

LIBRI-ADAPT: A NEW SPEECH DATASET FOR UNSUPERVISED DOMAIN ADAPTATION

Akhil Mathur^{†*} Fahim Kawsar^{*} Nadia Berthouze[†] Nicholas D. Lane[‡]

[†] University College London, United Kingdom

^{*}Nokia Bell Labs Cambridge, United Kingdom

[‡]University of Oxford, United Kingdom

ABSTRACT

This paper introduces a new dataset, Libri-Adapt, to support unsupervised domain adaptation research on speech recognition models. Built on top of the LibriSpeech corpus, Libri-Adapt contains 7200 hours of English speech recorded on mobile and embedded-scale microphones, and spans 72 different domains that are representative of the challenging practical scenarios encountered by ASR models. More specifically, Libri-Adapt facilitates the study of domain shifts in ASR models caused by a) different acoustic environments, b) variations in speaker accents, c) previously unexplored factors such as heterogeneity in the hardware and platform software of the microphones, and d) a combination of the aforementioned three shifts. We also provide a number of baseline results quantifying the impact of these domain shifts on the Mozilla DeepSpeech2 ASR model.

Index Terms— Unsupervised Domain Adaptation, Deep Learning, Microphone Variability, Speech Recognition

1. INTRODUCTION

Recent advancements in deep neural networks have led to significant improvements in the performance of automatic speech recognition (ASR). However, it is also known that minor deviations between the training and test domains (known as *domain shift*) can significantly degrade the generalizability of deep learning models [1]. In the case of speech recognition, such perturbations could be introduced in the data by several real-world factors including ambient noise [2], variations in acoustic environments [3, 4], accents of speakers [5], or even the microphones used to record the speech [6, 7] – and they represent a major challenge to the generalizability of ASR models in unconstrained settings. Clearly, the collection of labeled training data across all these heterogeneous scenarios is prohibitively expensive, therefore *unsupervised domain adaptation* has emerged as a promising technique towards adapting deep neural networks to different but related domains, even in the absence of labeled data from the target domain [8].

There are plethora of datasets to study and develop unsupervised domain adaptation techniques for visual recognition (e.g., [9, 1, 10]), however in speech recognition research, there is still a lack of large-scale realistic datasets which contains examples of various forms of *domain shifts* that are expected in real-world scenarios. In this work, we present a new dataset named *Libri-Adapt*, primarily developed to study unsupervised domain adaptation, which contains domain shifts caused by a) variability in recording microphones, b) variability in speaker accents and c) variability in acoustic environments. Libri-Adapt is mainly recorded on consumer-grade, embedded-scale microphones (please refer Table 1) that are becoming popular due to their easy-integration capabilities with microcom-

puters such as the Raspberry Pi to enable developers to create speech applications for wearable and embedded devices.

Libri-Adapt is built on top of the Librispeech-clean-100 dataset [11] which contains 100 hours of US English speech from users reading public-domain audio books. In Libri-Adapt, we record the Librispeech-clean-100 training corpus on six different microphones (five are embedded-scale or smartphone microphones and one is a Shure microphone to establish a higher-quality baseline), in three different accents (US English, British English, and Indian English) and under four synthesized background noise conditions (Clean, Rain, Wind, Laughter). As such, Libri-Adapt has a total of 7,200 hours (6 microphones x 3 accents x 4 environments x 100 hours) of data for training ASR models and can be used to study domain adaptation for potentially 72 different domains (6 microphones x 3 accents x 4 environments). For each domain, we also collect a 5.4 hour held-out test set derived from Librispeech test-clean corpus [11].

In §3, we describe the methodology used for collecting the dataset followed by an overview of how the dataset is structured. In §4, we demonstrate the potential of doing unsupervised domain adaptation research on this dataset by training a Mozilla DeepSpeech2 [12] ASR model on each domain in the dataset and testing it on another domains. Our results show that even state-of-the-art ASR models such as DeepSpeech2 suffer accuracy degradation due to domain shifts – e.g., the word error rate of a model trained on the Shure microphone increases from 10.19% (when it is tested on the same microphone) to 24.95% when it is tested on a Matrix Voice microphone.

