-likelihoods of ‘negative’ (noisy) samples to decide the direction for updating model parameters. Unlike language modelling, where NCE appears to work well with a wide choice of negative samples [7], TTS appears to be much more sensitive to the quality of negative samples and necessitates the development of elaborate sampling schemes [3].

Several alternative approaches have been proposed in the literature: score matching (SM) [8], denoising score matching (DSM) [9], and sliced score matching (SSM) [10]. All these alternatives attempt to predict the gradient of log-likelihood with respect to data (speech in TTS), or *score*. However, the use of SM would require computing the trace of the Hessian, which is expensive for automatic differentiation packages [11]. The introduction of noise to data in DSM training leads to inconsistency between inference and training [11]. Instead of matching scores, SSM matches their projections on randomly chosen directions, thus avoiding much expensive computation. SSM has been used with EBMs in many domains, e.g., imitation learn-

ing [12] and out-of-distribution detection [13, 14], but has not yet been investigated in TTS.

Score-based training of EBMs and currently popular diffusion style models [15, 16, 17] is typically followed by a variant of Markov Chain Monte-Carlo (MCMC) for performing inference. As a type of first-order optimisation, MCMC-style approaches are sensitive to the nature of objective function they optimise. For example, a highly non-convex shape may lead to issues in converging to an acceptable solution within acceptable time. Despite this, score-based approaches appear to focus on learning gradients as accurately as possible, regardless of the shape of the underlying objective. This paper introduces a new score-based learning objective, called *delta loss*, that penalises those scores that do not lie on linear paths between current speech estimates and reference speech samples. As such, delta loss is believed to be more appropriate for score-based models that utilise MCMC-style first-order optimisation approaches for inference. This paper also shows that although differently motivated the new loss function is closely linked with training objectives of currently popular flow matching approaches [16].

This paper makes the following key contributions:

1. first application of SSM to training EBMs in TTS;
2. a new score-based loss function for EBMs;

and the rest of it is organised as follows. Section 2 discusses EBMs and related work, including training methods and connections between EBMs and diffusion models. Section 3 describes existing score-based training approaches and introduces delta loss as a possible alternative. Experimental results and discussion are presented in Section 4. Conclusions drawn from this work are presented in Section 5.

2. Energy-Based Models

Given a sequence of text tokens \mathbf{x} , an energy-based model (EBM) of spectral feature sequences \mathbf{Y} could be defined as [3]

$$p_{\theta}(\mathbf{Y}|\mathbf{x}) = \frac{1}{Z_{\theta}(\mathbf{x})} \exp(-E_{\theta}(\mathbf{x}, \mathbf{Y})), \quad (1)$$

where $E_{\theta}(\mathbf{x}, \mathbf{Y})$ is an energy function with parameters θ that connects text and speech such that smaller values correspond to better matches between text and speech and *vice versa*, and $Z_{\theta}(\mathbf{x})$ denotes the normalising term, which is intractable to calculate. Although intractable, normalising terms cancel out in log-likelihood ratios, which is the fact exploited in Section 2.1, or disappear in the following gradients of log-likelihood

$$\nabla_{\mathbf{Y}} \log p_{\theta}(\mathbf{Y}|\mathbf{x}) = -\nabla_{\mathbf{Y}} E_{\theta}(\mathbf{x}, \mathbf{Y}) - \underbrace{\nabla_{\mathbf{Y}} \log Z_{\theta}(\mathbf{x})}_{=0}, \quad (2)$$

which is the fact exploited in Section 3 and by gradient-based inference schemes. In its simplest form, inference with EBMs




Figure 1: Two ways of computing scores for EBMs: analytic (top) and predictive (bottom)

could be performed using the following version of Langevin Markov Chain Monte-Carlo (MCMC) [3]

$$\mathbf{Y}^{(N+1)} \leftarrow \mathbf{Y}^{(N)} + \rho \nabla_{\mathbf{Y}} \log p_{\theta}(\mathbf{Y}^{(N)} | \mathbf{x}), \quad (3)$$

where ρ is an learning rate. Typically, a large number of iterations is required to yield accurate speech samples [3], which is consistent with other related models [15, 18, 19].

