

A DIFFERENTIABLE PERCEPTUAL AUDIO METRIC LEARNED FROM JUST NOTICEABLE DIFFERENCES

Pranay Manocha¹ Adam Finkelstein¹ Zeyu Jin² Nicholas J. Bryan² Richard Zhang² Gautham J. Mysore²

¹Princeton University

²Adobe Research

ABSTRACT

Assessment of many audio processing tasks relies on subjective evaluation which is time-consuming and expensive. Efforts have been made to create objective metrics but existing ones correlate poorly with human judgment. In this work, we construct a differentiable metric by fitting a deep neural network on a newly collected dataset of just-noticeable differences (JND), in which humans annotate whether a pair of audio clips are identical or not. By varying the type of differences, including noise, reverb, and compression artifacts, we are able to learn a metric that is well-calibrated with human judgments. Furthermore, we evaluate this metric by training a neural network, using the metric as a loss function. We find that simply replacing an existing loss with our metric yields significant improvement in denoising as measured by subjective pairwise comparison.

Index Terms— perceptual similarity, just noticeable difference, deep metric, speech enhancement, user study

1. INTRODUCTION

Humans have an innate ability to analyze and compare sounds. While efforts have been made to emulate human judgment into machine understandable metrics, the gap between human and machine judgment remains open. This is especially true in the context of recent advancements in deep learning [1], where synthetic audio has become so close to real recordings that most metrics fail to reflect human perception. This lack of a perceptually-consistent metric hinders the advancement of audio processing. Furthermore, many deep learning models rely on a metric to construct a loss function; when the loss function is misaligned with human judgment, artifacts are generated.

There exists a small number of objective metrics that, given a reference, evaluate sound quality and are constructed based on human assessment studies, e.g., PESQ [2], POLQA [3], ViSQOL [4]. However, there are two major drawbacks of these methods. First, these models have acknowledged shortcomings such as sensitivity to perceptually invariant transformations [5, 6], which hinders stability in more diverse tasks such as speech enhancement. Second, these metrics are non-differentiable, and thus cannot be directly exploited as a training objective within the context of deep learning.

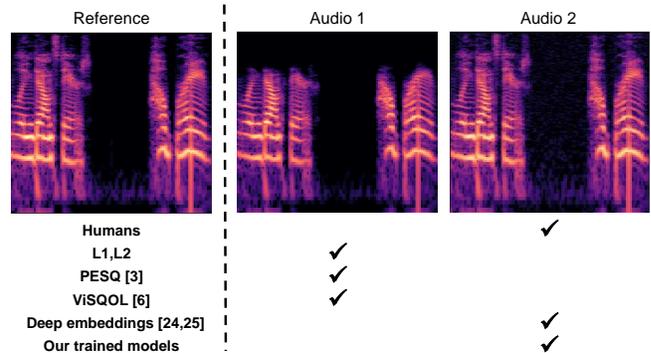


Fig. 1. Is the left or right recording “closer” to the reference? Conventional metrics (e.g. L1,L2) and various audio quality metrics (e.g. PESQ and ViSQOL) struggle to measure JNDs. Our metric, data and code can be found at https://gfx.cs.princeton.edu/pubs/Manocha_2020_ADP/

Another approach is to learn a loss function using adversarial learning. This approach has shown promising results in enhancement [7], synthesis [8], and source separation [9], and have been used for downstream tasks such as denoising [7]. Another line of research takes inspiration from the computer vision community by using representations learned via a different task to construct similarity metrics. It is often called a *deep feature loss* [10] and has been adopted in various audio tasks [11, 12]. However, these approaches are problem-specific [13] and still require human assessment for accurate evaluation, particularly when small perceptual difference need to be measured.

Therefore, we propose a new perceptual audio metric that is better aligned with human judgments based on Just-Noticeable-Differences (JND) – the threshold at which a difference is perceived. To do so, we first collect human judgment JND data via a large scale listening test in which subjects are asked whether two audio clips sound the same or different. We designed our experiment using active learning for more efficient data sampling and inject various perturbations of noises and effects to these clips so that the data has a coverage of potential degradation’s that appear in common audio processing tasks. Then, we train a neural network on the collected data to predict our human-labeled JND annotations and use the learned representation to construct a distance metric that measures how different two audio signals are from

one another. We then validate our metric by 1) showing our metric correlates well with three diverse third-party mean opinion score (MOS) datasets and 2) using our metric as a loss function to train a denoising neural network in which we showed non-trivial improvement from state-of-the-art methods via pairwise comparison listening test.

