

asya: Mindful verbal communication using deep learning

Pre-print Version, August 20, 2020

Evalds Urtans
Riga Technical University
Riga, Latvia
evalds@asya.ai

Ariel Tabaks
University of Central Lancashire
Preston, Lancashire, UK
ariel@asya.ai

ABSTRACT

asya is a mobile application that consists of deep learning models which analyze spectra of a human voice and do noise detection, speaker diarization, gender detection, tempo estimation, and classification of emotions using only voice. All models are language agnostic and capable of running in real-time. Our speaker diarization models have accuracy over 95% on the test data set. These models can be applied for a variety of areas like customer service improvement, sales effective conversations, psychology and couples therapy.

CCS CONCEPTS

- Theory of computation → Design and analysis of algorithms;
- Applied computing;

KEYWORDS

Deep Learning, Triplet loss, ConvNet, ResNet, DenseNet, Mel-Spectra, Speaker diarization, Emotion detection, NLP

1 INTRODUCTION

asya is a mobile application that listens to a person's voice and provides private feedback on a person's verbal communication. It gives metrics on how much a person listens and speaks, how long are a person's sentences and utterances, how fast a person speaks, how positive is the tone of a person, the confidence level of a person's voice based on the tone and other metrics. It is deployed as a web service and stand-alone application. These models are being applied to a variety of tasks starting from customer service evaluation to analysis of private conversations and couple's relationships therapy.

Neuroscientists Andrew Newberg, M.D., and Mark Waldman, have identified through brain scans and from other studies that if everyday verbal interactions is coupled with increased moment-to-moment awareness the results can lead to increased levels of trust building, resolved conflicts, increased intimacy and other benefits [3]. The findings show that people can benefit from speaking less, shorter, and slower as human brain short-term memory holds only about four "chunks" of information, which translates to speaking time under 30 seconds [3]. Furthermore, when people practice the 30 second rule, they can train themselves to increase awareness of

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

<https://asya.ai>, August 20, 2020, Riga, Latvia

© 2018 Copyright held by the owner/author(s).

filtering out lower quality information. In addition, when speaking is kept with brevity in mind, the emotional centers of the brain that can be triggered by certain words are less likely to lead the speaker to sabotage the conversation. The research also shows that exercising awareness not only helps to increase connection with other people, but also suppresses the brain's ability to generate feelings of anxiety, irritability or stress.

The tone of voice, emotions, and the way how a person speaks are as important as the content that a person speaks. For example, one could imagine how one could say the phrase "you are such a fool" in a way that could offend someone or in a way that could be even playful and funny.

2 RELATED WORK

In the last decade, there has been great progress in natural language processing (NLP) models to do the tasks like "text-to-speech", "speech-to-text", translation, semantic and syntactic understanding of language [21] [23] [1] [9]. However, only in recent years there has been emerging research to use deep learning models to analyze human's voice biometric features that are linked to psychology [12] [5]. Before it has been done using classical machine learning models that yielded in lesser results [13]. Voice features that typically are used for analysis are Mel-frequency filter banks, Log Mel spectrograms, Mel-frequency cepstrum (MFCC), or even raw waveforms in combination with audio envelopes. Mel-frequency filter banks are filters that are applied to spectrum calculated by Fast Fourier Transform (FFT) to simulate specific amplitudes of sounds at different frequencies that are audible to the human ear. The audio spectrum is much broader than that what human ear can perceive, but other frequencies of sound are less likely to contain useful information for NLP tasks.




Figure 1: Mel-frequencies filter bank to emulate human's ear perception.

Recently, also some progress has been done into recognizing emotions from a person's voice. Historically, there have been very little datasets available for voice alone, but recent techniques using transfer learning enabled to accumulate considerable datasets with unsupervised learning to analyze emotions in a human's voice [2].

Some of the recent works have used deep Convolutional Networks (ConvNet) [10] to extract features from human’s voice spectra and classify 8 basic emotions: Happiness, Sadness, Anger, Fear, Disgust, Surprise, Boredom and Neutrality. Even with the basic ConvNet model, it has been possible to surpass human reference accuracy on detecting emotions in a person’s voice. For example, humans on average were able to detect happiness in voice with 65% precision, whereas ConvNet model was able to detect it with 100% precision [19], [15].




Figure 2: Mel-spectrogram of a person speaking (processed input of asya models).

3 METHODOLOGY

asya uses deep learning models that take as an input Mel spectrograms and other features from raw voice recordings. Deep Residual Networks (ResNet) [7] and DenseNet [8] models have been applied as feature encoders. ConvNet models are deep artificial neural network models that have very similar results when experimentally compared to human retina natural neural networks [11]. At first layers of the model, they extract basic features like Gabor patches and edges, but the deeper they go they extract more general features. For example, for face detection task, first they would detect features like nose and eyes, but then in deeper layers faces as a whole. These models are very deep with usually more than 32 layers. ConvNet models with residual connections (ResNet) allow error to flow freely using back-propagation algorithm without vanishing gradient problem. In case of DenseNet, there are even more connections and better flow of gradient of error through the model. asya models have been trained on multiple large private datasets from different speakers and languages using DenseNet models and other proprietary models.