2.1. Noise contrastive estimation

Because of the intractable normalising term $\mathbf{Z}_{\theta}(\mathbf{x})$, training approaches that do not involve normalising terms, such as noise contrastive estimation (NCE) [20], could be adopted. As its name suggests, learning in NCE is performed by contrasting positive samples \mathbf{Y}^+ with negative, noisy, samples \mathbf{Y}^- . Specifically, the loss function in NCE is given by

$$\begin{aligned} \mathcal{L}_{\theta}^{\text{nce}}(\mathbf{x}, \mathbf{Y}^+, \mathbf{Y}^-) = & -\log \left(\frac{1}{1 + \exp(E_{\theta}(\mathbf{x}, \mathbf{Y}^+))} \right) \\ & -\log \left(\frac{1}{1 + \exp(-E_{\theta}(\mathbf{x}, \mathbf{Y}^-))} \right), \end{aligned} \quad (4)$$

where \mathbf{Y}^+ is a reference speech and \mathbf{Y}^- is a noisy speech. According to equation (4), optimal energy functions assign high energies to noisy samples and low energies to positive samples. Although NCE does not prescribe which noisy samples to use, the previous work in TTS found it sensitive to the quality of noisy speech [3]. This appears to be different to language modelling where simple negative sampling approaches were found to lead to better results [7]. Although high-quality noisy speech could be obtained by combining outputs from well-trained TTS models [21, 22, 23, 24] and data augmentation methods [3], such an approach would lead to an increased training cost.

2.2. Scores

The gradient of log-likelihood with respect to speech $\nabla_{\mathbf{Y}} \log p_{\theta}(\mathbf{Y} | \mathbf{x})$ is known in the literature as a *score* $\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y})$. Such scores appear not only in EBMs but also in closely related diffusion models where the following inference rule may be used to denoise speech

$$\mathbf{Y}^{(N+1)} \leftarrow \frac{1}{\sqrt{\alpha_N}} \{ \mathbf{Y}^{(N)} + \beta_N \mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}^{(N)}, N) \} + \sigma_N \mathbf{Z}^{(N)}, \quad (5)$$

where $\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}^{(N)}, N)$ denotes a time (iteration) dependent score, $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and α_N , β_N , and σ_N are hyperparameters. The close connection between EBMs in equation (3) and DMs becomes apparent by noticing that $\frac{1}{\sqrt{\alpha_N}}$ is a scale that combines iteration dependent weight decay parameters,

$1 - \frac{1}{\sqrt{\alpha_N}}$ with 1 from the current speech estimate. Thus, setting α_N to 1 and σ_N to 0 (removes stochastic term) would lead to the update rule in equation (3).

In addition to scores not depending on iteration index in EBMs ($\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y})$ vs $\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}, N)$), another interesting aspect of EBM scores is that they can either be modelled directly as in DMs or obtained by computing the gradient of log-likelihood. Figure 1 illustrates these two different approaches. Although more expensive the latter approach ensures that scores are true gradients of some true log-likelihood function.

3. Score-based Training

This section will first discuss score matching approaches and then will introduce a new criterion, which is shown to be linked with currently popular flow matching [16, 25, 26], that addresses some of score matching limitations.

3.1. Sliced score matching (SSM)

Rather than using NCE it is also possible to optimise EBMs by matching their scores to those of the true data distribution. If the true data distribution, $p_{\text{data}}(\mathbf{x}, \mathbf{Y})$ was known then EBM training could be performed by minimising the distance between model score $\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y})$ and data score $\nabla_{\mathbf{Y}} \log p_{\text{data}}(\mathbf{x}, \mathbf{Y})$. As the latter is seldom known, an equivalent, practically feasible, formulation that does not depend on data scores has been found by [27]. It's unbiased estimator is given by

$$\mathcal{L}_{\theta}^{\text{sm}}(\mathbf{x}, \mathbf{Y}^+) = \text{tr}(\nabla_{\mathbf{Y}^+} \mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}^+)) + \frac{1}{2} \|\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}^+)\|_2^2, \quad (6)$$

where the first term is the trace of the Hessian matrix, which is computationally expensive for high-dimensional data [10] even with modern hardware [28]. Hence, direct score matching (SM) is only suitable for simple EBMs.

Fortunately, a computationally efficient variant of SM called sliced score matching (SSM) [10] has been proposed and could be adopted with complex EBMs. Given a random projection vector \mathbf{v} , the loss function of SSM is given by:

$$\mathcal{L}_{\theta}^{\text{ssm}}(\mathbf{x}, \mathbf{Y}^+) = \mathbf{v}^{\top} \nabla_{\mathbf{Y}^+} \mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}^+) \mathbf{v} + \frac{1}{2} \|\mathbf{S}_{\theta}(\mathbf{x}, \mathbf{Y}^+)\|_2^2. \quad (7)$$

Computing the first term in equation (7) requires only 2 gradient operations rather than $\dim(\mathbf{Y}^+)$ gradient operations required in equation (6). Although it is possible to use K random projection rather than one, the single random projection version above is used in this work. Unlike NCE, which requires special care to crafting negative samples, SSM appears to be a straightforward, albeit an expensive, training approach.