2. PROPOSED FRAMEWORK

2.1. Data collection methodology

We collect a dataset of human judgments using modern crowdsourcing tools, which have been shown to perform similarly to expert, in-lab tests [14, 15]. We present a listener with two recordings, a reference x_{ref} and perturbed signal x_{per} , and ask if these two audio clips are *exactly same or different*, and record the binary response $h \in \{0, 1\}$.

For the reference recording x_{ref} , we first sample a speech recording from a large collection and then degrade it by randomly applying a set of perturbation (e.g. noise and reverb). To produce the perturbed recording x_{per} , we select a perturbation direction, or “axis” which can be one of several perturbation types or a combination applied sequentially. Figure 2 shows an example where the perturbation direction is a combination of two perturbation types. The types we study are further described in Section 3.1. The perturbed recording x_{per} is produced as a function \mathcal{H} of strength $\rho \in [0, 100]$, $x_{\text{per}} = \mathcal{H}(x_{\text{ref}}, \rho)$.

For values of ρ that are too large or small, the answer is “obviously” different or the same, respectively, and a downstream metric is unlikely to gain information from such data. As such, we employ an active learning strategy to more efficiently gather labelled data, in contrast to past approaches [16]. Our goal is to identify the Just Noticeable Difference (JND) threshold, ρ_{jnd} , such that a subject can *just* hear the difference between x_{ref} and x_{per} . We attempt to sample ρ to be close to the JND point, illustrated at a high-level in Figure 2.

We estimate the current subject’s most likely JND ρ_{jnd}^* , based on all past answers, and then produce the next test case by $x_{\text{per}} = \mathcal{H}(x_{\text{ref}}, \rho_{\text{jnd}}^*)$. We assume that human answers follow a Gaussian distribution with mean μ at the JND point and variance σ^2 , representing human error. Following this, we compute the likelihood of N past answers using $\mathcal{L}(\mu, \sigma^2) = \prod_{j=1}^N (1 - h_j)(1 - c(\rho_j|\mu, \sigma^2)) + h_j c(\rho_j|\mu, \sigma^2)$, where ρ_1, \dots, ρ_N are past perturbation strengths, h_1, \dots, h_N are the human judgments, and $c(v_j|\mu, \sigma^2)$ is the CDF of Gaussian $\mathcal{N}(\mu, \sigma^2)$. After computing μ and σ to maximize the above likelihood function, the next test case follows from $\rho_{\text{jnd}}^* = \mu$.

The ultimate product of our data collection is a database of triplets $\{x_{\text{ref}}, x_{\text{per}}, h\}$, which we leverage for training a perceptual metric.

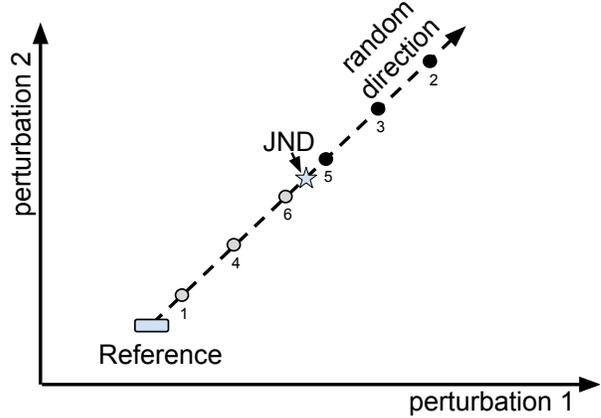


Fig. 2. An illustration of our active learning-based data collection. For a given reference, we probe the listener on one or more perturbation axis (dotted line). Grey circles are *same* and black circles are *different*.

2.2. Training a perceptual metric

A high quality perceptual distance metric D would provide a small distance $D(x_{\text{ref}}, x_{\text{per}})$ if human judges feel they are the same recording, and a larger distance if they are judged to be different. Here, we explore four separate strategies to learn such a metric. We then investigate how well each method correlates with human judgments. All models have the same architecture for comparison, described in Section 3.3.

