

Adversarial Black-Box Attacks on Automatic Speech Recognition Systems using Multi-Objective Evolutionary Optimization

Shreya Khare¹, Rahul Aralikkatte², Senthil Mani¹

¹IBM Research

²University of Copenhagen

skhare34@in.ibm.com, rahul@di.ku.dk, sentmani@in.ibm.com

Abstract

Fooling deep neural networks with adversarial input have exposed a significant vulnerability in the current state-of-the-art systems in multiple domains. Both black-box and white-box approaches have been used to either replicate the model itself or to craft examples which cause the model to fail. In this work, we propose a framework which uses multi-objective evolutionary optimization to perform both targeted and un-targeted black-box attacks on Automatic Speech Recognition (ASR) systems. We apply this framework on two ASR systems: Deepspeech and Kaldi-ASR, which increases the Word Error Rates (WER) of these systems by upto 980%, indicating the potency of our approach. During both un-targeted and targeted attacks, the adversarial samples maintain a high acoustic similarity of 0.98 and 0.97 with the original audio.

1. Introduction

Advancements in deep learning have improved the state-of-the-art systems in domains like computer vision, natural language processing and speech recognition. Recent studies [1, 2, 3] have shown that these systems can be easily fooled with carefully crafted inputs. For example, [4, 5] showed that an image classifier could be fooled into classifying an image to a label of choice by introducing small perturbations to the input image which are imperceptible to humans. ASR systems are becoming ubiquitous with the pervasiveness of smart devices. With the success of digital assistants, voice is replacing text as the main mode of communication with these devices. Along with automating mundane human activities, these digital assistants can now accomplish complex tasks which may require sensitive user data. For example, Alexa [6] can now access a user's bank details and perform actions like checking the balance and making credit-card payments. Therefore, it is essential that ASR systems are not susceptible to adversarial attacks. Attacks on ASR systems can broadly be classified based on two attributes: i) model transparency, and ii) intent of the attack.

Attacks based on model transparency: Attacks in which the adversary has access to the internals of an ASR system (network structure, weights, etc.) are known as *White Box* attacks [2]. Attacks of this kind are rare as commercial ASR systems seldom expose their internal working to users. On the other hand, *Black Box* attacks, where the attacker only has access to the input-output pairs of the model, are more probable. *Attacks based on intent:* The intention of an adversarial attack can be either to cause failure or manipulate the system towards an end-goal. The former is known as *un-targeted* attacks where the intent is to make the ASR system generate wrong output for a given audio sample. In the latter, known as *targeted* attacks, the aim is to perturb the input in such a way that the ASR outputs a desired text. While un-targeted attacks just degrade an ASR system's performance, targeted attacks are dangerous as they can manip-

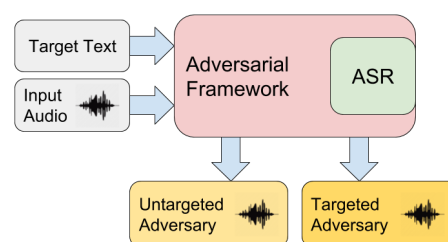


Figure 1: Overview of the adversarial example generation framework for ASR systems

ulate actions taken by digital assistants.

Previous works like [7] generate adversarial samples for ASR systems by adding perturbations to MFCC features and perform lossy reconstruction of the speech signal in the time domain. This lossy reconstruction makes the perturbations perceptible to the human ear. [8] proposes a method which enables the propagation of gradients to the MFCC reconstruction layer. They perform white box targeted attacks on Mozilla's Deepspeech model [9] and obtain adversarial samples which are 99.9% similar to the desired target text. Furthermore, [10, 11] utilise non-linearities of microphones in smart devices to generate spurious commands (adversarial examples with frequency 40 Khz.) and issue malicious commands to digital assistants. Another kind of attack proposed in [12], intends to generate sounds that are interpreted as voice commands by devices but are unrecognisable to humans. These attacks, although powerful, are usually easy to detect and can be prevented using filters. [13] use a genetic algorithm to perform black-box attacks, with an intent of mis-classifying audio samples containing short speech commands. [14] extended this approach to work on longer phrases and sentences. Their approach is limited to ASR systems which give access to the last layer (logits) and requires the knowledge of the model's loss function. Hence it can be considered as a *grey-box* attack at best.