3.2. Delta loss (Delta/ Δ)

Apart from requiring $1 + K$ gradient computations, another drawback of SM-style approaches is the focus on matching scores without paying much attention to the underlying log-likelihood. Consider, for example, the log-likelihood on the left-hand side of Figure 2. Even if (S)SM above would manage to perfectly match data scores with model scores it is unlikely that first-order optimisation schemes such as equation (5) would succeed in reaching its global minimum. Log-likelihoods such as those in the middle of Figure 2 are friendlier for first-order schemes though might require careful tuning and large number of iterations. Finally, log-likelihoods on the right-hand




Figure 2: *Different score functions: from left to right, the hypotheses list from the worst to the best.*

side of Figure 2 are optimal as they allow reaching global minima in one step provided correct learning rate has been selected. This suggests that pointing scores of noisy samples \mathbf{Y}^- generated by an acoustic model in the direction of $\Delta(\mathbf{Y}^+, \mathbf{Y}^-) = \mathbf{Y}^+ - \mathbf{Y}^-$ would yield model scores capable of reaching optima (\mathbf{Y}^+) in one step of simple first-order inference scheme

$$\mathbf{Y}^+ = \mathbf{Y}^- + \mathbf{S}_\theta(\mathbf{x}, \mathbf{Y}^-) \quad (8)$$

using any noisy speech sample as initialisation. Such an objective function called *delta loss* in this paper is given by

$$\mathcal{L}_\theta^\Delta(\mathbf{x}, \mathbf{Y}^+, \mathbf{Y}^-) = \frac{1}{2} \|\mathbf{S}_\theta(\mathbf{x}, \mathbf{Y}^-) - (\mathbf{Y}^+ - \mathbf{Y}^-)\|_2^2. \quad (9)$$

Compared to both NCE and (S)SM, delta loss is a very efficient training approach.

3.3. Connection to flow matching

Interestingly, the delta loss in equation (9) can be linked to the Flow Matching (FM) loss [16, 25, 26]

$$\mathcal{L}_\theta^{\text{fm}}(\mathbf{x}, \mathbf{Y}^+, \mathbf{Y}^0, t) = \frac{1}{2} \|\mathbf{V}(\mathbf{x}, \mathbf{Y}_t, t) - (\mathbf{Y}^+ - \mathbf{Y}^0)\|_2^2, \quad (10)$$

where $t \in [0, 1]$, \mathbf{Y}^0 is a source or anchor sample possibly influenced by \mathbf{x} , $\mathbf{Y}_t = t\mathbf{Y}^+ + (1-t)\mathbf{Y}^0$ and $\mathbf{V}(\mathbf{x}, \mathbf{Y}_t, t)$ is a function that transfers samples \mathbf{Y}_t between source/anchor and data distributions. Despite their similarity, the FM loss is motivated by the slow inference of diffusion style approaches [16], whereas delta loss is concerned with model scores being appropriate for first-order inference schemes.

4. Experiments

4.1. Experimental setup

4.1.1. Dataset

The dataset utilized in this work is LJSpeech [29], which consists of 13,100 audio clips totalling approximately 24 hours of speech from a single female speaker reading passages from public domain texts. The dataset is randomly divided into training (10,000 clips), validation (1,800 clips), and test (1,300 clips) sets. A randomly chosen subset of 600 validation set clips was used in some objective evaluations to streamline the process. Another randomly chosen subset of 100 test set clips was used in subjective evaluations.

4.1.2. Models

The pretrained acoustic model generating hypotheses of spectral features is an open-source Tacotron 2¹. These spectral features are used as-is as an initialisation of inference for all

¹<https://github.com/NVIDIA/tacotron2>

EBMs in this paper. For training these hypotheses are either not used (SSM), or adopted as-is (Δ) or additionally transformed to yield higher quality noisy samples (SSM). To ensure that references and hypotheses in training have the same duration (length), the pretrained acoustic model is forced to output spectral feature sequences with reference durations. No such constraint is enforced in inference. The non-score-based EBM trained with NCE is adopted from [3]. The score-based EBMs are based on the widely popular U-net architecture [30]. The specific architecture is identical to Matcha-TTS [24]² other than time/iteration was removed to yield time-independent scores. Each score-based EBM was trained for 50K steps using the Adam optimizer with a learning rate of 10^{-4} and a batch size of 10 on a single NVIDIA 3090 GPU. During inference, the score model is updating using a gradient descent algorithm with a learning rate of 3×10^{-5} . An open-source implementation of HiFi-GAN [31]³ is used as the vocoder to generate waveforms.