Using a pre-trained network. “Off-the-shelf” deep network embeddings have been used as a metric for training and have been shown to correlate well with human perceptual judgments in the vision setting [17], even without being explicitly trained on perceptual human judgments. We first investigate if similar trends hold in the audio setting. We describe the activation of layer l of an L -layer deep network embedding as $F_l(x) \in \mathbb{R}^{T_l \times C_l}$, where T_l and C_l are the time resolution and number of channels of the layer, respectively. A distance between two audio clips can be defined by averaging between the full feature activation stack,

$$D(x_{\text{ref}}, x_{\text{per}}) = \sum_l \frac{1}{T_l} \|\mathcal{F}_l(x_{\text{ref}}) - \mathcal{F}_l(x_{\text{per}})\|_1. \quad (1)$$

We train a model (*pre*) on two general audio classification tasks from DCASE 2016 [18], namely acoustic scene classification (ASC) and domestic audio tagging (DAT), following the strategy in [11].

Training a model on perceptual data. We take the above model and add linear weights over the model,

$$D(x_{\text{ref}}, x_{\text{per}}) = \sum_l \frac{1}{T_l} \|w_l \odot (\mathcal{F}_l(x_{\text{ref}}) - \mathcal{F}_l(x_{\text{per}}))\|_1, \quad (2)$$

where $w_l \in \mathbb{R}_l^C$ and \odot is the Hadamard product over channels. The linear weights can decide which channels are more or less “perceptual”. We present two variants. First, we keep the weights of all the layers \mathcal{F} fixed and only train the linear layers. This presents a “linear calibration” of an off-the-shelf network, denoted as *lin*. Second, *fin*, we initialize

from a pre-trained classification model (*pre*), and allow all the weights for network (\mathcal{F} and linear layer) to be fine-tuned. Next, we allow the network \mathcal{F} and the linear layer to both be trained from *scratch*.

3. EXPERIMENTAL SETUP

3.1. Dataset and perturbations

To demonstrate our proposed framework, we apply it to the broad field of speech telecommunication. In this domain, noises like packet losses, jitter, variable delay and other channel noise artifacts like channel noise, and sidetones are common. To simulate these issues, we divide our perturbation set into three categories:

- *Linear perturbations*: include noises like applause, blue noise, brown noise, crickets, pink noise, siren, violet noise, water drop and white noise taken from ESC50 [19].
- *Reverb perturbations*: we use a dataset of real impulse responses (IR) [20] and approximately modify the Direct-to-Reverberant Ratio (DRR) and Reverberation Time (RT60) of the sampled IR by multiplying it with a constant after the first direct response and time stretching respectively.
- *Compression perturbations*: we consider μ -law encoding, where we change the number of bits to encode the audio and MP3 compression, and vary the bit-rate.

For each listening test set, we select at most one instance from each category and sample a random order to apply these categories, *e.g.* white noise energy from linear, DRR from reverb and MP3 bit rate from compression as the three perturbation values (v_1, v_2, v_3). We permute the order to simulate different scenarios. For example, (reverb, linear, compress) simulates telecommunication while (compress, reverb, linear) simulates playback audio in a room environment.

3.2. Crowdsourcing

After determining the perturbation space, we crowdsource JND answers on Amazon Mechanical Turk (AMT). We require workers to have above 95% approval ratings. At the beginning of the Human Intelligence Task (HIT), the subject goes through a volume level calibration test in which loud and soft sounds are played alternatively. The participants are then asked not to change the volume in the middle of the HIT. Next, an attention test is presented where the participant is asked to identify a word heard in a long sentence. This removes participants that either do not understand English or were not paying attention. Upon successfully choosing the right word, the subject goes through two teaching tests, where we train the workers on what kind of differences to look for before we move on to the actual task. Each HIT contains 30 pairwise comparisons, 10 each for one randomly chosen reference and direction. Out of these 30 comparisons, 6 (20%) tests are sentinel questions in the form of obvious audio deformations. If the participant gets any of the 6 questions wrong, we discard their data. Each audio clip is roughly 2.5 seconds long, and the subjects can replay the files if they choose

to. On an average, it takes 7-8 minutes to complete a HIT. At the end, we also ask for comments/suggestions/reviews from the participants on their experience in doing this HIT. We launched 950 HITs and retained 740 after validation, collecting about 22k pairs of human subjective judgments.

3.3. Training and architecture

We use a network inspired by [11] consisting of 14 convolutional layers with 3×1 kernels, batch normalisation and leaky relu units, and zero padding to reduce the output dimensions by half after every step. The number of channels double after every 5 layers starting with 32 channels in the first layer. We also use dropout in all convolutional layers. The receptive field of the network is $2^{14} - 1$. We train the model using cross-entropy loss using a small classification model that maps distance to predicted human judgment.

We train this network for 1000 epochs, taking ≈ 3 days to complete using 1 GeForce RTX 2080 GPU. As part of online data augmentation to make the model invariant to delay, we decide randomly if we want to add a 0.25s silence to the audio at the beginning or the end and then present it to the network. This helps providing shift invariance property to the model, to disambiguate that in fact the audio is similar when time shifted.