Figure 3: Schematic illustration of ConvNet model for speaker feature embedding using Mel-spectrogram sample.

3.1 Speaker diarization

Speaker diarization (identification of different speakers in parts of utterances) has been done primary using i-vector, d-vector [17], x-vector [18] based models.

More recently, RNN based models also have been applied like UIS-RNN [22].

It has also been done using triplet loss or contrastive loss and cosine similarity of embedding vectors (fingerprint vectors of human

voice) [12] [5] [16]. asya uses Exponential Triplet Loss function and clustering of speaker embedding to achieve speaker diarization and speaker re-identification in one step [20].

asya models are utterance (phrase) and language independent, whereas, for example, Google Home recognize speaker by specific phrase like "Ok, Google". It means that asya models are capable of identifying a person’s voice at any point in natural conversations. To improve training results of speaker diarization models data have been split into multiple parts. asya has been trained as a set of hierarchical models that first predict if the audio in a given window is a noise or speech, then if it is a man or woman and finally does feature embedding of person’s voice. During the testing, we use the center of the mass of a person’s voice embedding vector to estimate a probability of voice sample belonging to a particular person.




Figure 4: Histograms of cosine distance between speaker embedding vectors acquired by asya male models. Left in the beginning of training, right after training have been completed. Green are distances between samples form same speaker. Red are distances between samples form different speakers.



Figure 5: Representation of speaker embedding vectors from asya models in 3D space (Spherical PCA). Colors denote different speaker samples in test data set.




Figure 6: Segmentation of audio sample for speaker diarization task for 2 male speakers. Red frames are furthest away from speaker feature embedding vector and green frames are closest thus identifying target speaker ¹.

3.2 Emotion classification

Speaker emotion classification task is also done using the same deep learning feature extractor as for speaker diarization task, but with softmax loss function at the end [19].

Standard academic datasets of classified emotional states of audio are read to validate approaches. For example, traditional German EmoDB dataset [6] contains only 500 samples of 10 speakers.

The methodology for our work initially has been based on EmoVox-Celeb [2]. EmoVoxCeleb is trained on FERPlus [4] dataset of still pictures of human FER (Facial Expressions) in different emotional states and then applied to a larger VoxCeleb dataset of videos [14]. These emotional states are classified as Paul Ekman’s 8 basic emotions.

Even though we can achieve state-of-art results in academic data-sets we had to create our own proprietary methodology and dataset to reach similar performance in production systems.

To acquire more training, data transfer learning and unsupervised learning have been used to scrape public data from video sites in the internet.

Table 1: Preliminary results of emotion classification using asya models

emotion	Human [14]	EmoVoxCeleb	asya
happiness	84	35	62
sadness	81	71	75
anger	97	72	75
fear	84	35	75
surprise	81	36	69
disgust	67	67	75
neutral	93	N/A	69

4 RESULTS

Asya models are currently in development and are being tested using mobile application in natural conversations to improve couples relationships through conversations in a similar manner, how does couple therapy would work.

Set of hierarchical models are executed in real-time in less than 500 ms. for every 1 sec. frame on consumer grade GPU server. It is also possible also to deploy and execute these models on flagship mobile phones with machine learning specialized processing units.

Table 2: Preliminary results of speaker re-identification and other voice classification tasks of asya models

model	accuracy
noise detection	99.2
gender detection	89.3
male speaker re-identification	87.4
female speaker re-identification	88.4

5 CONCLUSIONS

The proposed models are capable of analyzing human’s voice in real-time. asya is able to detect noise in audio samples and process only parts with a human’s voice. asya models are able to detect a human’s perception of the gender of the speaker with high precision. Finally, it is able to produce a unique embedding vector for each person’s voice to combine speaker diarization and reidentification tasks in a single step. Furthermore, asya models are able to detect Ekman’s basic human emotions from language and utterance independent data.

There can be a wide range of use cases for asya models. It has been successfully deployed to improve communication skills and encourage mindful communication in a commercial product https://asya.ai. Asya is being developed also to improve public speaking skills and provide feedback for psychologists about their sessions with patients. Asya models also can be used to monitor customer experience in customer service-centered businesses like phone hotlines, post offices, stores, telemarketing, etc. Finally, they could be used to identify persons of interest in large databases of audio recordings, but there are even more use cases than listed in this paper.

