



A Generative Model for Score Normalization in Speaker Recognition

Albert Swart and Niko Brümmer

Nuance Communications, Inc. (South Africa)

albert.swart@nuance.com, niko.brummer@nuance.com

Abstract

We propose a theoretical framework for thinking about score normalization, which confirms that normalization is not needed under (admittedly fragile) ideal conditions. If, however, these conditions are not met, e.g. under data-set shift between training and runtime, our theory reveals dependencies between scores that could be exploited by strategies such as score normalization. Indeed, it has been demonstrated over and over experimentally, that various ad-hoc score normalization recipes do work. We present a first attempt at using probability theory to design a generative score-space normalization model which gives similar improvements to ZT-norm on the text-dependent RSR 2015 database.

Index Terms: speaker recognition, score normalization, generative modelling

1. Introduction

Speaker recognition researchers have a love-hate relationship with score normalization. Purists regard it as a kludge.¹ Many researchers believe efforts are better spent elsewhere.² However, score normalization is often the only thing that helps in practice to reduce error-rates under non-ideal conditions, such as the data-set shift encountered in the most recent NIST SRE'16 evaluation [1, 2].

Various flavours of score normalization have been published, for example T-norm [3], adaptive T-norm [4], ZT-norm [5], S-norm [6] and adaptive S-norm [7]. The sharp intuition and clever engineering that have contributed to these recipes should not be undervalued—after all, they work and get the job done. However, we feel there is still a lack of theory for designing score normalization solutions. Indeed, it is often regarded as an art rather than a science—see for example “The awe and mystery of T-norm” [8]. The only informal theory seems to be based on the observation that the distributions of scores originating from different test segments (or different enrollments) can sometimes be mutually misaligned. Score normalization seeks to re-align these distributions via scaling and shifting.

In this paper, we make the following contributions. We explain why score normalization is not needed under ideal circumstances, when the recognition model is a good match for the data at runtime. We then motivate why score normalization could help under more challenging conditions, when there is a mismatch between model and data. Finally, we venture into a first attempt at using probability theory to design a new score normalization recipe. We do this by defining a simple generative score-space model that uses hidden variables to induce dependency between the trial-at-hand and some cohort

¹A kludge is a workaround or quick-and-dirty solution that is clumsy, inelegant, inefficient, difficult to extend and hard to maintain.

²Patrick Kenny, in a personal communication: “The score domain is an impoverished domain.”

scores. Once the model has been defined, probability theory does the rest of the job to find the normalization recipe. We conclude with experiments on RSR 2015 to demonstrate that our approach has practical merit.

2. Traditional score normalization

We shall confine ourselves to the canonical speaker recognition problem, where a *trial*, (e, t) , is scored via a function $s(e, t) \rightarrow \mathbb{R}$, where e represents the enrollment speech of some speaker of interest, and t represents the test-speech, which may or may not be of that speaker. The speaker recognition system processes enrollment and test speech to extract the representations e and t . It also implements the scoring function $s(e, t)$. In some systems, such as i-vector PLDA [6, 9], e and t have the same form (often just i-vectors) and the scoring function is symmetric: $s(e, t) = s(t, e)$. In other systems, the representations for e and t differ, which requires an asymmetric scoring function.

Symmetric scoring functions can be normalized with symmetric normalization recipes such as S-norm—while asymmetric normalization recipes, such as T-norm and ZT-norm can be applied to both symmetric and asymmetric scoring functions. In the interest of generality below, we assume asymmetric scoring. For later reference, we briefly summarize two well-known score normalization recipes.