In summary, Libri-Adapt is the largest public dataset targeted for consumer-grade embedded-scale microphones to facilitate the study and development of robust ASR models using domain adaptation techniques. The dataset will be released publicly under the CC BY 4.0 license¹ and can be accessed from <https://bit.ly/2UK1McG>.

2. RELATED WORK

Our work follows a rich tradition of making speech data public to support research on speech and audio recognition. For example, Gemmeke et al. [13] released AudioSet at ICASSP 2017 for the task of audio event recognition. Similarly, a number of speech and keyword datasets have been released, including LibriSpeech [11], Common Voices [14] and Speech Commands [15]. Researchers have also developed simulated datasets focused on the major challenges in ASR, e.g., handling reverberations [16, 17], noisy speech [2], speech separation and distant ASR [17, 2]. A prominent example has been the CHiME challenge [2, 3, 4] which has released datasets containing real and simulated multi-channel speech data in noisy real-world

¹<https://creativecommons.org/licenses/by/4.0/>

conditions as well as in multi-speaker scenarios. For the purpose of studying accent variability, prior works have used the Speecon [5] family of speech corpora recorded in multiple languages.

To the best of our knowledge, no prior work has developed a large-scale speech corpus for embedded-scale microphones (refer Table 1 for details) that are gaining popularity nowadays. Further, recent research [6] has shown that embedded microphones exhibit significant differences in their outputs owing to the variations in their hardware and on-board signal processing pipelines. As such, the primary goal of this paper is to collect speech datasets – from embedded-scale microphones – that are large enough to support training of deep ASR models. The variability across microphones can be interpreted as a form of domain shift, thus allowing for development and evaluation of unsupervised domain adaptation algorithms on this dataset. In addition, we are also interested in understanding how microphone variations interact with other challenges for ASR, namely variations in speaker accents and acoustic environments – as such, we record speech data from 3 speaker accents in 4 acoustic environments on 6 different embedded microphones.

Unsupervised Domain Adaptation (UDA). At a high-level, the idea behind UDA is as follows: given a pre-trained deep classifier for a source domain and an unlabeled dataset from a target domain, how can we adapt the weights of the source model such that it shows better performance in the target domain. This class of algorithms has been used recently to adapt ASR models from clean to noisy speech, for speech dereverberation and for speaker gender adaptation [18, 19, 20]. Our work aims to advance this line of research by introducing the large-scale multi-domain Libri-Adapt dataset.

3. METHODOLOGY AND DATASET OVERVIEW

In this section, we describe the different domains in the Libri-Adapt dataset and the methodology adopted to create them.

Microphones. One of the important goals for Libri-Adapt is to support ASR research on consumer-grade embedded-scale microphones. In the past few years, a number of embedded-scale microphones (both single- and multi-channel) have been released which can be used in conjunction with microcomputers such as Raspberry Pi and cloud-based speech models to develop custom speech analysis systems [21]. However, research [6] has shown that these microphones could potentially have significant variations in their hardware and software processing pipelines, which in turn could induce domain shift in the speech data obtained from them. Therefore, in order to develop robust ASR systems, we believe it is critical to collect speech datasets from various microphones and evaluate the generalizability of ASR models across them. In this vein, we collect training data (derived from Librispeech-clean-100) from six different microphones as listed in Table 1. Matrix Voice² and ReSpeaker³ are circular 7-channel microphone arrays that include on-device acoustic processing algorithms for de-reverberation and delay-and-sum beamforming to combine the outputs of the different channels. PlayStation-Eye is a digital camera used in Sony PlayStation gaming consoles and consists of a 4-channel microphone array used for speech user interaction with the gaming console. Next, we use a Google Nexus 6 as a representative smartphone microphone with which users are likely to interact while using speech applications. Further, we use a single-channel USB microphone as an example of low-end microphones popularly used with Raspberry Pi. Finally,

