

AUTOMATIC GRAMMAR AUGMENTATION FOR ROBUST VOICE COMMAND RECOGNITION

Yang Yang[†], Anusha Lalitha^{*}, Jinwon Lee[†], Chris Lott[†]

[†] Qualcomm AI Research, San Diego

^{*} Department of Electrical and Computer Engineering, University of California San Diego

ABSTRACT

This paper proposes a novel pipeline for automatic grammar augmentation that provides a significant improvement in the voice command recognition accuracy for systems with small footprint acoustic model (AM). The improvement is achieved by augmenting the user-defined voice command set, also called grammar set, with alternate grammar expressions. For a given grammar set, a set of potential grammar expressions (candidate set) for augmentation is constructed from an AM-specific statistical pronunciation dictionary that captures the consistent patterns and errors in the decoding of AM induced by variations in pronunciation, pitch, tempo, accent, ambiguous spellings, and noise conditions. Using this candidate set, greedy optimization based and cross-entropy-method (CEM) based algorithms are considered to search for an augmented grammar set with improved recognition accuracy utilizing a command-specific dataset. Our experiments show that the proposed pipeline along with algorithms considered in this paper significantly reduce the mis-detection and mis-classification rate without increasing the false-alarm rate. Experiments also demonstrate the consistent superior performance of CEM method over greedy-based algorithms.

Index Terms— voice command recognition, CTC, grammar augmentation, cross entropy method, statistical pronunciation dictionary

1. INTRODUCTION

Voice UI is becoming ubiquitous for all types of devices, from smartphones to automobiles. Although we have seen substantial improvement in speech recognition accuracy reported in the literature since the advent of deep neural network based solutions [1, 2, 3, 4], designing robust voice UI system for low memory/power footprint embedded devices without a cloud-based back-end still remains a challenging problem. Compared to its cloud-based counterpart, on-device inference, despite being limited by computation power, memory size, and power consumption, remains appealing for several reasons: (i) there are less privacy concerns as user voice data need not be uploaded to the cloud; (ii) it reduces the latency as it does not involve network access delay; (iii) its usage is not restricted by internet availability, and can be applied in devices with no built-in communication module.

In this work, we focus on improving the recognition accuracy of on-device voice UI systems designed to respond to a limited set of pre-defined voice commands. Such voice UI systems are commonly used in modern IoT/embedded devices such as bluetooth speaker, portable camcorder, hearables, home appliances, etc. Specially, we assume a fixed audio front-end and only look at the pipeline of mapping acoustic features to voice commands.

As illustrated in Fig. 1, we focus on the voice command recognition system composed of an acoustic model (AM) encoder that




Fig. 1. Voice command recognition pipeline

converts the acoustic features into phoneme/grapheme-based probabilistic output, followed by a decoder (e.g., FST) that maps the probabilistic output from AM to one of the voice commands. State of the art acoustic model utilizes either CTC [5], RNN-transducer [4], or Attention Model [6] (see [7, 8] for a good summary). They generate probabilistic outputs, which are fed to a decoder that generates the posterior probability of the corresponding phoneme or grapheme label. Even though these model architectures and training methodologies lead to satisfactory and even super-human transcription accuracy, the best models obtained are often too large for their deployment in small portable devices, e.g., even the smallest model considered in [9] (Table 11 therein) has 18M parameters.

In this work, we utilize a 211K parameter unidirectional-RNN-based acoustic model trained with CTC criterion using Librispeech [10] and a few other datasets, which output probabilities on grapheme targets. Due to the small model size, its transcription accuracy is low: the greedy decoding word-error-rate (WER) without any language model is 48.6% on Libri-speech test-clean dataset. Hence, one of the challenges addressed by our work is, given a small acoustic model trained with general speech dataset, how can one improve the command recognition accuracy utilizing limited command-specific data. Such small footprint AMs have been considered for keyword detection in [11] and [12]. Our work extends these by improving the command command recognition accuracy with a small footprint AM.