2.1. T-norm

We denote the *trial-at-hand* as (\tilde{e}, \tilde{t}) and the corresponding *raw score*, as $\tilde{s} = s(\tilde{e}, \tilde{t})$. The relatively simple T-norm [3] makes use of a *cohort*, $\mathcal{E} = \{e'_i\}_{i=1}^N$, populated by enrollments, e'_i , of N other speakers, which we assume to be all different from the speaker(s) present in the trial-at-hand. The *cohort scores*:

$$\mathcal{S}_{\tilde{t}} = \{s(e'_i, \tilde{t})\}_{i=1}^N \quad (1)$$

are formed by scoring \tilde{t} against the cohort. The *normalized score* is formed as:

$$s^*(\tilde{e}, \tilde{t}) = \frac{s(\tilde{e}, \tilde{t}) - \mu(\tilde{t})}{\sigma(\tilde{t})} \quad (2)$$

where $\mu(\tilde{t})$ and $\sigma(\tilde{t})$ are respectively the mean and standard deviation of $\mathcal{S}_{\tilde{t}}$.

2.2. ZT-norm

ZT-norm is more complex. It scores the trial-at-hand against two different cohorts and also scores the cohorts against each other. One cohort is populated by enrollments: $\mathcal{E} = \{e'_i\}_{i=1}^N$, and the other is populated by test representations: $\mathcal{T} = \{t'_i\}_{i=1}^M$. Three sets of cohort scores are required: \tilde{e} scored against \mathcal{T} and \tilde{t} against \mathcal{E} :

$$\mathcal{S}_{\tilde{e}} = \{s(\tilde{e}, t'_i)\}_{i=1}^M, \quad \mathcal{S}_{\tilde{t}} = \{s(e'_j, \tilde{t})\}_{j=1}^N,$$

and then also \mathcal{T} against \mathcal{E} for a matrix of *inter cohort scores*:

$$\mathcal{S}_I = \{s(e'_i, t'_j)\}_{i,j=1}^{M,N}$$

The zt-normalized score is computed by sequentially composing two normalization steps of the form (2). See [5, 10] for details. Of interest here, is that the normalized score is a function of \tilde{s} , as well as all the cohort scores, $\mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I$. Adaptive score normalizations use similar score sets as input [4, 7].

3. Theory and motivation

We present a theoretical motivation for doing score normalization, based on conditional independence analysis. The basic question is: *Given the raw score, are the cohort scores also relevant to better infer the speaker hypothesis for the trial-at-hand?*

For generality, let us consider the full set of cohort scores, $\mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}$ and \mathcal{S}_I as used by ZT-norm and adaptive score normalization.

In figure 1, we use graphical model notation [11] to reason about conditional independence relationships. The shaded nodes represent the various observed scores, the cohort scores, as well as the score for the trial-at-hand, \tilde{s} . These scores will be our inputs for computing the normalized score. Hypothesis labels for all cohort scores are assumed given (traditionally all non-targets).




Figure 1: *Graphical model analysis of score normalization.*

At normalization time, the original data for the trial-at-hand \tilde{e}, \tilde{t} and the cohorts, \mathcal{E}, \mathcal{T} are no longer available and are now hidden,³ as indicated by clear circles. The object of the whole exercise is to infer the value of the *hidden speaker hypothesis* for the trial-at-hand:⁴

$$\tilde{h} \in \{\text{target}, \text{non-target}\}$$

Let Θ denote the parameters of a probabilistic generative model that is assumed to have generated all of the data. For now, we consider Θ as given. A hypothesis prior,

$$\pi = P(\tilde{h} = \text{target} \mid \pi)$$

is also given. Let us now apply the rules for dependency in graphical models [11] to see if we can find dependency between \tilde{h} and the cohort scores. That is, we want to know if the computation $P(\tilde{h} \mid \tilde{s}, \pi)$ suffices for inference of \tilde{h} , or if we need *everything* that is given: $P(\tilde{h} \mid \tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \pi, \Theta)$.

³If we consider the data as given, this blocks dependency between \tilde{h} and the cohort scores and score normalization no longer applies.

⁴Under *target*, \tilde{e} and \tilde{t} come from the same speaker. Under *non-target*, they come from different speakers.