In this work, we consider adversarial audio generation as a multi-objective optimization problem which achieves a balance between conflicting objectives: (i) increasing text dissimilarity, and (ii) maintaining acoustic similarity. We propose a framework which is algorithm agnostic and generates near-optimal solutions in both un-targeted and targeted settings. We show that our approach works well with two genetic optimizers: (i) a simple Multi-Objective Genetic Algorithm (MOGA) [15], and (ii) Elitist Non-dominated Sorting Genetic Algorithm (NSGA-II) [16]. We perform black-box attacks on two popular ASR systems: Deepspeech [9] and Kaldi-ASR (TDNN chain model) [17], and present both empirical and human evaluation results ¹

¹The project web page is at <https://shreyakhare.github.io/audio-adversarial/>

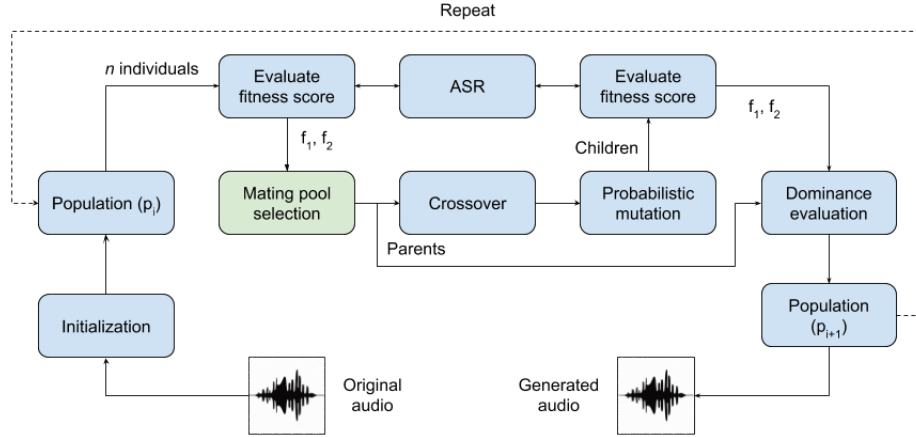


Figure 2: System Architecture

2. Approach

Multi-objective optimization is a process of finding solutions which must satisfy more than one conflicting objectives. In such cases, often there is no single optimal solution. The interaction among different objectives gives rise to a set of compromised solutions known as trade-off, non-dominated, non-inferior or *Pareto-optimal* solutions [18].²

Before moving forward, we define the following terms: (i) *dominance*: a is said to dominate b iff a is no worse than b in every fitness objective (blue points in Fig. 3). If this constraint does not hold between a and b , then they are called non-dominating individuals, (ii) *pareto front*: a set of mutually non-dominating individuals all of which are pareto-optimal (individuals on the red line in Fig. 3). There can be multiple pareto fronts in the search space, and (iii) *crowding distance*: an individual is said to have high crowding distance if there are many individuals in its vicinity. It indicates the density of the area around the individual in the search space. For more details, the reader is advised to refer [19].

As with any evolutionary system, in each generation our framework scores candidate solutions to determine which individuals are the fittest. They are allowed to reproduce using crossover and mutation which results in fitter individuals (better solutions) over time. The approach to generate adversarial samples is shown in Fig. 2. The blue boxes represent algorithm agnostic parts of the framework. Only the *mating pool selection* changes with the choice of algorithm.

Evaluation of fitness score is arguably the most important step in evolutionary algorithms. It decides which individuals survive and therefore prunes the search space. Our goal is to generate audibly similar samples which maximize errors in ASR systems. Parameterizing a fitness function to evaluate this criterion is hard. Therefore we breakdown the fitness function into two components, each of which quantifies a desired property: (i) acoustic similarity – the Euclidean distance between MFCC [20] of the original and generated audio samples, and (ii) text dissimilarity – *edit distance* between the texts generated by the ASR system when the original and generated samples are provided as inputs. To the best of our knowledge, we are the first to use such a multi-objective fitness function to optimize an adversarial audio generator.