4.1.3. Evaluation

The Mel cepstral distortion (MCD), log-scale F0 root mean square error (log f0), SpeechBERTScore [32], and UTMOSv2 [33] are used as objective evaluation metrics. While the former two metrics are well known in the literature, the latter two, SpeechBERTScore and UTMOSv2, are model-based metrics that were previously found to be highly correlated with human perception. An implementation of these objective evaluation metrics is available in the `DiscreteSpeechMetrics` toolkit⁴. Mean Opinion Score (MOS) is used as a subjective evaluation metric to assess the perceived naturalness of synthesized and reference speech. Five predictive models took part in listening tests (1xTacotron 2, 2xSSM, 2xDelta). Three listeners, who are all native English speakers, were asked to rate the perceived naturalness of each waveform on a scale from 1 (bad) to 5 (excellent) on the Amazon Mechanical Turk platform, with the average score representing the overall perceived naturalness.

4.2. Objective evaluation

Table 1 compares the performances of EBMs trained using SSM loss, delta loss or SSM and delta loss to the pretrained acoustic model, Tacotron 2. The number of inference steps, N , for

SSM	Delta	MCD ↓	log f0 ↓	S-BERT ↑	UTMOSv2 ↑
✓	✗	5.312	0.296	0.905	3.710
✗	✓	5.399	0.299	0.900	2.959
✓	✓	5.312	0.296	0.905	3.718
Tacotron 2		5.770	0.300	0.882	3.409

Table 1: *Objective evaluation of score-based training, 100 inference steps, full validation (1800) set*

all EBMs was set to 100. Despite its simplicity, the delta loss trained EBM is competitive to the SSM trained EBM in all metrics other than UTMOSv2. The delta loss also appears to underperform compared to Tacotron 2 in the same metric while being superior based on all others. Combining SSM and delta losses does not appear to yield significant gains.

²<https://github.com/shivammehta25/Matcha-TTS>

³<https://github.com/jik876/hifi-gan>

⁴<https://github.com/Takaaki-Saeki/DiscreteSpeechMetrics>

Table 2 shows how the number of inference steps affects objective evaluation metrics for SSM and delta loss trained EBMs. These results suggest that for SSM loss the best results are ob-

Step	MCD↓	$\log f_o$ ↓	S-BERT ↑	UTMOSv2 ↑
0	5.765	0.297	0.882	3.412
1	5.294	0.299	0.904	3.717
10	5.294	0.299	0.904	3.708
50	5.294	0.300	0.904	3.719
100	5.294	0.300	0.904	3.704

(a) SSM

Step	MCD↓	$\log f_o$ ↓	S-BERT ↑	UTMOSv2 ↑
0	5.765	0.297	0.882	3.412
1	5.292	0.299	0.904	3.703
10	5.277	0.299	0.904	3.674
50	5.273	0.297	0.903	3.400
100	5.389	0.299	0.900	2.937

(b) Delta

Table 2: Objective evaluation of inference for score-based EBMs, partial validation (600) set

tained after just 1 step, which is in contrast to 100s of steps that may be required with NCE-trained EBMs and related diffusion models. Despite its simplicity, the delta loss appears to achieve similar results with also just 1 step. Although additional steps could improve some metrics they appear to lead to a consistent drop in UTMOSv2.

Table 3 compares NCE [3] and score-based trained EBMs. These results suggest that score-based training is a powerful al-

Loss	Step	MCD↓	$\log f_o$ ↓	S-BERT ↑	UTMOSv2 ↑
NCE*	100	5.762	0.295	0.875	2.572
SSM	1	5.294	0.299	0.904	3.717
Delta	1	5.292	0.299	0.904	3.703

Table 3: Objective evaluation of NCE (full validation set) and score-based (partial validation set) trained EBMs

ternative that yields significantly better MCD, S-BERT and UTMOSv2 values but may be insignificantly worse in $\log f_o$.