Dependency can ‘flow’ along or against any arrow, but can be ‘blocked’ at some nodes, depending on whether the nodes are observed or not. Dependency at any observed node, O , is blocked along paths of the forms $\rightarrow O \rightarrow$ and $\leftarrow O \rightarrow$. Dependency is also blocked along paths of the form $\rightarrow N \leftarrow$ for any node N , if neither N , nor any of its descendants are observed. When these rules are applied to figure 1, we find that: *Yes, there are dependency paths from \tilde{h} to Θ and to all the cohort scores.* At first glance it therefore looks as if we should always be doing score normalization. But this ignores the special nature of the scores that are computed under ideal circumstances.

3.1. Ideal circumstances

Perusal of figure 1 shows that as long as Θ is given, we have:

$$P(\tilde{h} \mid \tilde{e}, \tilde{t}, \tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \pi, \Theta) = P(\tilde{h} \mid \tilde{e}, \tilde{t}, \pi, \Theta) \quad (3)$$

Intuitively, cohorts \mathcal{E} and \mathcal{T} (and any scores computed from them) are independent observations of speakers different from the speaker(s) in the trial-at-hand. If Θ is given, data from the cohorts are not needed to learn more about Θ , and the cohorts will be of no further help in the trial-at-hand. Now consider the *ideal* score, namely the *log-likelihood-ratio*:

$$\tilde{s} = s(\tilde{e}, \tilde{t}) = \log \frac{P(\tilde{e}, \tilde{t} \mid \tilde{h} = \text{target}, \Theta)}{P(\tilde{e}, \tilde{t} \mid \tilde{h} = \text{non-target}, \Theta)} \quad (4)$$

in which case we can apply Bayes’ rule to see:

$$\begin{aligned} P(\tilde{h} = \text{target} \mid \tilde{e}, \tilde{t}, \pi, \Theta) \\ &= \sigma(\tilde{s} + \text{logit } \pi) \\ &= P(\tilde{h} = \text{target} \mid \tilde{s}, \pi) \end{aligned} \quad (5)$$

where $\sigma(s) = \frac{1}{1+e^{-s}}$ and $\text{logit } \pi = \log \frac{\pi}{1-\pi}$. Combining (3) and (5), we find the *normalization killer equation*:

$$P(\tilde{h} \mid \tilde{s}, \pi) = P(\tilde{h} \mid \tilde{e}, \tilde{t}, \tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \pi, \Theta) \quad (6)$$

We have shown that if the scoring function properly computes the likelihood-ratio, using the *same* model parameters that are assumed to have generated all the data, then everything but \tilde{s}, π is irrelevant to inferring \tilde{h} . In this case, the raw score, \tilde{s} , is also the final score.

When i-vector PLDA is trained on large amounts of in-domain data [12, 13], it seems we are close to this ideal, because score normalization usually does not improve accuracy under those circumstances, although calibration may still be required.

3.2. Non-ideal circumstances

It is not hard to come up with excuses for score normalization. The above ideal circumstances are fragile. If, for any reason, the scoring function is different from (4), then the special property (5) does not apply and dependency flows from \tilde{h} through \tilde{e} and \tilde{t} to the cohort scores. There are several reasons why the scoring function can be different from (4), such as:

- The recognition model is too complex and we have to resort to approximate scoring.
- The recognition model is a poor fit to the data.
- The recognition model has been trained on too little data and the parameter estimate is inaccurate.

- There is data-set shift between the training data and the trial-at-hand, where we may encounter, for example, new languages and recording channels.

In our experiments on RSR 2015 below, the first and possibly the second reasons explain the need for score normalization. All of these circumstances mean that, in practice, the true Θ that generated the data is not available for scoring and should therefore be considered hidden. A principled inference for \tilde{h} must then be of the form:

$$P(\tilde{h} \mid \tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \pi) \quad (7)$$

The take-home message is simply that the cohort scores *should* be used in some way.

4. Probabilistic score normalization

Once committed to score normalization, we need to find a way to compute $P(\tilde{h} \mid \tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \pi)$. This generalizes score calibration [14], which merely computes $P(\tilde{h} \mid \tilde{s}, \pi)$. For reasons of simplicity and computational efficiency—just as in calibration—we will be ignoring the real, complex, mechanisms that produce the scores and instead resort to a simple score-space model.