In Table 1, we list a few samples of the greedy decoding results from the 211K parameter acoustic model. It is worth noting that even though the word-error-rate is high, the error that it makes tends to be a phonetically plausible rendering of the correct word [1]. Running through a large dataset, we also observe that the error patterns tend to be consistent across different utterances. This leads to a useful insight: for the recognition of a limited set of voice commands (a.k.a. grammar of the decoder), one could improve recognition accuracy

AM greedy decoding	Ground truth
the recter pawsd and den shaking his classto hands before him went on	the rector paused and then shaking his clasped hands before him went on
tax for wone o thease and itees he other	facts form one of these and ideas the other

Table 1. Greedy decoding samples from the acoustic encoder. Word errors are labeled in bold.

by adding variations that capture common and consistent errors from the acoustic model to original command set. We define grammar as a set of valid voice commands (e.g., the grammar can be $\{\text{play music, stop music, ...}\}$) and we refer to this technique of adding variations to the original grammar as *grammar augmentation*. Effective grammar augmentation is the focus of this work.

The main contribution of this paper is the design of effective grammar augmentation framework which provides significant improvement over the baseline system. Next, we highlight our main contributions in detail: (a) For any given set of original voice commands, we propose the design of a candidate set of all grammar variations which captures the consistent errors for a given AM (b) We propose a technique for fast evaluation of command recognition accuracy along with false-alarm and mis-detection rate for any augmented grammar set and finally (c) We devise various algorithms to automatically identify an improved augmented grammar set by suitably adding variations from the candidate set to the original grammar.

Our novel pipeline using the above techniques is illustrated in Fig. 2. The rest of the paper is organized as the following: In Section 2, we give an overview of the proposed grammar augmentation pipeline and dive into the generation of a candidate set and fast grammar evaluation techniques. In Section 3 algorithms via greedy optimization and CEM algorithm are utilized to automate the grammar augmentation process. The experiment results are presented in Section 4 and we discussion on future directions in Section 5.

2. PIPELINE FOR AUTOMATIC GRAMMAR AUGMENTATION

Our AM is trained with CTC loss [5], and can thus assign a posterior probability $\mathbb{P}_{\text{CTC}}(g|u)$ for each command g in a command set, for an input utterance u . For a given test utterance, our system picks the command with the highest probability, or rejects the utterance if the highest probability is below a pre-defined confidence threshold (see Section 2.3) [13][14].

Command decoding errors happen if the AM output deviates from the ground truth to the extent that it can no longer successfully discriminate against other grammar sequences. The idea behind grammar augmentation is to restore the discriminative power of the acoustic model by including in the grammar the sequence variations that capture pronunciation variations or consistent AM error patterns. To do that, we begin with generation of a candidate set containing meaningful variations.

2.1. AM-specific statistical pronunciation dictionary

The augmentation candidates should ideally capture consistent error patterns from the AM, induced by variations in pronunciation, pitch, tempo, accent, ambiguous spellings, or even inherent mistakes made by the AM. For example, if any command includes words that have homophones, then it is necessary to consider adding those homophones into the grammar. To capture these word-level variations, we introduce a novel concept named *AM-specific statistical pronunciation dictionary*, obtained by the following steps: First, we run the AM through a large general speech dataset (e.g., the training set of AM). For each utterance, we obtain its greedy decoding sequence by outputting the character with the maximum probability at each time frame, followed by the CTC squashing function [5] to collapse repeated output graphemes and remove blanks. Given that most utterances from a general speech dataset correspond to a sentence rather than a single word, we use Levenshtein algorithm to find

the minimum-edit-path of the ground-truth to the decoding, and by doing so obtain a mapping of each word to its corresponding maximum probability decoding. For each word, we gather the statistics regarding the frequencies of its maximum-probability decoding outputs. Here we sample a few entries from the dictionary obtained using our 211K-parameter AM:

set		pause		two	
set	32.2%	pause	15.7%	to	53.3%
said	16.6%	pose	14.9%	two	34.7%
sat	11.4%	pase	7.68%	do	1.0%
sait	8.15%	porse	7.31%	tu	0.7%
sed	4.71%	pas	7.31%	too	0.3%

