



Large-scale Speaker Retrieval on Random Speaker Variability Subspace

Suwon Shon¹, Younggun Lee², Taesu Kim²

¹Massachusetts Institute of Technology, Cambridge, MA, USA

²Neosapience, Inc., Seoul, South Korea

swshon@csail.mit.edu {yg,taesu}@neosapience.com

Abstract

This paper describes a fast speaker search system to retrieve segments of the same voice identity in the large-scale data. A recent study shows that Locality Sensitive Hashing (LSH) enables quick retrieval of a relevant voice in the large-scale data in conjunction with i-vector while maintaining accuracy. In this paper, we proposed Random Speaker-variability Subspace (RSS) projection to map a data into LSH based hash tables. We hypothesized that rather than projecting on completely random subspace without considering data, projecting on randomly generated speaker variability space would give more chance to put the same speaker representation into the same hash bins, so we can use less number of hash tables. Multiple RSS can be generated by randomly selecting a subset of speakers from a large speaker cohort. From the experimental result, the proposed approach shows 100 times and 7 times faster than the linear search and LSH, respectively.

Index Terms: speaker retrieval, speaker search, hashing, speaker variability subspace

1. Introduction

Today, so to speak, we are living in the flood of online videos and social media. It is known that hundreds of hours of video are uploaded every minute on Youtube. Since most of the online videos include speech segments, collecting and abusing someones voice becomes much easier than ever before. Moreover, recent significant advances in speech synthesis technology made it easier to clone voice identity and it becomes harder to distinguish between original speech and synthesized one [1, 2]. In the area of multimedia retrieval, music and video search technology have been developed to protect the copyright of the contents. However, it is yet difficult to search someone's voice identity from a large amount of multimedia data. Voice identity search would play an important role to prevent illegal cloning and abusing. One might apply conventional speaker identification or verification algorithms for such purpose. It could be done by training the speaker model and inferring the model score for every speech segment available online. But, it is practically not doable because of the large search space. Thus, it is inevitable to develop faster algorithms without performance degradation to prompt to a user as quickly as possible [3].

After audio segment can be represented in a single low-dimensional latent vector such as i-vector [4], hashing approach was proposed on the speaker search problem to maximize the speed of retrieval while minimizing the performance degradation [5, 6, 7, 8]. Locality Sensitive Hashing (LSH) [9] is a powerful tool to approximate the nearest neighbor search in high dimension and successfully adopted on top of the i-vector for a large-scale fast speaker identification.

For a large scale application, long hash bits are needed to reduce the size of each hash bin and to speed up a retrieval time. However, since the LSH is data-independent unsupervised hashing method rely on the random projection, it suffers from the redundancy of the hash bits. Also, long hash bits decrease the collision probability between similar samples. Consequently, we are encouraged to use a lot of multiple hash tables and this increases the query time and storage consumption. To reduce this redundancy, a supervised hashing method can be considered. A spectral hashing [10] is a representative method in a supervised hashing method. It produces a very efficient hash code by learning a hash function from the training data. Ideally, all the true neighbor of an arbitrary query can be found in the specific hash bucket. However, this is not feasible using only a single hash table for a practical situation such as large-scale data lies on the high-dimensional space.

In this study, we propose a Random Speaker-variability Subspace (RSS) projection approach based on the Linear Discriminant Analysis (LDA) to give maximum efficiency on the hash table of LSH. Providing weak supervision by selecting random subset speakers in the training dataset when we generate random projection matrix for LSH, we observed it could dramatically save the hash bits (hyperplanes) and hash tables by removing the redundancy. We also re-define the speaker search task into more detail sub-task as speaker identification task and speaker retrieval task considering the query and search space. Experiments were conducted using both i-vector and neural network based speaker embedding, x-vector [11]. We used Voxceleb 1 and 2 dataset [12, 13] and prepared a various combination of dataset categories to measure the performance of speaker search sub-tasks.

2. Speaker search in the wild

Speaker search can be divided into two sub-categories as speaker identification and speaker retrieval. The main difference between two sub-tasks is whether the query is from arbitrary speech segment or target speaker segment.

2.1. Speaker Identification

For speaker identification, the query x is a latent vector from an arbitrary speech segment and its class (speaker index) is C_x . Search space is $D = \{\omega_1, \omega_2, \dots, \omega_S\}$ where S is total number of Speaker and ω_s is a latent vector from a speaker speech segment, s . y_s is the score between x and ω_s . For closed-set speaker identification which query speech is assumed to be spoken by one of the speakers in the search space, the system forced the input query to match the best matching speaker in the search space as $C_x = C_{y^*}$ where $y^* = \max_{s=1 \dots S} y_s$. Previous study on large-scale speaker identification mostly supposed this

closed-set identification to measure the performance [5, 6].