A brute-force solution could discriminatively train a non-linear binary classifier with a cross-entropy criterion. This would be a non-linear logistic regression that processes all of the score inputs to produce a log-likelihood-ratio output score. The disadvantages include: (i) difficulty in designing the functional form, especially for cohorts of variable sizes, (ii) how to avoid specialization to the cohorts used in training, (iii) a large parameter count, which increases vulnerability to over-fitting and to data-set shift.

We choose the generative alternative, where we explicitly model dependencies between the scores and between scores and hypotheses. This has the advantage over the discriminative solution, that once the model has been defined, the functional form for the score can be derived by following the rules of probability theory. Variable sizes of cohorts present no problem. As we will show, we can define a model with very few trainable parameters. The danger of specializing on a given cohort choice at training time is not encountered because there is no notion of a cohort at training time.

One way to induce a general dependency between scores is to simply do Bayesian score calibration [15], where the score calibration parameters are marginalized out. This model, however, is too simple because it has no concept of *trial sides* that induce structured dependencies between subsets of scores.

We decided, instead, to roughly try to emulate the look-and-feel of ZT-norm, with a model that associates a hidden variable, \mathbf{x}_i with every enrollment representation, e_i , and another hidden variable, \mathbf{y}_j , with every test speech representation t_j . The hidden variables can be scalars, or smallish vectors. At training time this model can be represented as in figure 2.

The training data is a K -by- L matrix of scores, $\{s_{ij}\}_{i,j=1}^{K,L}$, with associated hypothesis labels h_{ij} , obtained by scoring K enrollments against every one of L test representations. The model is defined via the joint distribution:

$$\prod_{i=1}^K \prod_{j=1}^L P(\mathbf{x}_i) P(\mathbf{y}_j) P(s_{ij} \mid \mathbf{x}_i, \mathbf{y}_j, h_{ij}) \quad (8)$$

In our experiments below we generalize training to make use of multiple independent score matrices. The training proce-




Figure 2: Normalization model at training time.

dures depends on the model complexity. For the linear-Gaussian model presented below, we use an EM algorithm.

4.1. Scoring

At runtime, the model is rearranged as in figure 3, where the cohort score sets are $\mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I$, as defined above. The hidden variables are represented as $\mathcal{X} = \{\mathbf{x}_i\}$ and $\mathcal{Y} = \{\mathbf{y}_j\}$. The observed data can still be viewed as a rectangular score matrix, where $\mathcal{S}_{\tilde{e}}$ and $\mathcal{S}_{\tilde{t}}$ occupy the same row and column, respectively, as \tilde{s} . We assume all hypothesis labels are given except for \tilde{h} , which must be inferred as:

$$P(\tilde{h} = \text{target} \mid \tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \pi) = \sigma(s^* + \text{logit } \pi) \quad (9)$$

where s^* is the *normalized score*, computed as the log-likelihood-ratio:

$$\begin{aligned} s^* &= \log \frac{P(\tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I \mid \tilde{h} = \text{target})}{P(\tilde{s}, \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I \mid \tilde{h} = \text{non-target})} \\ &= \log \frac{P(\tilde{s} \mid \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \tilde{h} = \text{target})}{P(\tilde{s} \mid \mathcal{S}_{\tilde{e}}, \mathcal{S}_{\tilde{t}}, \mathcal{S}_I, \tilde{h} = \text{non-target})} \end{aligned} \quad (10)$$

Evaluation of this expression requires marginalization over the hidden variables, the complexity of which depends on the model. For the linear-Gaussian model presented below, we find a closed-form solution.




Figure 3: Normalization model at runtime.