4. RESULTS

4.1. Correlation to MOS

We use previously published large-scale third-party MOS studies to see if our trained metric correlates well on their task. We show results of our models, and compare these with embeddings obtained from self-supervised models like OpenL3 [21] and large scale pretrained models like VG-Gish [22] trained on Audioset [23] as well as more conventional objective metrics like MSE, PESQ¹ and ViSQOL². To find correlation, we check Spearman’s Rank order correlation as well as Pearson’s correlation coefficient by evaluating on a per speaker level where we average scores for each speaker for each condition. We choose three distinct classes of publicly available datasets for our correlation analysis:

1. **VoCo** [25]: consists of MOS tests to verify quality of 6 different word synthesis and insertion algorithms, hence misaligned data.
2. **FFTnet** [26]: consists of MOS tests done to evaluate the quality of synthetic audio generated by 5 different type of speech generation algorithms. It introduces artifacts specific to SE (speech enhancement), and may also be misaligned.
3. **Bandwidth Expansion** [27]: consists of MOS tests to verify quality of 3 different bandwidth expansion algorithms. The objective here is to obtain higher audio perceptual features. These audio consist of very subtle local level changes.

¹implementation from [24]

²implementation from <https://qslab.ucd.ie/index.php/speech-quality-metrics/>

Type	Name	VoCo [25]		FFTnet [26]		BWE [27]	
		SC	PC	SC	PC	SC	PC
Ours	Pre	0.60	0.90	0.60	0.90	0.60	0.90
	Lin	0.30	0.45	0.30	0.45	0.30	0.45
	Fin	0.46	0.71	0.30	0.45	0.30	0.45
	Scratch	0.71	0.94	0.63	0.39	0.61	0.47
Self-sup	VGGish	0.10	0.23	-0.41	-0.44	0.51	0.50
	OpenL3	0.27	0.36	0.12	0.17	0.53	0.53
Conv	MSE	0.18	0.80	0.18	0.15	0.00	0.26
	PESQ	0.43	0.85	0.49	0.56	0.21	0.18
	ViSQOL	0.50	0.75	0.02	0.35	0.13	0.09

Table 1. Spearman (SC) and Pearson (PC) correlations for MOS experiments. Models include: ours, (self)-supervised embeddings, and conventional metrics. \uparrow is better.

	1 (Hard)	2 (Medium)	3 (Easy)	4 (Very easy)
Ours > DeepFeatures	56.6	61.3	72.6	71.3
Ours > WaveNet	87.3	82.0	70.6	67.3
Ours > OMLSA	74.6	85.3	82.6	76.6
Ours > Wiener	95.3	99.3	97.3	94.0

Table 2. Denoising pairwise comparison listening test results. Majority vote says our method is better than the baseline. Each row lists a specific *ours vs baseline* experiment. Results are divided into subsets based on difficulty. Chance is 50%.

The result is displayed in Table 1. As we can see our proposed method “**scratch**” has the best performance overall. In addition, there are some notable observations, listed below:

- Neural-network-based metrics are more robust on un-aligned data - we see that all our models, including *pre*, perform better than conventional metrics on the first two datasets in which new speech is synthesized.
- Though the *pre* model learns “time alignment” for free, it has problems on distinguishing between high frequency subtle difference in the third task. It is likely because this model is trained on a task that deemphasizes these frequency bands. On the contrary *VGGish* and *OpenL3* perform relatively well as they are trained on much larger-scaled tasks and thus have a notion of higher frequency perceptual features. However, since these are not trained on speech, they perform worse on the first two tasks.
- Conventional metrics such as PESQ and ViSQOL perform better in the first two cases than the last, indicating they are less accurate when measuring subtle differences.
- Methods like *VGGish* rely on spectrogram difference which is negatively correlated with MOS to begin with which makes recovery of useful information impossible. This shows that (for some cases) losing out on phase information leads to a decrease in audio quality [28], making it one of the disadvantages of using magnitude spectrogram as an input.
- *OpenL3* performs better than *VGGish* across tasks. It is interesting to note that *OpenL3* was trained using self-supervision and *VGGish* was trained using supervised learning.