However, in practical situation, the system needs to consider that the speech query may not be spoken by the speaker in the search space. This can be represented as open-set speaker identification. Open-set speaker identification is also known as multi-target detection [14], so we can borrow the top-1 stack detector performance measurement term in [15] about miss and false alarm probability.

2.2. Speaker Retrieval

For speaker retrieval, the query ω is a latent variable from a speech segment of the target speaker we are interest in and its class is C_ω . The search space is $D = \{x_1, x_2, \dots, x_N\}$ where N is the total number of segments to search and x_n is a latent vector from an arbitrary speech segment from an unknown speaker. Compared to the speaker identification task, the query and search space are changed each other. The aim of this task is to retrieve all relevant documents. Thus, speaker retrieval is 1 to N verification task. The performance measurement is the same as the 1 to 1 verification task which is a general speaker verification task. However, the trials to measure the performance need to be prepared by full combination between the query and search space. For instance, if we have 10 target speakers for an input query, the number of trials is $10 \times N$.

3. Speaker Dataset

To make up speaker search environment, we used Voxceleb 1 and 2 dataset [12, 13]. The Voxceleb dataset is composed of automatically collected audio and video data for large scale speaker identification. The collecting pipeline includes face detection and active speaker verification to verify all video clips have detected faces with synchronized speaking voice by the conservative threshold to maximize the precision. These processes are reasonable pre-processing to our speaker search scenario, so we used the audio part of the Voxceleb dataset without modification. Voxceleb 1 and 2 have more than 1,281,352 utterances from 7,365 speakers. Utterances were extracted from a video clip and each clip has roughly 10 to 50 utterances. We used Voxceleb 1 development set (147,935 utterances from 1,211 speakers) for training on both identification and retrieval task. For the test, we combined the utterances in a different way for each task.

Speaker identification task: For search space in the speaker identification task, we used the utterances from first 3 video clips in Voxceleb 1 test set, plus entire utterances from Voxceleb 2 development set. We will use all utterances from the same speaker at once for speaker representation, so each entity of search space represent a unique speaker and include total 6,034 speakers. For the query, we used the rest of the utterances in the Voxceleb1 test set. Since we already picked the utterances from first 3 clips for search space, all the query must have the same identity in search space and this means closed-set identification task as previous studies [5, 6].

Speaker retrieval task: For search space, we used 100,000 utterances from Voxceleb 2 development set and 3,776 utterances from the first 3 clips in Voxceleb 1 test set. We randomly selected 100,000 utterances in the Voxceleb 2 development set from 1,092,009 utterances, so total 103,776 utterances became the search space. For the query, we used utterances from first 3 clips in Voxceleb 1 test set (total 40 speakers, a target class which has relevant utterance in search space) and Voxceleb 2 test set (total 120 speakers, a non-target class which does not

Table 1: *Speaker verification performance by speaker representation method on Voxceleb1 test set*

Speaker representation	EER
i-vector (cosine)	8.1
i-vector (PLDA)	5.4
x-vector (cosine)	9.9
x-vector (PLDA)	6.0

have relevant utterances in search space). Thus, 160 speaker representations were used as query.

We mainly focused on the speaker retrieval task, but we also conducted an experiment on speaker identification task as the previous study reported [6].

4. Fast large-scale speaker retrieval

4.1. Speaker recognition system

For performance comparison, we considered two speaker representation, i-vector [4] and x-vector [11]. To train i-vector and x-vector, we used the voxceleb 1 development set as depicted in the section 3. We followed other training detail same as [16]. Table 1 shows the speaker verification performance of i-vector and x-vector using Equal Error Rate (EER) measurement. We used the same system as reported in [16]. Note that since the Voxceleb 1 contains only 1,211 speakers which is a small number to training DNN, x-vector shows slightly worse performance than i-vector.

4.2. LSH

The goal of our study is that to retrieve similar speakers efficiently with minimum performance loss. In the previous study [6], LSH was adopted for this task and showed very promising performance on the large-scale speaker identification task combining with i-vector.

LSH projects the i-vector into random hyperplane. This hash operation maps close vectors into the same bins(i.e. buckets) with high probability [17]. Suppose r is d -dimension random projection vector drawn from a standard normal distribution, d is dimension of original i-vector w . The hash function maps i-vector w as :

$$w_r = h_r(w) = \text{sgn}(w^T r) = \begin{cases} 1 & \text{if } w^T r \geq 0 \\ 0 & \text{if } w^T r < 0 \end{cases} \quad (1)$$

We can concatenate several hash functions and use multiple, independent hash functions to boost performance, by using $d \times k$ dimensional random projection matrix R_l where k is the number of hyperplane per hash table, d is dimension of original speaker representation, l , $1 \leq l \leq L$, is the hash tables index and L is the total number of hash tables. Thus, the parameter k and L need to be chosen carefully since the parameter decides the trade-off between performance and computational efficiency.