²We use the terms ‘Solution’ and ‘Individual’ interchangeably

In the un-targeted setting, the generated sample is considered fitter if it has high acoustic similarity with the input audio and if the text generated has high edit distance with respect to the ASR output. Whereas in the targeted setting, a fit solution, while having high acoustic similarity with the input audio, should also have low edit distance when compared with the desired text.

Once fitness values are computed for every candidate solution, we use an ensemble of techniques for **mating pool selection**. For MOGA, we define three selection schemes: (i) *Same Rank Selection*: In this selection scheme, we first create two lists l_1 and l_2 by ranking the individuals based on each of the two fitness criteria mentioned above. Let $l_x \leftarrow i_{x_r}$ where i is an individual in list x with rank r . We pair two ‘ i ’s which have the same r in the two lists, (ii) *Inverse Rank Selection*: This scheme is the opposite of the previous one. One list is ranked from best to worst and the other vice-versa. This makes sure that diversity is maintained in the population and we do not constrain the search space too soon, and (iii) *Roulette Wheel Selection* [21]: In this scheme, the fitness score is used to compute a probability of selection for each individual. If f_i is an individual i ’s fitness score, then its probability of being selected is:

$$p_i = \frac{f_i}{\sum_{i=1}^N f_i} \quad (1)$$

where N is the number of individuals in the population. While unlikely, individuals with high fitness scores may still be eliminated ensuring there is no premature convergence to local optima.

Now, the NSGA-II selection procedure can be described as follows: (i) select the k best individuals and move them to the next generation without change, (ii) sort the other individuals by their dominance and pick $N - k$ individuals, and (iii) if candidates of one pareto front are dominating, use crowding distance to re-sort the individuals and maintain diversity

Once the mating pairs are identified using the above selection strategies, we use arithmetic recombination operators to perform a single-point **crossover** [22]. If (p_1, p_2) is a mating pair, then three children (C) are created where one has the characteristics of the parents in equal proportion, and two are dominated by each parent.

$$C = \frac{p_1 + p_2}{2} \cup \frac{2 \times p_1 + p_2}{3} \cup \frac{p_1 + 2 \times p_2}{3} \quad (2)$$

Table 1: Texts generated by ASR systems on randomly selected audio samples and their adversarial counterparts

Attack	Original text	Generated Text (original)	Target Text	Generated Text (attacked)
Un-targeted	the one you are blocking	the one you are blocking	N.A.	the money of locking
Un-targeted	never mind about that	never mind about that	N.A.	no i about that
Un-targeted	what's your name he asked	whats yo name he asked	N.A.	one of late ask
Targeted	i've got to go to him	ive got to go to him	a cat	a cat
Targeted	follow the instruction here	follow the instruction here	all of these	all of these shapes
Targeted	do you know what	did you know that	that I love you	I love

Thereafter, we perform **mutation** on each parameter of every individual by adding Gaussian noise with a probability $prob_m$ so that the new population is scattered over a large area in a different part of the search space thus reducing crowding distance. Once a set of offspring is created, their fitness scores are evaluated in the same way as described previously. The combined set of parents and children are ranked using (i) dominance scores in MOGA, (ii) dominance and crowding distance in NSGA-II. The top N individuals are carried forward to the next generation. This process is repeated until we reach termination.

3. Experiments and Results

We use the standard Mozilla Common Voice dataset (CVS) [23] to test the effectiveness of our approach. Following [8, 14], we take 100 instances from CVS and use them to perform both un-targeted (make the ASR output as erroneous as possible) and targeted (make the ASR generate a desired output) attacks. Table 1 shows some examples of the original text generated by the ASR systems (on the original audio sample), the target text (the end goal) and the text generated from the adversarial input. In this section, we present our implementation details along with evaluation methods. We evaluate our approach in two ways: i) *quantitative* – where we empirically prove the effectiveness of the samples using WER and correlation coefficient (CC), and ii) *qualitative* – where we show that humans find it difficult to distinguish between the original and adversarial samples by conducting a survey.