4.2. Linear-Gaussian normalization model

To enjoy closed-form training and scoring, we implemented a simple, linear-Gaussian [16] version of the general model described above. The hidden variables \mathbf{x}_i and \mathbf{y}_j are multivariate random variables of dimension D and have independent standard Gaussian priors. In the likelihood, the score means are linear functions of the hidden variables:

$$P(s_{ij} \mid \mathbf{x}_i, \mathbf{y}_j, h_{ij}) = \mathcal{N}(s_{ij} \mid \mu_{ij} + \boldsymbol{\alpha}_{ij}^T \mathbf{x}_i + \boldsymbol{\beta}_{ij}^T \mathbf{y}_j, \sigma_{ij}^2)$$

where all of $\mu_{ij} \in \{\mu_{tar}, \mu_{non}\}$, $\alpha_{ij} \in \{\alpha_{tar}, \alpha_{non}\}$, $\beta_{ij} \in \{\beta_{tar}, \beta_{non}\}$ and $\sigma_{ij} \in \{\sigma_{tar}, \sigma_{non}\}$ agree with $h_{ij} \in \{\text{target}, \text{non-target}\}$. The model therefore has a total of $4 + 4D$ trainable parameters. Notice if $\alpha_{ij} = \beta_{ij} = \mathbf{0}$, the model simplifies to a Gaussian calibration model [14].

In training, the hidden-variable posterior is essential to the EM algorithm. Due to *explaining away*, the hidden variables are dependent in the posterior. If we denote by \mathbf{z} , the vector of stacked hidden variables, \mathbf{x}_i followed by the \mathbf{y}_j , the posterior can be written as the multivariate Gaussian:

$$P(\mathcal{X}, \mathcal{Y} | \mathcal{S}, \mathcal{H}) = \mathcal{N}(\mathbf{z} | \boldsymbol{\mu}_z, \boldsymbol{\Lambda}_z^{-1}) \quad (11)$$

where \mathcal{S} and \mathcal{H} denote all scores and labels and where

$$\boldsymbol{\Lambda}_z = \begin{bmatrix} \mathbf{A} & \mathbf{C} \\ \mathbf{C}' & \mathbf{B} \end{bmatrix}, \quad \boldsymbol{\mu}_z = \boldsymbol{\Lambda}_z^{-1} \begin{bmatrix} \gamma_x \\ \gamma_y \end{bmatrix}$$

where \mathbf{A}, \mathbf{B} are block-diagonal and where $\mathbf{A}, \mathbf{B}, \mathbf{C}, \gamma_x, \gamma_y$ have elements:

$$\begin{aligned} \mathbf{A}_{ii} &= \mathbf{I} + \sum_j \frac{\alpha_{ij} \alpha_{ij}^T}{\sigma_{ij}^2}, & \mathbf{C}_{ij} &= \frac{\alpha_{ij} \beta_{ij}^T}{\sigma_{ij}^2}, \\ \mathbf{B}_{jj} &= \mathbf{I} + \sum_i \frac{\beta_{ij} \beta_{ij}^T}{\sigma_{ij}^2} \\ \gamma_{x,i} &= \sum_j \frac{s_{ij} - \mu_{ij}}{\sigma_{ij}^2} \alpha_{ij}, & \gamma_{y,j} &= \sum_i \frac{s_{ij} - \mu_{ij}}{\sigma_{ij}^2} \beta_{ij} \end{aligned}$$

The model is trained using an EM algorithm with minimum divergence [17].

The runtime scoring formula can be obtained conveniently in terms of the hidden-variable posterior, by application of the *candidate's formula* [18]. The numerator and denominator of (10) are the marginal distributions for all scores ($K \times L$ cohort scores + trial at hand \tilde{s}), \mathcal{S} , given all labels (cohort + hypothesized \tilde{h}), \mathcal{H} :

$$\begin{aligned} \log P(\mathcal{S} | \mathcal{H}) &= -\frac{1}{2} \sum_i^{K+1} \sum_j^{L+1} \left(\frac{(s_{ij} - \mu_{ij})^2}{\sigma_{ij}^2} + \log 2\pi\sigma_{ij}^2 \right) \\ &\quad + \frac{1}{2} \boldsymbol{\mu}_z^T \boldsymbol{\Lambda}_z \boldsymbol{\mu}_z - \frac{1}{2} \log |\boldsymbol{\Lambda}_z|. \end{aligned}$$

In situations where all scores between the cohort and the trial-at-hand are non-targets, we only need to pre-compute and Cholesky-factorize the posterior precision matrix twice, once for each value of \tilde{h} . This gives fast, vectorized scoring that can be applied to large sets of evaluation trials.