In this study, we basically follow the same i-vector retrieval with LSH algorithm in [6] and also used cosine distance which guarantees its performance in conjunction with LSH. We measured the baseline performance of closed-set speaker identification and speaker retrieval tasks using i-vector and x-vector as shown in table 2.

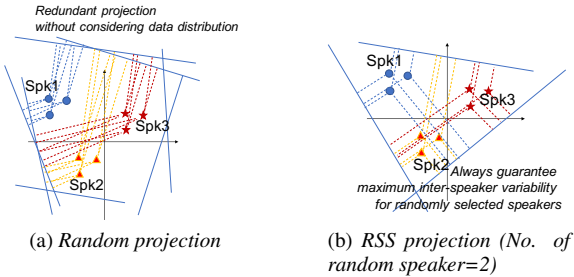


Figure 1: Example of projections for 3 speakers. Same color represents different utterances for same speaker.

Table 2: Speaker search performance using *i*-vector and *x*-vector

Speaker Representation	Speaker retrieval, EER(%)	Speaker identification, Acc.(%)
<i>i</i> -vector	4.36	79.23
<i>x</i> -vector	5.60	76.74

4.3. Random Speaker-variability Subspace (RSS) Projection

Since LSH is an unsupervised hashing method that uses random projection matrix drawn from a d -dimensional standard normal distribution, it suffers from redundancy of the hyperplanes (hash bits) as represented in figure 1. Thus, lots of hash tables are needed to access enough points for the satisfactory recall. In this situation, we can give a weak supervision to generate the random subspace projection matrix using speaker labels to make the same speaker’s voice mapped into the same bin more efficiently. We generate Linear Discriminant Analysis (LDA) transformation matrix directly by using utterances from the random subset of the speakers and use this matrix for projection matrix R_l .

Suppose S_l is a randomly selected speaker subset from training dataset where the number of speaker is N_s . Using the speaker representation in the subset S_l , we can obtain the between-class scatter matrix and the within-class scatter matrix. Then, we can calculate LDA matrix which maximizes the ratio of the two scatter matrices. This LDA matrix projects speaker representation into RSS and we can generate it L times randomly to substitute the random projection matrix R_l in LSH. For hyperparameter N_s , if we choose too many speakers, it will generate a similar LDA matrix which have many redundancy between projection matrices. Also we have to choose N_s more than the length of hash bits k to project into k dimension. To balance the size of each hash bin, we modified the hash function as $h_r(w) = \text{sgn}(w^T r + b)$ where b is the mean of projected data, i.e, $b = -\frac{1}{N} \sum_{i=1}^N w_i^T r$

Similar work was reported in the face recognition study [18] that uses random sampling LDA. They generated multiple projection matrix by selecting eigenvectors randomly. Then LDA was applied on each projected subspace. Since they have only d eigenvectors to be selected randomly, the redundancy between the projected subspaces would be increased if we increase the number of hash tables. This approach is not actually designed for hashing. However, projecting samples onto random subspace to enforce weak classifiers is similar scheme to our approach, so we also use this random sampling LDA to generate projection matrix R_l and compare it with other approaches.




Figure 2: Average Hamming distance of same identities divided by the distance of different identities. Lower is better.

In the projected subspace, hamming distance approximates the cosine distance as the number of random hyperplanes k increases [19]:

$$\cos(w_i, w_j) \approx \cos\left(\frac{H(h_r(w_i), h_r(w_j))}{k}\pi\right) \quad (2)$$

where $H(\cdot)$ is a hamming distance. Thus, if the projection matrix effectively approximates the original distance, the speaker representation from the same identity has high possibility to be in the same or near bucket and we can use less number of hash functions to approximate the original distance. It means that the hamming distance of speaker representation from the same identity would be closer than others. Taking this into consideration, the approximation capability of a new projection matrix can be measured by the averaged hamming distance of the same identities and different identities. We checked this averaged hamming distance using LSH, random sampling LDA, and the proposed method. As shown in figure 2, the proposed method shows that the distance between the same speakers is close, but the different speakers result in far distance. From this observation we can expect the proposed method would give significant efficiency than others.