4.2. Speech enhancement using trained loss

We show utility of our trained metric as a loss function for the task of SE. We use the dataset available in [29], which to our knowledge is the largest available dataset for denoising, consisting of around 11,572 files for training and 824 files for validation. This dataset consists of 28 speakers equally split between male and female speakers containing 10 unique background types across 4 different training SNR’s. More information on this dataset can be found in [29]. Our denoising network (separate from the learned loss network) is a 16 layer fully convolutional context-aggregation network inspired from [11]. We keep the size of the SE model same for fair comparison. We compute losses from all 14 layers of our trained loss function and add them to get the net loss. We use Adam Optimiser [30] with a learning rate of 10^{-4} , and train for 400 epochs on one GeForce RTX 2080 GPU.

We compare our loss function trained SE system with the state-of-the-art method based on deep feature loss [11]. We also compare our result with a few other baselines including Speech Denoising Wavenet [31], OMLSA [32] and Wiener Filter [24]. We randomly select 600 noisy clips from the validation set of [29] and denoise these using the above algorithms. We perform A/B preference tests on AMT, consisting of *Ours vs baseline* pairwise comparisons. Each pair is rated by 15 different turkers and then majority voted to see which method performs better. To analyse these results, we divide our results into 4 subsets: from having highest input noise (Hard) to lowest input noise (Very easy). All results are statistically significant with $p < 10^{-4}$.

Results are shown in Table 2. We observe that our model is more preferred on all subsets across baseline methods. Specifically looking at *ours vs deepfeatures*, we observe that our model does best for high SNR recordings, where the degradation from the background is less noticeable. This highlights the importance of training our perceptual loss on JND data. Given that our loss function is trained on JND data, it is able to better correlate with local, subtle differences than other loss functions.

5. CONCLUSION AND FUTURE WORK

We propose a framework to collect human “just noticeable difference” judgments on audio signals. Directly *learning* a perceptual metric from our data produces a metric that correlates better with MOS tests than traditional metrics, such as PESQ [3]. Furthermore, we show that the metric can be directly *optimized* as a loss function, in the task of speech enhancement. A similar story has emerged in the computer vision literature, where trained networks have been shown to both correlate well with human perceptual judgments [17] and serve well as an optimization objective [33], compared to traditional metrics such as SSIM [34]. In the future, we would like to extend this dataset to other applications like music and voice conversion. Our data and code is released to the public.