2.2. Candidate set for grammar augmentation

Utilizing this statistical dictionary, we build a candidate set containing potential grammar variations by repeatedly replacing each word in the original grammar by its top- k likely max-decoding outputs. Consider a voice UI application for a small bluetooth player, one could have the following five commands forming the original grammar.

command (\mathcal{C})	original grammar ¹	candidate set for grammar augmentation (\mathcal{G})
play music	play music	pla music, ply music, play mesic, ...
stop music	stop music	stap music, stup music, stup mesic, ...
pause music	pause music	pose music, pase mesic, pause mesic, ...
previous song	previous song	previs song, previous son, ...
next song	next song	nex song, lext song, nex son, ...

By looking up in the statistical dictionary the words contained in the original grammar, one can form an array of alternate expressions for the original commands as shown above. For each command, the set of candidates is the cartesian product of the top- k decoding list from the statistical pronunciation dictionary for each word in the command. The value of k can be different for different words, and is chosen to capture at least a certain fraction of all the variations.

2.3. Evaluation of command recognition accuracy

Let us denote the set of commands as \mathcal{C} , the set of all grammar candidates \mathcal{G} , and the mapping function from \mathcal{G} to \mathcal{C} as f . A grammar G is a subset of \mathcal{G} . For the purpose of evaluating the recognition accuracy of any grammar, we need a command-specific dataset containing audio waveforms and the corresponding target commands. We denote such dataset as $(u, t) \in \mathcal{D}$ with u and t denoting an utterance and its corresponding target command. To evaluate the false alarm rate, we also need an out-of-domain dataset $u \in \mathcal{D}_{\text{ood}}$ that contains a set of utterances that do not correspond to any of the commands.

As mentioned before, the acoustic decoder compares the posterior probabilities $\mathbb{P}_{\text{CTC}}(g|u)$ of all the grammar candidates g included in grammar set $G \subset \mathcal{G}$ given the audio waveform u , and output the command $f(g^*)$ where $g^* = \arg\max_{g \in G} \mathbb{P}_{\text{CTC}}(g|u)$. This calculation is done by running a forward-only dynamic programming algorithm on the AM output. In order to avoid having to repeat the calculation of the probability scores for every choice of grammar

¹Here we assume that AM is trained with grapheme as target, and as a result the grammar is exact the same as the command. Note that the same grammar augmentation pipeline introduced here can be applied to AM trained with phoneme target as well, in which case the grammar is a set of phoneme sequences, and the statistical pronunciation dictionary contains variations of each word in phoneme representation.




Fig. 2. Grammar augmentation pipeline.

set $G \subseteq \mathcal{G}$, we pre-compute and store the probability scores for all the candidate grammar, and all the utterances in both command-specific dataset \mathcal{D} and out-of-domain dataset \mathcal{D}_{ood} . Precisely, as a pre-processing steps of the grammar augmentation search algorithm, we obtain the following probability scores:

$$\mathbb{P}_{\text{CTC}}(g|u), \quad \forall g \in \mathcal{G}, \forall u \in \mathcal{D} \cup \mathcal{D}_{\text{ood}}. \quad (1)$$

To achieve a false alarm rate (FAR) target of α , the confidence threshold for the probability score can be computed as below,

$$\tau(G, \alpha) = \min_{\tau} \left\{ \tau : \frac{|\{u \in \mathcal{D}_{\text{ood}} : \max_{g \in G} \mathbb{P}_{\text{CTC}}(g|u) > \tau\}|}{|\mathcal{D}_{\text{ood}}|} < \alpha \right\}.$$

The decoded command for an utterance u is

$$d(G, \alpha, u) = \begin{cases} \phi, & \text{if } \max_{g \in G} \mathbb{P}_{\text{CTC}}(g|u) < \tau(G, \alpha), \\ \operatorname{argmax}_{c \in \mathcal{C}} \max_{g \in G, f(g)=c} \mathbb{P}_{\text{CTC}}(g|u), & \text{otherwise,} \end{cases}$$

where ϕ denotes decoding being out-of-domain.