5. Experimental results

We conducted two tasks of experiments, speaker identification and speaker retrieval, and the experimental environments were fully described in the section 3. For speaker retrieval, EER was used for performance measurement, for the speaker identification task (closed-set), accuracy was used for as described in section 2. *i*-vector and *x*-vector were extracted in 600 and 512 dimensions respectively, then reduced into 150 dimensions using LDA. We measured the retrieval time to return candidates for a given query and excluded the time to extract *i*-vector and *x*-vector. For the baseline, we generate L random projection matrix which drawn from d -dimensional standard normal distribution. For RSS projection, the optimal operating N_s is varied by the number of hyperplanes. Rather than optimizing, we set to d , equals to *i*-vector or *x*-vector dimension, so the subset has more than d samples at least to avoid within-class scatter matrix become singular by small sample size problem [20]. We repeated L times to generate RSS projection matrix by selecting random speakers for each matrix. For the random sampling LDA [18], we used randomly 100 eigenvectors from PCA for random subspace, then repeat L times to generate random sampling LDA matrix. We used these three projection matrix R_l for LSH.

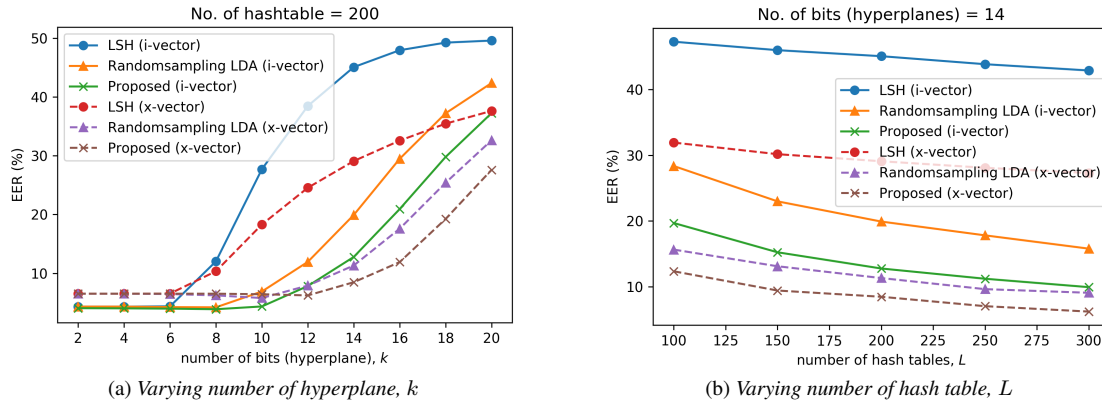


Figure 3: EER measurement on speaker retrieval task.

Figure 3 (a) and (b) shows the result on various number of hyperplanes and hash tables. The proposed RSS projection shows significant efficiency compared to others in the EER measurement. Meanwhile, it is interesting that the x-vectors mapped into hash bins more efficiently than i-vectors in all methods. We surmise that this phenomenon is because the x-vector extracting DNN was trained discriminatively with one hot speaker label while the i-vector framework assumes that the i-vector distributed in Gaussian distribution. Thus, the distance between the same speaker x-vectors is more likely to become close than the i-vector. We believed that any speaker embeddings extracted from DNN would take this advantage, not only x-vector. Figure 4 (a) and (b) shows trade-off plot between retrieval speed and performance. We conducted the experiments by varying the parameter k and L and scattered in speed and performance axis. On both speaker retrieval and identification task, the proposed approach shows remarkable performance improvement. For example, to speed up while maintaining above 95% speaker identification performance relative to linear search, the proposed approach has 100 times faster than linear search and 7 times faster than LSH as shown in table 3. Note that the EER in figure 3 and figure 4 (a) is absolute value, not relative value to linear search. To give an intuition on the parameters, we specified pairs of parameters on the plot in $[L, k]$ format by varying a parameter while the other was fixed. As shown in the figure, k is more sensitive to performance than L . Thus, L can be used for fine-tune to satisfy required performance.

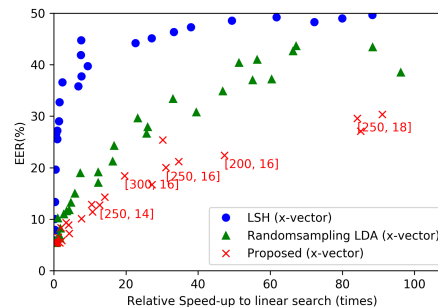
Table 3: Performance summary of hashing methods on speaker identification

Hashing	LSH	Random-sampling LDA	Proposed
Baseline Accuracy		76.74 %	
Hashing Accuracy	74.10%	74.14%	74.65%
Relative speed	14×	7×	100×
No. of Hyperplanes	10	6	12
No. of Hash tables	300	150	150