3.1. Implementation Details

For each audio sample from the dataset, a set of 100 individuals are created to form the initial population. They are initialized by adding random uniform noise to the original audio signal with

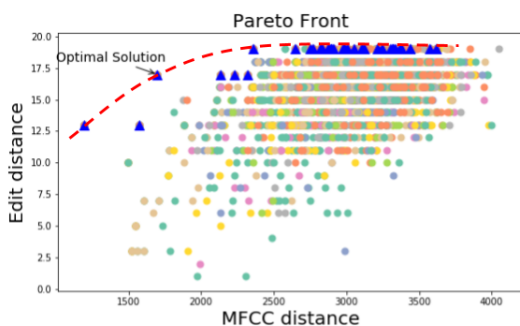


Figure 3: A pareto front showing the dominance of individuals for a random audio sample

a sampling rate of 16,000. Therefore, a two-second audio clip will result in individuals having 32,000 genes. Intuitively, this helps in faster convergence when compared to random initialization.

After computing the fitness scores based on acoustic (MFCC) and textual (edit distance) similarities, the population is converted to a list of mating pairs. The extraction of MFCC features is performed with window size of $25ms$ and a stride of $10ms$. It is to be noted that an individual may appear in more than one mating pair. The three crossover operations described previously are applied to each pair to produce three children. Probabilistic mutation is applied to every gene of every individual with $prob_m = 0.005$.

Next, the parents and children are combined and re-ranked based on their dominance (and crowding distance in case of NSGA-II) and the top 30 pareto-optimal individuals are propagated to the next generation. This process is repeated for a $maxiters = 50$, or until we get the same set of pareto-optimal solutions over two successive generations which indicates convergence. The most dominant individual from the last generation is the adversarial counterpart for our input audio signal. After termination, there may be multiple, equally good dominant solutions to pick from (with equal dominance). In this case, we choose one at random as the best solution.

3.2. Quantitative Evaluation

For quantitative evaluation, we first compute the WER of the original samples using an ASR System. For the selected 100 samples the WER of Deepspeech is 0.5 and Kaldi-ASR is 1.0. After both un-targeted and targeted attacks are performed, the new WER and CC are obtained for the adversarial samples which can be seen in Table 2. For Deepspeech (Kaldi-ASR), the WER degrades by 980% (368%) during un-targeted attacks.

Table 2: Mean word error rates and correlation coefficients of the ASR systems on original and adversarial samples. MOS stands for multi-objective selection. (\uparrow - higher the better, \downarrow - lower the better).

Attack	Metric	ASR	MOS+ MOGA	MOS+ NSGA
Untargeted	WER \uparrow	Kaldi-ASR	3.9	4.68
		Deepspeech	5.4	3.7
	CC \uparrow	Kaldi-ASR	0.97	0.97
		Deepspeech	0.98	0.98
Targeted	WER \downarrow	Kaldi-ASR	3.33	3.33
		Deepspeech	2.14	1.5
	CC \uparrow	Kaldi-ASR	0.98	0.97
		Deepspeech	0.98	0.97




Figure 4: Word Error Rate comparison between human and ASR transcribed text on adversarial samples

For targeted attacks, the target text snippets are generated as follows: if the word length of the original sample text is n , then a phrase whose length is the range of $[2, n + 1]$ is chosen randomly from the data to form the target text. In this situation, the WER between the original and targeted text will be in the range $(n, n + 1) \approx (5, 6)$. In this setting, the text generated by DeepSpeech (Kaldi-ASR) for the adversarial audio samples, are on average 1.0 (1.83) words away from their target texts. This suggests that even after only 50 generations, the generated samples are quite close to their target. Table 2 also shows that average CC is extremely high for both ASR systems, thus indicating acoustic similarity between the original and adversarial samples.