4.3. Subjective evaluation

Finally, a range of listening tests was conducted to get an insight into score-based EBMs following the protocol outlined in Section 4.1.3. Table 4 summarises the conducted MOS study. The first block in Table 4 shows MOS scores for original reference, ground truth, waveforms (line 1), ground truth waveforms resynthesised from log Mel spectrograms using HiFi-GAN vocoder (line 2) and the pretrained acoustic model, Tacotron 2, used in this paper (line 3). As expected, these results suggest that some level of degradation could be attributed to the imperfect vocoding process and that the pretrained acoustic model lags 0.41 in MOS behind that result. The second block confirms subjectively that the delta loss (line 5) proposed in this paper offers a computationally efficient powerful alternative to a more expensive and elaborate SSM loss (line 4). The delta loss trained EBM reduces the lag in MOS from resynthesised

#	Method	MOS
1	Ground Truth	4.29±0.04
2	Ground Truth (log Mel + HiFi-GAN)	4.15±0.06
3	Tacotron 2 (log Mel + HiFi-GAN)	3.74±0.08
4	SSM (1 steps) (log Mel + HiFi-GAN)	3.90±0.06
5	Delta (1 steps) (log Mel + HiFi-GAN)	3.95±0.05
6	SSM (10 steps) (log Mel + HiFi-GAN)	3.82±0.09
7	Delta (10 steps) (log Mel + HiFi-GAN)	3.76±0.07

Table 4: Subjective evaluation

ground truth to 0.20. The final, third, block shows that although for SSM objective metrics appear to show little discrimination, the additional inference steps lead to a degradation in MOS. For delta loss, MOS and UTMOSv2 appear to be correlated.

The breakdown of MOS score counts in Figure 3 shows that although SSM and delta loss trained EBMs appear to be comparable in terms of the overall MOS scores, there are in fact important differences between these EBM systems. Figure 3

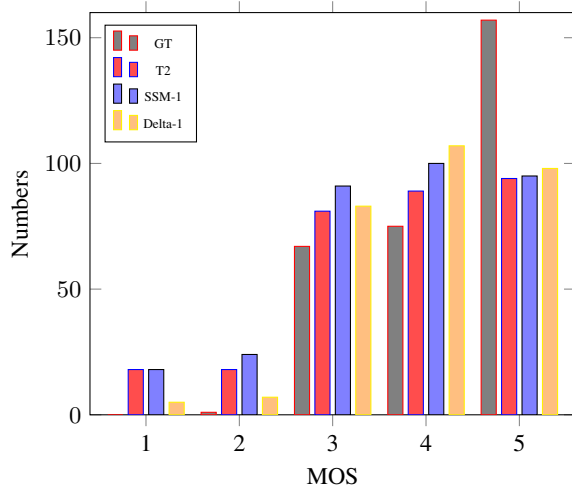


Figure 3: Detailed breakdown of MOS score counts

suggests that the delta loss trained EBM has much fewer 1, 2 and 3 scores and more 4 and 5 scores than SSM trained EBM.

5. Conclusion

This paper presented two new score-based approaches to training energy-based models (EBMs) for text-to-speech (TTS) as an alternative to more challenging Noise Contrastive Estimation (NCE). The first approach is sliced score matching (SSM) that has been previously used with EBMs in other areas but not in TTS. The second approach is a new objective function called *delta loss* which aims to promote log-likelihoods more suitable for first-order inference schemes used by EBMs and is simpler than NCE and SSM. This new loss was shown to be connected to losses used in currently popular flow matching. Experiments on the LJSpeech dataset show superior performance of score-based approaches compared to NCE. Despite its simplicity, the new delta loss function appears to show comparable performance in objective evaluations and better performance in subjective evaluations (less low scores and more high scores).