6. Conclusion

We have proposed a RSS based hash algorithm to search and retrieve someone’s voice identity, which is applicable to a large scale speech dataset. Previous studies have focused on the speaker identification, but we have redefined the search problem

both in the speaker identification and retrieval. The proposed approach has shown significant efficiency to save retrieval time and storage consumption compared to “vanilla” LSH that uses random projection. We have also observed that the speaker embedding is more advantageous to be mapped into a hash table compared to traditional i-vector. To the best of our knowledge, this is the first study to use the class supervision on hashing for acoustic information retrieval. For future work, we would explore speaker retrieval methods further in many other domains as well as real-world applications. It would be also interesting to investigate how the system works in the presence of cloned voice by recent speech synthesis algorithms.



(a) Speaker retrieval task. Y-axis represent EER(%) value.



(b) Speaker identification task. Y-axis represent accuracy relative to linear search accuracy.

Figure 4: Speedup vs. performance trade-off. The number of k varied from 2 to 20 in step of 2 and L from 100 to 300 in steps of 50. A pair of numbers on the plot represent $[L, k]$ as examples.

7. References

- [1] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomyriannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards End-to-end Speech Synthesis," *ArXiv e-prints arXiv:1703.10135*, 2017.
- [2] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. L. Moreno, and Y. Wu, "Transfer Learning from Speaker Verification to Multispeaker Text-To-Speech Synthesis," *ArXiv e-prints arXiv:1806.04558*, 2018.
- [3] H. Aronowitz and D. Burstein, "Efficient Speaker Identification and Retrieval," in *Interspeech*, 2005, pp. 2433–2436.
- [4] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-End Factor Analysis for Speaker Verification," *IEEE Trans. on Audio, Speech, and Lang. Process.*, vol. 19, no. 4, pp. 788–798, may 2011.
- [5] W. Jeon and Y.-M. Cheng, "Efficient Speaker Search over Large Populations using Kernelized Locality-sensitive Hashing," in *ICASSP*, 2012, pp. 4261–4264.
- [6] L. Schmidt, M. Sharifi, and I. Lopez Moreno, "Large-scale Speaker Identification," in *ICASSP*, 2014, pp. 1669–1673.
- [7] R. Leary and W. Andrews, "Random projections for large-scale speaker search," pp. 66–70, 2014.
- [8] D. Sturim and W. M. Campbell, "Speaker linking and applications using non-parametric hashing methods," in *Interspeech*, 2016, pp. 2170–2174.
- [9] P. Indyk and R. Motwani, "Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality," *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pp. 604–613, 1998.
- [10] Y. Weiss, A. Torralba, and R. Fergus, "Spectral Hashing," *Nips '08*, no. 1, pp. 1–8, 2008.
- [11] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, and Y. Carmiel, "Deep Neural Network Embeddings for Text-Independent Speaker Verification," in *Interspeech*, 2017, pp. 999–1003.
- [12] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Interspeech*, 2017, pp. 2616–2620.
- [13] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep Speaker Recognition," in *Interspeech*, 2018, pp. 1086–1090.
- [14] S. Shon, N. Dehak, D. Reynolds, and J. Glass, "MCE 2018: The 1st Multi-target Speaker Detection and Identification Challenge Evaluation," in *Interspeech*, 2019.
- [15] E. Singer and D. Reynolds, "Analysis of Multitarget Detection for Speaker and Language Recognition," in *ODYSSEY The Speaker and Language Recognition Workshop*, no. 4, 2004, pp. 301–308.
- [16] S. Shon, H. Tang, and J. Glass, "Frame-level Speaker Embeddings for Text-independent Speaker Recognition and Analysis of End-to-end Model," in *IEEE Spoken Language Technology Workshop (SLT)*, 2018.
- [17] A. Gionis, P. Indyk, and R. Motwani, "Similarity Search in High Dimensions via Hashing," *VLDB '99 Proceedings of the 25th International Conference on Very Large Data Bases(VLDB)*, vol. 99, no. 1, pp. 518–529, 1999.
- [18] X. Wang and X. Tang, "Random sampling LDA for face recognition," *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, vol. 2, pp. 259–265, 2004.
- [19] A. Jansen and B. Van Durme, "Indexing Raw Acoustic Features for Scalable Zero Resource Search," *Conference of the International Speech Communication Association*, vol. 4, pp. 2466–2469, 2012.
- [20] R. Huang, Q. Liu, H. Lu, and S. Ma, "Solving the small sample size problem of LDA," in *International Conference on Pattern Recognition*, vol. 3, 2002, pp. 29–32.