To compare our results with a previous state-of-the-art, we use the model proposed in [14]. This *grey-box* system (it requires access to the model’s loss function and final layer logits) can only perform targeted attacks. Using the same 100 audio samples and target texts, the mean WER of this system is observed to be 5.0 and the CC is 0.99. Compared to this, our framework generates adversarial samples which are much closer to their targets while having comparable acoustic similarity.

3.3. Qualitative Evaluation

While quantitative evaluation empirically proves the goodness of our framework using WER and CC, it is necessary to determine if humans can be fooled with this approach. For this purpose, a qualitative evaluation has been carried out by conducting subjective listening tests with 27 participants.

These tests comprised of questions for evaluating an audio sample on 4 aspects: (i) comprehensibility of content, (ii) listening effort, (iii) presence of perturbations, and (iv) naturalness of generated samples. Each participant is presented with 24 random audio samples (from both targeted and un-targeted attacks) and is required to transcribe the text and rate the listening effort required (from 1-4, with 4 denoting highest effort) to comprehend the text. This provides an indirect measure of the noise/aberrations present in the samples. The WER is calculated for the transcribed text and compared with the ASR generated text, with the original text being the ground truth. We obtained 225 data points to infer our insights from. Fig. 4 shows the WER for the two ASR systems and human transcribed text. The mean WER for the human transcribed text is 1.5 which clearly sug-

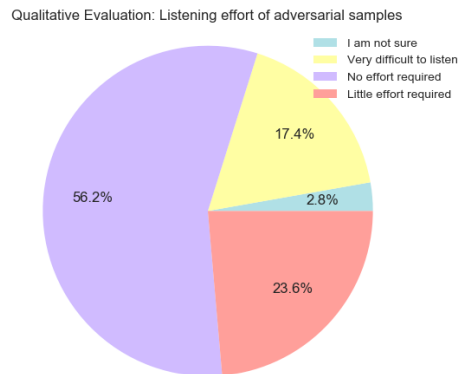


Figure 5: Listening effort required by the survey participants to understand and transcribe the adversarial samples

gests that humans could not find any significant differences between the adversarial and the original samples.

Fig. 5 depicts the listening effort required by the participants to comprehend the audio. More than 80% of the participants required little or no effort to comprehend the text, indicating that there was either no or minimal audible noise in the generated adversarial samples.

To evaluate the naturalness of the generated samples, we asked the participants to distinguish between the original and adversarial samples. They were able to correctly identify both original and adversarial samples 50% of the time which is equivalent to a random guess. To verify if humans can determine if the generated audio sample has been manipulated, the participants were asked to listen to the generated samples and determine if they were artificially generated or altered in some manner. 73% said the audio sounded natural, whereas 27% said that they could hear some ticks in the background which made them believe that the audio could be synthesized.

Lastly, we conducted a *transferability* experiment to study if the samples generated to fool one ASR system can do the same to another system. The adversarial samples generated for DeepSpeech in an un-targeted setting were tested on Kaldi-ASR and vice versa. In this setting, the mean WER of Kaldi-ASR and DeepSpeech are 1.9 and 3.6 respectively. This indicates that though there are some common weaknesses between the two ASR systems which can be exploited, there are other model specific weaknesses which may be a result of the model’s inductive bias or bias introduced by the training data. This is an interesting research direction to pursue in the future.

4. Conclusion and Future Work

In this work, we introduce an algorithm agnostic framework for attacking ASR systems using evolutionary multi-objective optimization. The framework is tested on two ASR systems and is used for both un-targeted and targeted attacks. Our fitness function decomposes adversarial audio quality evaluation into two objectives: Euclidean distance of MFCC features to get the acoustic similarity between audio samples and edit distance to measure the generated text similarity. We show that MOGA and NSGA-II can be plugged into the framework to generate adversarial samples with high WER and CC. To the best of our knowledge, this is the first time a multi-objective fitness criterion is used to optimize auditory and textual features jointly.