REFERENCES

- [1] Karim Ahmed, Nitish Shirish Keskar, and Richard Socher. [n. d.]. Weighted Transformer Network for Machine Translation. abs/1711.02132 ([n. d.]). arXiv:1711.02132 <http://arxiv.org/abs/1711.02132>
- [2] Samuel Albanie, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. [n. d.]. Emotion Recognition in Speech Using Cross-Modal Transfer in the Wild. In *2018 ACM Multimedia Conference on Multimedia Conference, MM 2018, Seoul, Republic of Korea, October 22-26, 2018* (2018). 292–301. <https://doi.org/10.1145/3240508.3240578>
- [3] Andrew Newberg M.D and Mark Robert Waldman. [n. d.]. *Words Can Change Your Brain: 12 Conversation Strategies to Build Trust, Resolve Conflict, and Increase Intimacy.*
- [4] Emad Barsoum, Cha Zhang, Cristian Canton Ferrer, and Zhengyou Zhang. [n. d.]. Training Deep Networks for Facial Expression Recognition with Crowd-Sourced Label Distribution. In *ACM International Conference on Multimodal Interaction (ICMI)* (2016).
- [5] Hervé Bredin. [n. d.]. TristouNet: Triplet Loss for Speaker Turn Embedding. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2017, New Orleans, LA, USA, March 5-9, 2017* (2017). 5430–5434. <https://doi.org/10.1109/ICASSP.2017.7953194>
- [6] Felix Burkhardt, Astrid Paeschke, Melissa A Rolfes, Walter F. Sendlmeier, and Benjamin Weiss. [n. d.]. A Database of German Emotional Speech. In *INTERSPEECH* (2005).
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. [n. d.]. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016* (2016). 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [8] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. [n. d.]. Densely Connected Convolutional Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017* (2017). 2261–2269. <https://doi.org/10.1109/CVPR.2017.243>
- [9] Takuhiro Kaneko and Hirokazu Kameoka. [n. d.]. Parallel-Data-Free Voice Conversion Using Cycle-Consistent Adversarial Networks. abs/1711.11293 ([n. d.]). arXiv:1711.11293 <http://arxiv.org/abs/1711.11293>

- [10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. [n. d.]. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105. <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>
- [11] Jonas Kubilius, Stefania Bracci, and Hans P. Op de Beeck. [n. d.]. Deep Neural Networks as a Computational Model for Human Shape Sensitivity. 12, 4 ([n. d.]). <https://doi.org/10.1371/journal.pcbi.1004896>
- [12] Chao Li, Xiaokong Ma, Bing Jiang, Xiangang Li, Xuewei Zhang, Xiao Liu, Ying Cao, Ajay Kannan, and Zhenyao Zhu. [n. d.]. Deep Speaker: An End-to-End Neural Speaker Embedding System. abs/1705.02304 ([n. d.]). arXiv:1705.02304 <http://arxiv.org/abs/1705.02304>
- [13] Benoît Mathieu, Slim Essid, Thomas Fillon, Jacques Prado, and Gaël Richard. [n. d.]. YAAFE, an Easy to Use and Efficient Audio Feature Extraction Software. In *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR 2010, Utrecht, Netherlands, August 9-13, 2010* (2010). 441–446. <http://ismir2010.ismir.net/proceedings/ismir2010-75.pdf>
- [14] Arsha Nagrani, Joon Son Chung, and Andrew Senior. [n. d.]. VoxCeleb: A Large-Scale Speaker Identification Dataset. In *INTERSPEECH (2017)*.
- [15] Yafeng Niu, Dongsheng Zou, Yadong Niu, Zhongshi He, and Hua Tan. [n. d.]. A Breakthrough in Speech Emotion Recognition Using Deep Retinal Convolution Neural Networks. abs/1707.09917 ([n. d.]). arXiv:1707.09917 <http://arxiv.org/abs/1707.09917>
- [16] Florian Schroff, Dmitry Kalenichenko, and James Philbin. [n. d.]. FaceNet: A Unified Embedding for Face Recognition and Clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015* (2015), 815–823. <https://doi.org/10.1109/CVPR.2015.7298682>
- [17] David Snyder, Daniel Garcia-Romero, Daniel Povey, and Sanjeev Khudanpur. [n. d.]. Deep Neural Network Embeddings for Text-Independent Speaker Verification. In *INTERSPEECH (2017)*.
- [18] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur. [n. d.]. X-Vectors: Robust DNN Embeddings for Speaker Recognition. ([n. d.]), 5329–5333.
- [19] Somayeh Shahsavari. [n. d.]. Speech Emotion Recognition Using Convolutional Neural Networks.
- [20] Evalds Urtans, Agris Nikitenko, and Valters Vecins. [n. d.]. Exponential Triplet Loss. In *ICDDA (2020)*.
- [21] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. [n. d.]. WaveNet: A Generative Model for Raw Audio. In *The 9th ISCA Speech Synthesis Workshop, Sunnyvale, CA, USA, 13-15 September 2016* (2016). 125. http://www.isca-speech.org/archive/SSW_2016/abstracts/ssw9_DS-4_van_den_Oord.html
- [22] Aonan Zhang, Quan Wang, Zhenyao Zhu, John W. Paisley, and Chong Wang. [n. d.]. Fully Supervised Speaker Diarization. ([n. d.]), 6301–6305.
- [23] Cong Zhou, Michael Horgan, Vivek Kumar, Cristina Vasco, and Dan Darcy. [n. d.]. Voice Conversion with Conditional SampleRNN. In *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*. (2018). 1973–1977. <https://doi.org/10.21437/Interspeech.2018-1121>