²<https://www.matrix.one/products/voice>

³https://respeaker.io/usb.6+1_mic.array/

Device	Channels	Maximum Sampling Rate (kHz)	Advertised Signal Processing Capabilities
Matrix Voice	7	8 to 96	Beamforming, dereverberation, noise cancellation
Respeaker	7	8 to 48	Voice activity detection, beamforming, background noise suppression
PlayStation Eye	4	8 to 48	Voice location tracking, echo cancellation, background noise suppression.
USB Mic	1	8 to 48	-
Google Nexus 6	3	8 to 48	-
Shure MV5	1	8 to 48	Auto gain, Equalization

Table 1. Technical specification of the microphones used in the Libri-Adapt dataset.

to establish a higher-quality microphone baseline, we employ a Shure MV5 Condenser microphone⁴ which can also be used as a plug-and-play microphone with mobile and wearable devices.

We use a Raspberry Pi as the host device for Matrix, ReSpeaker, USB and PlayStation Eye, while the Shure microphone uses a MacBook Pro as the host device. Each microphone in the dataset can be considered as a unique domain, which allows researchers to perform a variety of domain adaptation experiments.

Speaker Accents. Libri-Adapt consists of speech data in three English accents, namely US English, British English and Indian English. To facilitate domain adaptation experiments for countering the domain shift due to variations in speaker accents, we made a design decision to create accented speech that share the same underlying text corpus. In other words, the accented speech subsets have exactly the same sentences being spoken, but in different English accents. For US English, we use the original Librispeech-clean-100 dataset that was recorded with 251 US speakers. To create its variants in British English and Indian English accents, we employed the text-to-speech (TTS) API⁵ from Google Cloud Platform which uses a pretrained WaveNet model to generate accented speech in multiple voices. More specifically, we passed the transcripts of the Librispeech-clean-100 corpus to the Google TTS API along with the target accent (British and Indian) and obtained two accented versions of the Librispeech-clean-100 dataset, while ensuring that the underlying spoken sentences remain same in each accent.

Data Recording Setup. We adopted the methodology of replay-and-record to collect the Libri-Adapt dataset. Each accented dataset (US English, British English and Indian English) was played on a JBL LSR305 monitor speaker connected to a laptop, and recorded simultaneously on all six microphones in a quiet room. The micro-

⁴<https://www.shure.com/en-GB/products/microphones/mv5/>

⁵<https://cloud.google.com/text-to-speech/>

Domain Types	Count	Domain Names
Microphones	6	Matrix Voice, ReSpeaker, PlayStation Eye, USB, Shure, Nexus 6
Accents	3	US English, Indian English, British English
Background Noise	4	Clean (or None), Rain, Wind, Laughter

Table 2. Overview of the Libri-Adapt dataset

phones were kept at a distance of 15cm from the speaker. To ensure that the recordings are time-synchronized across microphones, we implemented a network-based synchronization scheme based on MQTT [22] messaging protocol. Before a speech file is played, the laptop sends a MQTT message to the recording device hosting each microphone to initiate an audio recording process for the duration of the speech file. Thereafter, the speech file is played on the JBL LSR305 monitor speaker and is recorded simultaneously by all six microphones. After the recording completes, each device sends a confirmation message to the laptop, which proceeds to play the next file in the dataset.

Indeed, this approach of replay-and-record has its pros and cons. It allows for collecting large-scale datasets (recall that Libri-Adapt has 7200 hours of data) with various domain shifts cheaper than recruiting human subjects – however on the downside, as the speech is replayed through a speaker, it is impacted by the transfer function of the speaker. To counter this issue, we use a high-quality monitor speaker with a reasonably flat frequency response to replay the audios, which ensures that there is minimal impact on the replay audio quality due to the speaker. More importantly, we argue that the replayed dataset still allows for studying the domain shift problem caused by different microphones because the replay process could be interpreted as a form of covariate shift on the input distribution, that is shared across all microphones.