5. References

- [1] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Srndic, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," *CoRR*, vol. abs/1708.06131, 2017. [Online]. Available: <http://arxiv.org/abs/1708.06131>
- [2] N. Carlini and D. A. Wagner, "Towards evaluating the robustness of neural networks," *CoRR*, vol. abs/1608.04644, 2016. [Online]. Available: <http://arxiv.org/abs/1608.04644>
- [3] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *CoRR*, vol. abs/1312.6199, 2013. [Online]. Available: <http://arxiv.org/abs/1312.6199>
- [4] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and Harnessing Adversarial Examples," *ArXiv e-prints*, Dec. 2014.
- [5] X. Yuan, P. He, Q. Zhu, R. R. Bhat, and X. Li, "Adversarial examples: Attacks and defenses for deep learning," *CoRR*, vol. abs/1712.07107, 2017. [Online]. Available: <http://arxiv.org/abs/1712.07107>
- [6] Amazon alexa. [Online]. Available: <https://developer.amazon.com/alexa>
- [7] D. Iter, J. Huang, and M. Jermann, "Generating adversarial examples for speech recognition," 2017.
- [8] N. Carlini and D. A. Wagner, "Audio adversarial examples: Targeted attacks on speech-to-text," *CoRR*, vol. abs/1801.01944, 2018. [Online]. Available: <http://arxiv.org/abs/1801.01944>
- [9] A. Y. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng, "Deep speech: Scaling up end-to-end speech recognition," *CoRR*, vol. abs/1412.5567, 2014. [Online]. Available: <http://arxiv.org/abs/1412.5567>
- [10] N. Roy, H. Hassanieh, and R. Roy Choudhury, "Backdoor: Making microphones hear inaudible sounds," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '17. New York, NY, USA: ACM, 2017, pp. 2–14. [Online]. Available: <http://doi.acm.org/10.1145/3081333.3081366>
- [11] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, "DolphinAttack: Inaudible Voice Commands," *ArXiv e-prints*, Aug. 2017.
- [12] N. Carlini, P. Mishra, T. Vaidya, Y. Zhang, M. Sherr, C. Shields, D. Wagner, and W. Zhou, "Hidden voice commands," in *Proceedings of the 25th USENIX Conference on Security Symposium*, ser. SEC'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 513–530. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3241094.3241135>
- [13] M. Alzantot, B. Balaji, and M. B. Srivastava, "Did you hear that? adversarial examples against automatic speech recognition," *CoRR*, vol. abs/1801.00554, 2018. [Online]. Available: <http://arxiv.org/abs/1801.00554>
- [14] R. Taori, A. Kamsetty, B. Chu, and N. Vemuri, "Targeted adversarial examples for black box audio systems," 2018. [Online]. Available: <http://arxiv.org/abs/1805.07820>
- [15] C. M. Fonseca and P. J. Fleming, "Genetic algorithms for multiobjective optimization: Formulation discussion and generalization," in *Proceedings of the 5th International Conference on Genetic Algorithms*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1993, pp. 416–423. [Online]. Available: <http://dl.acm.org/citation.cfm?id=645513.657757>
- [16] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, April 2002.
- [17] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The kaldi speech recognition toolkit," in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society, Dec. 2011, iEEE Catalog No.: CFP11SRW-USB.
- [18] C. M. Fonseca and P. J. Fleming, "An overview of evolutionary algorithms in multiobjective optimization," *Evolutionary Computation*, vol. 3, no. 1, pp. 1–16, 1995. [Online]. Available: <https://doi.org/10.1162/evco.1995.3.1.1>
- [19] C. A. Coello Coello, "A comprehensive survey of evolutionary-based multiobjective optimization techniques," *Knowledge and Information Systems*, vol. 1, no. 3, pp. 269–308, Aug 1999. [Online]. Available: <https://doi.org/10.1007/BF03325101>
- [20] L. R. Rabiner, B. Gold, and C. K. Yuen, "Theory and application of digital signal processing," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 8, no. 2, pp. 146–146, Feb 1978.
- [21] L. Eshelman and J. Schaffer, "Real-coded genetic algorithms and interval-schema," 11 1992.
- [22] A. Eiben and J. Smith, *Introduction To Evolutionary Computing*, 01 2003, vol. 45.
- [23] Mozilla common voice. [Online]. Available: <https://voice.mozilla.org/en>