With a fixed false-alarm rate, there are two types of error event: mis-detection and mis-classification. Mis-detection refers to the case where a voice command is issued but not detected (i.e., decoded as being out-of-domain), whereas mis-classification happens where a voice command is issued and detected, but the wrong command is decoded. Precisely, the mis-detection-rate (MDR) and the mis-classification-rate (MCR) are defined as below

$$\begin{aligned} \text{MDR}(G, \alpha) &= |\{(u, t) \in \mathcal{D} : d(G, \alpha, u) = \phi\}| / |\mathcal{D}|, \\ \text{MCR}(G, \alpha) &= |\{(u, t) \in \mathcal{D} : d(G, \alpha, u) \notin \{\phi, t\}\}| / |\mathcal{D}|. \end{aligned}$$

3. AUGMENTATION SEARCH ALGORITHMS

The grammar augmentation algorithms we consider search for the grammar set G among all subsets of a candidate set \mathcal{G} that minimizes a weighted sum of the mis-detection-rate and mis-classification-rate with a fixed false-alarm target α ,

$$\min_{G \subseteq \mathcal{G}} \text{MCR}(G, \alpha) + \beta \text{MDR}(G, \alpha). \quad (2)$$

Here the weight factor β controls the significance of mis-detection versus mis-classification. Since we pre-compute the probabilities as shown in Equation (1), for each grammar $G \subseteq \mathcal{G}$ the objective function can be evaluated without invoking the AM, which significantly speeds up the search algorithms.

It is important to note that adding candidate to the grammar does not always improve performance: (i) With a fixed false-alarm target, adding more candidates only increase the confidence threshold $\tau(G, \alpha)$, which could potentially result in degraded mis-detection rate. (ii) distinguishability of the commands has a complex inter-dependency, hence adding grammar candidate for one command may reduce the recognition rate of other commands, as it may alter the classification boundary amongst the set of commands.

3.1. Augmentation via greedy optimization methods

We consider the following three methods based on greedy optimization:

Naive greedy search: Start with the original grammar, iteratively go through all the candidates from \mathcal{G} . In each iteration, add the candidate that best improves the objective function and update the confidence threshold to maintain target FAR, until no candidate can improve further.

Greedy search with refinement: This algorithm is similar to greedy search except for every time a candidate is added to the grammar, we remove those candidates among the remaining ones which contain the added candidate as a subsequence. For example, for *pause music* command, if candidate *pose music* is added to the grammar, then *porse music* is removed from subsequent iterations. Trimming the candidate set in this manner increases the diversity of variations in the grammar.

Beam-search: In each iteration a list of l best grammar sets is maintained. This degenerates to the naive greedy algorithm when $l = 1$.

3.2. Augmentation via cross entropy method (CEM)

Cross entropy method (CEM) is a widely used combinatorial optimization algorithm and has been successfully applied in some reinforcement learning problems [15, 16]. The main idea is rooted from rare event sampling, for which the algorithm tries to minimize the KL divergence between a proposed sampling distribution and the optimal zero-variance importance sampling distribution [16]. Going back to the grammar augmentation objective function in Equation (2), the search space is the power set of the candidate set \mathcal{G} , which can be represented by $\{0, 1\}^{|\mathcal{G}|}$, with each grammar choice represented by a $|\mathcal{G}|$ -dimensional binary vector.

Applying the idea of CEM, we start with an initial probability distribution on $\{0, 1\}^{|\mathcal{G}|}$, and iteratively tune its parameter so that it assigns most of the probability mass in the region towards the minimization of the objection function. In our design, the distribution on this discrete space is induced by the sign of a $|\mathcal{G}|$ -dimensional independent Gaussian distributions, parameterized by their mean and variance in each dimension and initialized as stand Gaussian. For each iteration, we start with a population of s samples from the current distribution, each representing a feasible candidate choice. We evaluate the objective function of $\text{MDR} + \beta \text{MCR}$ for each of sample candidate choice, and keep the best γ fraction. We then update the parameter of the distribution using the sample mean and variance of the top γs candidates (also called elite set), and iterate the procedure by obtaining s samples from the updated distribution.