Ambient Noise Augmentation. Finally, we augment the speech datasets collected from all microphones with three types of background noise, namely ‘Wind’, ‘Rain’ and ‘Laughter’. To this end, we use the audio samples for ‘Wind’, ‘Rain’ and ‘Laughter’ from the Environmental Sound Classification (ESC-50) [23] dataset and record them on each of the microphone following our replay-and-record methodology. In total, we record 20 noise samples for each noise type on each microphone. Finally, to create the noisy variants of the dataset, each speech file in the original (Clean) dataset is overlaid with a randomly picked noise sample from the same microphone (at SNR of 20dB). As such, we end up with 4 background noise conditions in the dataset: Clean, Rain, Wind, Laughter.

Dataset Overview. Table 2 provides an overview of the Libri-Adapt dataset. In total, the dataset has 6 microphone domains, 3 accent domains and 4 background noise domains, resulting in a total of $6 \times 3 \times 4 = 72$ domains. Each domain consists of 28,540 speech files (sampling rate 16KHz) with a duration of 100.6 hours as labeled training data, and 2600 speech files with a duration of 5.4 hours as labeled test data. The dataset facilitates a wide variety of domain adaptation experiments, both supervised and unsupervised. For example, one could train an ASR model with {US English speakers on a Nexus 6 microphone in Clean background noise} and test its performance on {Indian English speakers using a ReSpeaker microphone with Rain background noise}.

	Training Microphones					
	Matrix Voice	ReSpeaker	USB	Shule	Nexus 6	PlayStation Eye
Matrix Voice	13.05	28.0	20.64	24.95	26.21	24.18
ReSpeaker	16.06	12.66	14.49	16.09	15.93	15.21
USB	13.30	14.44	11.40	14.41	15.39	15.38
Shule	12.58	13.10	11.61	10.19	12.96	13.07
Nexus 6	13.09	12.81	11.99	12.73	12.27	15.28
PlayStation Eye	14.01	13.65	14.24	16.16	16.52	11.39

Table 3. WER of a fine-tuned DeepSpeech2 model trained and tested on various microphone pairs for Clean, US English speech. The columns correspond to the training microphone domain and rows correspond to the test microphone domain.

4. EFFECT OF DOMAIN SHIFT ON ASR MODELS

In this section, we analyze the Libri-Adapt dataset to demonstrate its feasibility in supporting unsupervised domain adaptation research for ASR models. To this end, we perform a number of experiments where we train an ASR model on a particular domain and evaluate its performance on another domain. More specifically, we use the Mozilla DeepSpeech2 pre-trained model [12] (release 0.5.0) as the base model which has a low word error rate (WER) of 8.22% on the LibriSpeech dataset. We fine-tune the base model on the training data from each domain and thereafter, the fine-tuned model for each domain is tested on the held-out test set from other domains to compute the WER.

WER across Microphones. First, we analyze the impact of domain shift caused by microphone variability on the accuracy of ASR models. Table 3 shows the WER of the DeepSpeech2 model when it is fine-tuned for different ‘training’ microphones (in columns) and tested on other microphones (in rows). As expected, in most cases, it can be observed that when data from the same microphone is used for training and testing the model, it achieves the smallest WER, for e.g., 11.39% in the case of PlayStation Eye. However, in cases of domain shift caused by microphone mismatch, the WER increases to 24.18% when the PlayStation Eye model is tested on Matrix Voice. We also note that the WER obtained when models from other microphones are tested on Matrix Voice is significantly high (above 20% in all cases), however when the model is fine-tuned on Matrix Voice, we can achieve a reasonable WER of 13.05%. This highlights the opportunity for domain adaptation algorithms to adapt or align the speech feature representations across microphones when ASR models are deployed in scenarios with microphone mismatch.