6. REFERENCES

- [1] A. Oord and S. Dieleman et al., “WaveNet: A generative model for raw audio,” *arXiv preprint*, 2016.
- [2] A. W. Rix et al., “Perceptual evaluation of speech quality PESQ—a new method for speech quality assessment of telephone networks and codecs,” in *2001 ICASSP*. IEEE.
- [3] J. G. Beerends et al., “Perceptual objective listening quality assessment POLQA, the third generation itu-t standard,” *Journal of AES*, 2013.
- [4] A. Hines et al., “ViSQOL: an objective speech quality model,” *EURASIP*, 2015.
- [5] A. Hines et al., “Robustness of speech quality metrics to background noise and network degradations: Comparing ViSQOL, PESQ and POLQA,” in *2013 ICASSP*. IEEE.
- [6] T. Manjunath, “Limitations of perceptual evaluation of speech quality on voip systems,” in *BMSB*, 2009.
- [7] S. Pascual et al., “SEGAN: Speech enhancement generative adversarial network,” *arXiv preprint*, 2017.
- [8] C. Donahue, J. McAuley, and M. Puckette, “Adversarial audio synthesis,” *arXiv preprint*, 2018.
- [9] D. Stoller et al., “Adversarial semi-supervised audio source separation applied to singing voice extraction,” in *2018 ICASSP*. IEEE.
- [10] L. A. Gatys, A. S. Ecker, and M. Bethge, “A neural algorithm of artistic style,” *arXiv preprint*, 2015.
- [11] F. G. Germain et al., “Speech denoising with deep feature losses,” *arXiv preprint*, 2018.
- [12] I. Ananthabhotla et al., “Towards a perceptual loss: Using a neural network codec approximation as a loss for generative audio models,” in *ACM ICM*, 2019, pp. 1518–1525.
- [13] K. Kilgour et al., “Fr’echet audio distance: A metric for evaluating music enhancement algorithms,” *arXiv:1812.08466*, 2018.
- [14] M. Cartwright et al., “Fast and easy crowdsourced perceptual audio evaluation,” in *2016 ICASSP*. IEEE.
- [15] M. Cartwright et al., “Crowdsourced pairwise-comparison for source separation evaluation,” in *2018 ICASSP*. IEEE.
- [16] D. McShefferty et al., “The just-noticeable difference in speech-to-noise ratio,” *Trends in hearing*, 2015.
- [17] R. Zhang et al., “The unreasonable effectiveness of deep features as a perceptual metric,” in *IEEE CVPR*, 2018.
- [18] A. Mesaros et al., “Detection and classification of acoustic scenes and events: Outcome of the dcase 2016 challenge,” *IEEE/ACM TASLP*, 2018.
- [19] K. J. Piczak, “ESC: Dataset for environmental sound classification,” in *23rd ACM - Multimedia*, 2015.
- [20] J. Traer and J. H. McDermott, “Statistics of natural reverberation enable perceptual separation of sound and space,” *PNAS*, 2016.
- [21] J. Cramer, H. Wu, J. Salamon, and J. P. Bello, “Look, listen, and learn more: Design choices for deep audio embeddings,” in *2019 ICASSP*. IEEE.
- [22] S. Hershey et al., “CNN architectures for large-scale audio classification,” in *2017 ICASSP*. IEEE.
- [23] J. F. G. et al., “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. ICASSP 2017*.
- [24] Y. Hu and Philipos C. L., “Subjective comparison of speech enhancement algorithms,” in *2006 ICASSP*.
- [25] Z. Jin et al., “Voco: Text-based insertion and replacement in audio narration,” *ACM TOG*, 2017.
- [26] Z. Jin, A. Finkelstein, G. J. Mysore, and J. Lu, “FFTNet: A real-time speaker-dependent neural vocoder,” in *2018 ICASSP*. IEEE.
- [27] B. Feng et al., “Learning bandwidth expansion using perceptually-motivated loss,” in *2019 ICASSP*. IEEE.
- [28] H. Purwins et al., “Deep learning for audio signal processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 206–219, 2019.
- [29] C. Valentini-Botinhao et al., “Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks,” in *Interspeech*, 2016.
- [30] Diederik P. Kingma et al., “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [31] D. Rethage, J. Pons, and X. Serra, “A WaveNet for speech denoising,” in *2018 ICASSP*. IEEE.
- [32] I. Cohen and B. Berdugo, “Speech enhancement for non-stationary noise environments,” *Signal processing*, 2001.
- [33] J. Johnson et al., “Perceptual losses for real-time style transfer and super-resolution,” in *ECCV*, 2016.
- [34] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al., “Image quality assessment: from error visibility to structural similarity,” *TIP*, 2004.

7. SUPPLEMENTARY MATERIAL

7.1. Details about framework

Refer to section 2.1. Note that our goal is to measure JND which is to look for ρ_{jnd} that makes the difference between x_{ref} and x_{per} just noticeable. Additionally, we put priors on μ and σ to make the first several tests less susceptible to human error. This has several advantages:

1. **Information Maximization:** one good way to achieve maximum information gain is to ask questions around JND, which is where the answers are the least obvious and most challenging. Also, this strategy has the advantage of inherently creating a “balanced” dataset³ where you have an almost equal number of “same” or “different” answers.
2. **Extra added bias:** We also encourage an equal chance of saying *same* or *different* by using $v_{jnd}^* = \mu + q\sigma$, where $q > 0$ when we have collected more “same” than “different”, and vice versa. This is done so that the participant is likely to break the trend of giving same answers. If the same trend still continues, we discard participants’ data as he is not paying attention.
3. **Additional Priors:** Our model also starts out with a prior that focuses on exploration early on in the test. As more data is acquired, the model becomes more confident and the prior is deemphasized. This procedure (a) stochastically covers a wide range of the sample space and (b) can recover from wrong answers, as participants may provide noisier responses earlier in the test while gaining familiarity.

7.2. Training the perceptual model

7.2.1. Training Objective

Refer to section 2.2 and 3.3. Our network has a small classification network \mathcal{G} at the end of our perceptual distance model \mathcal{F} , which maps this distance $D(x_{ref}, x_{per})$ to a predicted human judgment \hat{h} . We minimize the binary cross-entropy between this predicted value and ground truth human judgment h .

$$\mathcal{L}(G, D) = \text{BCE}(G(D(x_{ref}, x_{per})), h) \quad (3)$$

7.2.2. Training Data

Refer to section 3.1. We focus only on speech audio samples rather than other audio samples (like music instruments etc.). This is because using other audio samples would introduce

multiple sound sources which might have different effects at different levels of perturbations.

Also note that all the audio samples are normalised to constant power before adding external perturbations.

³ N. Roy and A. McCallum - Toward optimal active learning through monte carlo estimation of error reduction - ICML 2001