6. References

- [1] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, F. Huang *et al.*, “A tutorial on energy-based learning,” *Predicting structured data*, vol. 1, no. 0, 2006.
- [2] L. Wenliang, D. J. Sutherland, H. Strathmann, and A. Gretton, “Learning deep kernels for exponential family densities,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 6737–6746.
- [3] W. Sun, Z. Tu, and A. Ragni, “Energy-based models for speech synthesis,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 667–12 671.
- [4] Q. Li, Y. Zhang, B. Li, L. Cao, and P. C. Woodland, “Residual energy-based models for end-to-end speech recognition,” *arXiv preprint arXiv:2103.14152*, 2021.
- [5] S. Bhattacharyya, A. Rooshenas, S. Naskar, S. Sun, M. Iyyer, and A. McCallum, “Energy-based reranking: Improving neural machine translation using energy-based models,” *arXiv preprint arXiv:2009.13267*, 2020.
- [6] A. Bakhtin, Y. Deng, S. Gross, M. Ott, M. Ranzato, and A. Szlam, “Residual energy-based models for text,” *Journal of Machine Learning Research*, vol. 22, no. 40, pp. 1–41, 2021.
- [7] X. Chen, X. Liu, M. J. Gales, and P. C. Woodland, “Recurrent neural network language model training with noise contrastive estimation for speech recognition,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5411–5415.
- [8] A. Hyvärinen and P. Dayan, “Estimation of non-normalized statistical models by score matching,” *Journal of Machine Learning Research*, vol. 6, no. 4, 2005.
- [9] P. Vincent, “A connection between score matching and denoising autoencoders,” *Neural computation*, vol. 23, no. 7, pp. 1661–1674, 2011.
- [10] Y. Song, S. Garg, J. Shi, and S. Ermon, “Sliced score matching: A scalable approach to density and score estimation,” in *Uncertainty in Artificial Intelligence*. PMLR, 2020, pp. 574–584.
- [11] Y. Song and D. P. Kingma, “How to train your energy-based models,” *arXiv preprint arXiv:2101.03288*, 2021.
- [12] K. Kim, A. Jindal, Y. Song, J. Song, Y. Sui, and S. Ermon, “Imitation with neural density models,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 5360–5372, 2021.
- [13] S. Elflein, B. Charpentier, D. Zügner, and S. Günnemann, “On out-of-distribution detection with energy-based models,” *arXiv preprint arXiv:2107.08785*, 2021.
- [14] X. Liu, D. Staudt, C.-T. Lin, and C. Zach, “Effortless training of joint energy-based models with sliced score matching,” in *2022 26th International Conference on Pattern Recognition (ICPR)*. IEEE, 2022, pp. 2643–2649.
- [15] Y. Song and S. Ermon, “Generative modeling by estimating gradients of the data distribution,” *Advances in neural information processing systems*, vol. 32, 2019.
- [16] X. Liu, C. Gong, and Q. Liu, “Flow straight and fast: Learning to generate and transfer data with rectified flow,” *arXiv preprint arXiv:2209.03003*, 2022.
- [17] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-tts: A generative flow for text-to-speech via monotonic alignment search,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 8067–8077, 2020.
- [18] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [19] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [20] M. Gutmann and A. Hyvärinen, “Noise-contrastive estimation: A new estimation principle for unnormalized statistical models,” in *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2010, pp. 297–304.
- [21] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [22] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural speech synthesis with transformer network,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 6706–6713.
- [23] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” *arXiv preprint arXiv:2006.04558*, 2020.
- [24] S. Mehta, R. Tu, J. Beskow, É. Székely, and G. E. Henter, “Matcha-tts: A fast tts architecture with conditional flow matching,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 341–11 345.
- [25] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le, “Flow matching for generative modeling,” *arXiv preprint arXiv:2210.02747*, 2022.
- [26] E. Heitz, L. Belcour, and T. Chambon, “Iterative α -(de) blending: A minimalist deterministic diffusion model,” in *ACM SIGGRAPH 2023 Conference Proceedings*, 2023, pp. 1–8.
- [27] A. Hyvärinen, J. Hurri, P. O. Hoyer, A. Hyvärinen, J. Hurri, and P. O. Hoyer, “Estimation of non-normalized statistical models,” *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*, pp. 419–426, 2009.
- [28] J. Martens, I. Sutskever, and K. Swersky, “Estimating the hessian by back-propagating curvature,” *arXiv preprint arXiv:1206.6464*, 2012.
- [29] K. Ito and L. Johnson, “The lj speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [30] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [31] J. Kong, J. Kim, and J. Bae, “Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Advances in neural information processing systems*, vol. 33, pp. 17 022–17 033, 2020.
- [32] T. Saeki, S. Maiti, S. Takamichi, S. Watanabe, and H. Saruwatari, “SpeechBERTScore: Reference-aware automatic evaluation of speech generation leveraging nlp evaluation metrics,” in *InterSpeech 2024*, 2024, pp. 4943–4947.
- [33] K. Baba, W. Nakata, Y. Saito, and H. Saruwatari, “The t05 system for the voicemos challenge 2024: Transfer learning from deep image classifier to naturalness mos prediction of high-quality synthetic speech,” in *2024 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2024, pp. 818–824.