5. Experiments

We evaluate our score model on the RSR 2015 corpus [19]. Part I of the corpus consists of 143 female and 157 male speakers speaking 30 prompted phrases during 9 recording sessions. The dataset is split equally into background, development and evaluation sets with non-overlapping speakers. Our text-dependent speaker recogniser is a GMM-UBM system with eigenchannel compensation, similar to [20]. Gender-dependent systems are trained on the background and development sets.

We train a score normalization model on the development set for each gender. Each model is trained on a collection of 30 phrase-specific score matrices to exclude cross-phrase trials. At runtime, the same development sets are used for gender-dependent score normalization cohorts. We also report results only on the same-phrase trials.

It is interesting to note that the hidden variables account for a significantly higher portion of the target score variance ($\nu_{tar} = \alpha_{tar}^T \alpha_{tar} + \beta_{tar}^T \beta_{tar}$) than for non-target scores:

	σ_{non}^2	ν_{non}	σ_{tar}^2	ν_{tar}
Female	23.4	3.5	58.7	55.0
Male	18.9	4.1	55.7	66.3

In Figure 4 we compare raw, un-normalized scores and ZT-norm against our score normalization with scalar (LGSM 1D) and 2-dimensional hidden variables (LGSM 2D). Our score model improves results everywhere relative to raw scores and performs similar or better than ZT-norm in spite of the limitation that our model can only shift scores.

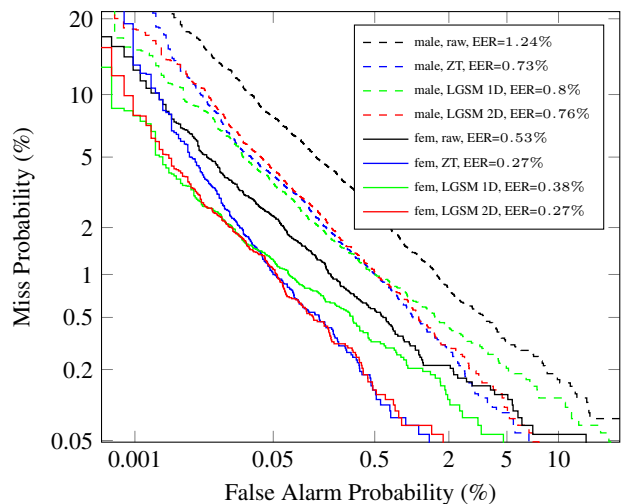


Figure 4: Score normalization results the RSR 2015 corpus

5.1. Calibration

Alert readers may wonder whether this model naturally gives calibrated scores. Since the normalized scores are log likelihood-ratios, this should theoretically be the case. As a sanity check, we found this to hold for synthetically generated data. On RSR data, our normalized scores were better calibrated than raw and ZT-norm scores in the low false-alarm region, but this was not consistent at all operating points. This suggests our model may not have enough capacity to properly model all aspects of real scores. We hope future work, with more sophisticated models, may improve this.

6. Conclusion

We proposed a framework for reasoning about, and designing score normalization models in a principled, probabilistic manner. Our first implementation of a simple linear-Gaussian score model exercises this approach on real data and achieves results comparable to existing score normalization algorithms.

Future work includes developing non-linear score models in which hidden variables can have non-linear effect on the scores.

7. Acknowledgements

We thank the Brno University of Technology Speech Group for helpful discussions during our collaboration in NIST SRE'16.