4. EXPERIMENTS

In this section, we present some experiments which illustrate the improvement that can be obtained in recognition accuracy by applying our grammar augmentation algorithm. All the results are ob-

tained with a dataset containing 5 commands: *play music*, *pause music*, *stop music*, *next song* and *previous song*. This dataset contains utterances with varying gender, pitch, volume, noise types and accents, and are split into training, validation, and testing datasets. The training dataset is used to train the augmentation search algorithms to minimize the objective defined in (2). The validation dataset is used to compare performances of grammar sets obtained and decide which one to take. Finally, we report the results of the final grammar set on a test dataset. For the training objective function in Equation (2), we pick $\beta = 1$, in which case minimizing the sum of MDR and MCR is equivalent to maximizing the command success rate $1 - \text{MCR}(G, \alpha) - \text{MDR}(G, \alpha)$. A candidate set is obtained from running the 211K parameter AM with a 2000-hour dataset using steps discussed in Section 2.1 and 2.2. We consider 150 grammar candidates ($|\mathcal{G}| = 150$) using our statistical pronunciation dictionary.

4.1. Performance Evaluation

We analyze the grammar augmentation algorithms described in Section 3 with a fixed FAR target of $\alpha = 0.1\%$ and compare the augmented grammar output by each algorithm. Fig. 3 shows the command success rate and the decomposition of the error in terms of mis-detection and mis-classification. Note that CEM algorithm provides most improvement in command success rate unlike greedy-optimization based algorithms which may commit to sub-optimal grammar sets early on. As discussed previously, adding more variations to the grammar set makes it more susceptible to mis-detection errors. In fact, adding all 150 grammar expression reduces the command success rate to 80% and increases the MDR to 13.76%. However, Fig 3 shows that performing augmentation in a principled manner can greatly reduce the mis-classification error without increasing the mis-detection errors.




Fig. 3. Performance of grammar augmentation algorithms.

4.2. Complexity of Grammar Augmentation Algorithms

We evaluate the complexity of the augmentation algorithms considered in Section 3. The most computationally expensive step in implementing our augmentation algorithms is the evaluation of MCR and MDR for any candidate grammar set. Hence, we measure the complexity of our augmentation algorithms in terms of number of grammar evaluations needed to output their best augmented grammar set. Fig. 4 illustrates the variation/improvement in command success rate ($1 - \text{MDR} - \text{MCR}$) as the number of grammar evaluations increases. Note that CEM takes only marginally more evaluations while providing the maximum reduction in the sum of MCR and MDR. While beamsearch explores more and requires more grammar evaluation, it provides only marginally better improvement over

naive greedy. The greedy algorithm refinement reaches its best performance in the least number of grammar evaluations. This suggests that incentivizing diversity over exploration may provide better improvement in command success rate and in fewer evaluations.




Fig. 4. Test dataset performance vs. number of grammar evaluations

4.3. Effect of Candidate Set Size on Performance

So far we considered a candidate set size of 150 ($|\mathcal{G}| = 150$). Next, we investigate the effect of varying the candidate set size on the performance of the augmentation algorithms. We vary the candidate size by varying the number of words k we choose from the top- k likely max-decoding outputs for every word in the statistical pronunciation dictionary. Hence, a larger candidate size captures a larger probability of max-decoding outputs. We repeat our experiments by altering the candidate set size from 25 to 150. Table 2 shows the performance the augmentation algorithms for various candidate set sizes. In particular, it shows that CEM improves as we increase the candidate set and is consistently better than greedy based algorithms.

Candidate set size $ \mathcal{G} $	Greedy	Greedy (refinement)	Beamsearch (width 5)	CEM
25	92.31	91.79	92.31	93.25
50	93.16	92.74	93.50	93.59
75	93.08	93.16	92.99	93.68
100	92.82	92.65	92.05	94.02
150	93.76	94.02	93.85	94.44

Table 2. $1 - \text{MDR} - \text{MCR}$ (%) for different algorithms with different candidate set size $|\mathcal{G}|$.