WER across Accents. Next, we evaluate the impact of accent changes on ASR models. For this experiment, we set the microphone domain to *ReSpeaker* and background noise to *Clean*, which enables us to isolate the impact of accent variability on the model. As shown in Figure 1, there is a significant impact on WER caused by variations in speaker accents – in particular, the generalizability of US and British English accent models to Indian English accent is quite poor, for instance, the WER of a model trained on US accent increases from 11.4% to 29.5% when it is tested on data with an Indian accent. Interestingly, we also observe that the WER of the British and Indian accent models on their own domains is very low (~5%) as compared to the US accent model (11.4%) – a possible explanation of this result is that British and Indian accent datasets




Fig. 1. WER of a fine-tuned DeepSpeech2 model trained and tested on various accent datasets recorded on a *ReSpeaker* microphone in a *Clean* background.

were generated by the Google WaveNet TTS model which supports a limited number (5-6) of user voices. As such, it was easier to fine-tune the model for these accents, as compared to US accent which had a total of 251 unique users. From a domain adaptation perspective, this represents a practical, albeit challenging, scenario: how to adapt a model trained on a small user group to a larger user group with different accents.

	Training Domains			
	Clean	Rain	Wind	Laughter
Clean	11.70	19.07	16.22	13.84
Rain	30.32	21.63	23.45	23.56
Wind	17.20	22.35	14.62	15.34
Laughter	22.14	21.93	21.13	14.52

Table 4. WER of a fine-tuned DeepSpeech2 model trained and tested on various background noise domains.

WER across Background Noise. In Table 4, we show the cross-domain results when ASR models are trained and tested on different background noise pairs. In this experiment, we set the microphone domain to USB and accent to US English, which helps in isolating the effect of background noise on model performance. The results in Table 4 confirm that domain shifts caused by ambient noise can degrade the performance of state-of-the-art ASR models, hence making it a pertinent problem for domain adaptation research.

Studying multi-domain shifts. So far, we have analyzed the effect of a single form of domain shift on ASR models. However, Libri-Adapt can also be utilized to study the effect of multiple forms of domain shifts on ASR models. As shown in Table 5, we can combine microphones, accents and noise conditions in various ways to form a domain; and then evaluate the cross-domain generalizability of ASR models. For instance, we observe in Table 5 that when a model trained on US accented clean speech collected on a ReSpeaker microphone (*Domain 1*) is tested on Indian accented speech with Rain background noise recorded on a PlayStation Eye (*Domain 2*), it suffers a WER increase of nearly 29.8%. In summary, Libri-Adapt allows researchers to create various complex adaptation scenarios to test the generalizability of their domain adaptation algorithms.

5. CONCLUSION

We presented a speech dataset to facilitate the study and development of unsupervised domain adaptation algorithms for ASR models. The dataset, derived from the Librispeech corpus, contains 7,200

	Training Domains			
	Domain 1	Domain 2	Domain 3	Domain 4
Domain 1	12.86	47.01	37.23	15.64
Domain 2	42.68	6.37	59.25	43.05
Domain 3	27.09	42.73	4.80	25.14
Domain 4	19.07	52.64	46.18	14.26

Table 5. Libri-Adapt facilitates a variety of domain adaptation experiments by combining microphones, accents and noise conditions. This table shows the WER for various domain adaptation tasks. *Domain 1* = ReSpeaker-US-Clean; *Domain 2* = PlayStationEye-Indian-Rain; *Domain 3* = Shule-British-Laughter; *Domain 4* = USB-US-Wind

hours of training data from 72 different domains spanning across 6 microphones, 3 speaker accents and 4 background noise conditions. One limitation of this work is that in order to create a large-scale and multi-domain dataset without incurring significant costs, we adopted the methodology of replay-and-record which may introduce minor speaker-related effects in the dataset. As a future work, we will develop unsupervised domain adaptation algorithms to tackle the domain shifts in the dataset. The Libri-Adapt dataset will be released publicly under the CC BY 4.0 license and can be accessed from <https://bit.ly/2UK1McG>.