8. References

- [1] D. Colibro, C. Vair, E. Dalmaso, K. Farrell, G. Karvitsky, S. Cumani, and P. Laface, “Nuance - Politecnico di Torino’s 2016 NIST Speaker Recognition Evaluation system,” in *Interspeech, Stockholm*, 2017.
- [2] O. Plchot, P. Matějka, A. Silnova, O. Novotný, M. Diez, J. Rohdin, O. Glembek, N. Brummer, A. Swart, J. Jorrín-Prieto, P. García, L. Buera, P. Kenny, J. Alam, and G. Bhattacharya, “Analysis and Description of ABC Submission to NIST SRE 2016,” in *Interspeech, Stockholm*, 2017.
- [3] R. Auckenthaler, M. Carey, and H. Lloyd-Thomas, “Score normalization for text-independent speaker verification systems,” *Digital Signal Processing*, vol. 10, pp. 42–54, 2000.
- [4] D. Sturim and D. Reynolds, “Speaker adaptive cohort selection for tnorm in text-independent speaker verification,” in *ICASSP*, 2005.
- [5] R. Vogt, B. Baker, and S. Sridharan, “Modeling session variability in text-independent speaker verification,” in *EUROSPEECH, Lisbon*, 2005.
- [6] P. Kenny, “Bayesian speaker verification with heavy-tailed priors,” in *Odyssey, The Speaker and Language Recognition Workshop, Brno*, 2010.
- [7] S. Cumani, P. D. Batzu, D. Colibro, C. Vair, P. Laface, and V. Vasilakakis, “Comparison of speaker recognition approaches for real applications,” in *INTERSPEECH, Florence*, 2011.
- [8] J. Navrátil and G. N. Ramaswamy, “The awe and mystery of tnorm,” in *EUROSPEECH, Geneva*, 2003.
- [9] D. Garcia-Romero and C. Y. Espy-Wilson, “Analysis of i-vector length normalization in speaker recognition systems,” in *Interspeech, Florence*, 2011.
- [10] S. Shum, N. Dehak, R. Dehak, and J. Glass, “Unsupervised Speaker Adaptation based on the Cosine Similarity for Text-Independent Speaker Verification,” in *Odyssey The Speaker and Language Recognition*, Brno, Czech Republic, 2010.
- [11] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [12] D. Garcia-Romero, A. McCree, S. Shum, and C. Vaquero, “Unsupervised domain adaptation for i-vector speaker recognition,” in *Odyssey, The Speaker and Language Recognition Workshop Joensuu, Joensuu*, 2014.
- [13] S. H. Shum, D. A. Reynolds, D. Garcia-Romero, and A. McCree, “Unsupervised clustering approaches for domain adaptation in speaker recognition systems,” in *Odyssey, The Speaker and Language Recognition Workshop Joensuu, Joensuu*, 2014.
- [14] N. Brümmer, A. Swart, and D. van Leeuwen, “A comparison of linear and non-linear calibrations for speaker recognition,” in *Odyssey, The Speaker and Language Recognition Workshop, Brno*, 2014.
- [15] N. Brümmer and A. Swart, “Bayesian calibration for forensic evidence reporting,” in *INTERSPEECH, Singapore*, 2014.
- [16] S. Roweis and Z. Ghahramani, “A unifying review of linear gaussian models,” *Neural Comput.*, vol. 11, no. 2, pp. 305–345, Feb. 1999.
- [17] N. Brümmer, “A Minimum Divergence Recipe for VBEM,” October 2010. [Online]. Available: <https://sites.google.com/site/nikobrummer/VBEMandMINDIV.pdf>
- [18] J. Besag, “A candidate’s formula: A curious result in bayesian prediction,” *Biometrika*, vol. 76, pp. 183–183, 1989.
- [19] A. Larcher, K. A. Lee, B. Ma, and H. Li, “The RSR2015: Database for text-dependent speaker verification using multiple pass-phrases,” in *INTERSPEECH 2012*, 2012.
- [20] P. Kenny, T. Stafylakis, J. Alam, P. Ouellet, and M. Kockmann, “Joint factor analysis for text-dependent speaker verification,” in *Odyssey, The Speaker and Language Recognition Workshop, Brno*, 2014.