5. CONCLUSION AND FUTURE WORK

In this work, we focus on a small-footprint voice command recognition system composed of a CTC-based small-capacity acoustic encoder, and a corresponding maximum a posteriori decoder for the recognition of a limited set of fixed commands. With a command specific dataset, we proposed a novel pipeline that automatically augments the command grammar for improved mis-detection and mis-classification rate. We achieved this by adapting the decoder to the consistent decoding variations of the acoustic model. An important direction of future work is to extend our grammar augmentation pipeline to provide personalization, i.e., to improve the recognition accuracy for a specific user by adapting the decoder to better fit both the AM and the user's pronunciation pattern.

6. REFERENCES

- [1] Awni Y. Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng, “Deep Speech: Scaling up end-to-end speech recognition,” 2014, vol. abs/1412.5567.
- [2] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, Jie Chen, Jingdong Chen, Zhijie Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, and et al, “Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin,” in *Proceedings of The 33rd International Conference on Machine Learning*, Maria Florina Balcan and Kilian Q. Weinberger, Eds., New York, New York, USA, 20–22 Jun 2016, vol. 48 of *Proceedings of Machine Learning Research*, pp. 173–182, PMLR.
- [3] Ronan Collobert, Christian Puhrsch, and Gabriel Synnaeve, “Wav2Letter: an End-to-End ConvNet-based Speech Recognition System,” 2016, vol. abs/1609.03193.
- [4] Alex Graves, “Sequence Transduction with Recurrent Neural Networks,” 2012, vol. abs/1211.3711.
- [5] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, “Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks,” in *Proceedings of the 23rd International Conference on Machine Learning*, New York, NY, USA, 2006, ICML ’06, pp. 369–376, ACM.
- [6] Jan Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, KyungHyun Cho, and Yoshua Bengio, “Attention-Based Models for Speech Recognition,” 2015, vol. abs/1506.07503.
- [7] Rohit Prabhavalkar, Kanishka Rao, Tara N. Sainath, Bo Li, Leif Johnson, and Navdeep Jaitly, “A Comparison of Sequence-to-Sequence Models for Speech Recognition,” in *INTER-SPEECH*, 2017.
- [8] E. Battenberg, J. Chen, R. Child, A. Coates, Y. G. Y. Li, H. Liu, S. Satheesh, A. Sriram, and Z. Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Dec 2017, pp. 206–213.
- [9] Dario Amodei, Rishita Anubhai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Jingdong Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, and et al, “Deep Speech 2: End-to-End Speech Recognition in English and Mandarin,” 2015, vol. abs/1512.02595.
- [10] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2015, pp. 5206–5210.
- [11] Tara N. Sainath and Carolina Parada, “Convolutional neural networks for small-footprint keyword spotting,” in *INTER-SPEECH*, 2015, pp. 1478–1482, ISCA.
- [12] G. Chen, C. Parada, and G. Heigold, “Small-footprint keyword spotting using deep neural networks,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 4087–4091.
- [13] Y. Miao, M. Gowayyed, and F. Metze, “EESSEN: End-to-end speech recognition using deep RNN models and WFST-based decoding,” in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, Dec 2015, pp. 167–174.
- [14] Naoyuki Kanda, Xugang Lu, and Hisashi Kawai, “Maximum a posteriori based decoding for ctc acoustic models,” in *INTER-SPEECH*, 2016.
- [15] István Szita and András Lörincz, “Learning Tetris Using the Noisy Cross-entropy Method,” Cambridge, MA, USA, Dec. 2006, vol. 18, pp. 2936–2941, MIT Press.
- [16] Pieter-Tjerk de Boer, Dirk P. Kroese, Shie Mannor, and Reuven Y. Rubinstein, “A Tutorial on the Cross-Entropy Method,” Feb 2005, vol. 134, pp. 19–67.

Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc. and/or its subsidiaries