6. REFERENCES

- [1] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell, "Adversarial discriminative domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7167–7176.
- [2] Jon Barker, Ricard Marxer, Emmanuel Vincent, and Shinji Watanabe, "The third 'chime' speech separation and recognition challenge: Dataset, task and baselines," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 504–511.
- [3] Emmanuel Vincent, Shinji Watanabe, Aditya Arie Nugraha, Jon Barker, and Ricard Marxer, "An analysis of environment, microphone and data simulation mismatches in robust speech recognition," *Computer Speech & Language*, vol. 46, pp. 535–557, 2017.
- [4] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, "The fifth 'chime' speech separation and recognition challenge: Dataset, task and baselines," in *Interspeech*, 2018.
- [5] Dorota Iskra, Beate Grosskopf, Krzysztof Marasek, H Heuvel, Frank Diehl, and Andreas Kiessling, "Speecon-speech databases for consumer devices: Database specification and validation," 2002.
- [6] Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Nadia Berthouze, and Nicholas D. Lane, "Mic2mic: Using cycle-consistent generative adversarial networks to overcome microphone variability in speech systems," in *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, New York, NY, USA, 2019, IPSN '19, pp. 169–180, ACM.
- [7] Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Robert Smith, Nicholas D. Lane, and Nadia Berthouze, "On robustness of

- cloud speech apis: An early characterization,” in *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, New York, NY, USA, 2018, UbiComp '18, p. 1409–1413, Association for Computing Machinery.
- [8] Hal Daume III and Daniel Marcu, “Domain adaptation for statistical classifiers,” *Journal of artificial Intelligence research*, vol. 26, 2006.
- [9] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell, “Cycada: Cycle-consistent adversarial domain adaptation,” in *International Conference on Machine Learning*, 2018, pp. 1994–2003.
- [10] Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Nadia Berthouze, and Nicholas D. Lane, “Flexadapt: Flexible cycle-consistent adversarial domain adaptation,” 2019.
- [11] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [12] “Pre-trained Mozilla DeepSpeech2 ASR model,” <https://github.com/mozilla/DeepSpeech/releases/tag/v0.5.0>, 2019.
- [13] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [14] “Mozilla Voice,” <https://voice.mozilla.org/en>, 2019, Accessed: February 13, 2020.
- [15] Pete Warden, “Speech commands: A dataset for limited-vocabulary speech recognition,” *arXiv preprint arXiv:1804.03209*, 2018.
- [16] Keisuke Kinoshita, Marc Delcroix, Takuya Yoshioka, Tomohiro Nakatani, Emanuël Habets, Reinhold Haeb-Umbach, Volker Leutnant, Armin Sehr, Walter Kellermann, Roland Maas, et al., “The reverb challenge: A common evaluation framework for dereverberation and recognition of reverberant speech,” in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2013, pp. 1–4.
- [17] Mirco Ravanelli, Piergiorgio Svaizer, and Maurizio Omologo, “Realistic multi-microphone data simulation for distant speech recognition,” in *Interspeech 2016*, 2016.
- [18] Daniel Michelsanti and Zheng-Hua Tan, “Conditional generative adversarial networks for speech enhancement and noise-robust speaker verification,” *Proc. Interspeech 2017*, pp. 2008–2012, 2017.
- [19] Ehsan Hosseini-Asl, Yingbo Zhou, Caiming Xiong, and Richard Socher, “A multi-discriminator cyclegan for unsupervised non-parallel speech domain adaptation,” *arXiv preprint arXiv:1804.00522*, 2018.
- [20] Ke Wang, Junbo Zhang, Sining Sun, Yujun Wang, Fei Xiang, and Lei Xie, “Investigating generative adversarial networks based speech dereverberation for robust speech recognition,” *arXiv preprint arXiv:1803.10132*, 2018.
- [21] “DIY hardware to emulate Amazon Echo,” <http://www.instructables.com/id/Build-DIY-Amazon-Alexa-With-a-MATRIX-Creator-on-Ha/>, 2015.
- [22] “MQTT Messaging Protocol,” <http://mqtt.org/>, 2019.
- [23] Karol J Piczak, “Esc: Dataset for environmental sound classification,” in *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 2015, pp. 1015–1018.