

Technical Note

Not peer-reviewed version

TIPAA-SSL: Text Independent Phone-to-Audio Alignment based on Self-Supervised Learning and Knowledge Transfer

[Noé Tits](#)^{*}, Prernna Bhatnagar, Thierry Dutoit

Posted Date: 2 July 2024

doi: 10.20944/preprints202406.2082.v1

Keywords: speech recognition; phoneme recognition; deep learning; transfer learning



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Technical Note

TIPAA-SSL: Text Independent Phone-to-Audio Alignment based on Self-Supervised Learning and Knowledge Transfer

Noé Tits *, Prernna Bhatnagar and Thierry Dutoit

University of Mons; prernna.bhatnagar@student.umons.ac.be (P.B.); thierry.dutoit@umons.ac.be (T.D.)

* Correspondence: noe.tits@umons.ac.be

Abstract: In this paper, we present a novel approach for text independent phone-to-audio alignment based on phoneme recognition, representation learning and knowledge transfer. Our method leverages a self-supervised model (wav2vec2) fine-tuned for phoneme recognition using a Connectionist Temporal Classification (CTC) loss, a dimension reduction model and a frame-level phoneme classifier trained thanks to forced-alignment labels (using Montreal Forced Aligner) to produce multi-lingual phonetic representations, thus requiring minimal additional training. We evaluate our model using synthetic native data from the TIMIT dataset and the SCRIBE dataset for American and British English, respectively. Our proposed model outperforms the state-of-the-art (charsiu) in statistical metrics and has applications in language learning and speech processing systems. We leave experiments on other languages for future work but the design of the system makes it easily adaptable to other languages.

Keywords: speech recognition; phoneme recognition; deep learning; transfer learning

1. Introduction

Pronunciation is pivotal for language acquisition and effective communication. However, it often receives insufficient attention in second language (L2) education, leading to persistent challenges for learners in achieving clarity. Technology, particularly Computer-Assisted-Language-Learning (CALL) has emerged as a significant aid in supporting L2 pronunciation development across formal and informal learning environments [1]. Recent advancements in deep learning offer promising avenues for enhancing language learning by providing unlimited targeted and immediate feedback, addressing the class time and size limitations preventing educators from differentiating their pedagogy on such an individual skill.

To improve pronunciation, integrating speech processing technologies with traditional teaching methods can provide a comprehensive approach to language learning. One fundamental task in this domain is text independent phone-to-audio alignment, a process that involves aligning phonetic representations with corresponding audio signals without relying on pre-determined text. This task is essential for accurately mapping the phonemes to their acoustic representations, contributing to the development of precise and effective speech processing technologies in language learning [2,3]. Text independent phone-to-audio alignment faces a significant challenge due to the difficulty in obtaining extensive and well-annotated datasets. A possible solution to this challenge entails using established systems (like text dependent phone-to-audio alignment systems) to extract temporal information from publicly available speech datasets. This approach can be refined through the application of transfer learning and self-supervised learning methodologies in the development of a solution.

Transfer learning [4], is a method within the field of deep learning. Its approach involves pre-training a model on a large dataset and subsequently fine-tuning it on a smaller dataset tailored to the specific task. This methodology has demonstrated considerable success across various domains, particularly in the context of self-supervised learning.

In Self-supervised learning [5] a model is trained to autonomously learn representations of input data without the need for explicit supervision. This proves particularly advantageous when labeled data is either limited or entirely unavailable. Within the field of speech technology, the application of self-supervised learning through transfer learning has proven invaluable in addressing several complex

scenarios. For example, in Automatic Speech Recognition (ASR) for low resource languages [6], where annotated data may be scarce, transfer learning allows the utilization of pre-trained models on more data-abundant languages, adapting them effectively to the target language. Likewise, in emotional or expressive speech synthesis [7–9], speech emotion recognition [10], voice conversion [11] and pronunciation assessment [12].

In this paper, we take full advantage of state-of-the-art methodologies in deep learning, self-supervised learning, and phonetic representation to present a novel approach to text independent phone-to-audio alignment. Most state-of-the-art self-supervised systems perform well on American English which means that other varieties of English get penalised and the model is biased towards American English. The rationale to develop the system proposed in the paper was the pedagogical needs to create a system that performs equally well on other variants of English as it does on American English. For our system, we use the self-supervised model Wav2Vec2, fine-tuned for phoneme recognition using CTC loss. We integrate this with a dimensional reduction model based on Principal Component Analysis (PCA) and a frame-level phoneme classifier. The resulting model pipeline produces a vector of probabilities for every audio frame. From the same model we also extract predicted phonemes and their boundaries and hence use this information for a text independent phone alignment system.

The major contributions of this paper are as follows: First, we propose a text independent phoneme alignment system using self-supervised learning. This not only advances the state-of-the-art in this specific task but also opens avenues for broader applications in language learning and speech processing systems. Second, this system functions effectively with diverse English language variations (British) and is capable of accommodating various languages, making it language-independent.

2. Related Work

2.1. Phone-to-Audio Alignment

Text independent phone-to-audio alignment involves predicting a sequence of phones and their temporal locations within speech signal without prior linguistic information, such as a known text or phone sequence.

In contrast, text dependent phone-to-audio alignment utilizes text information to align a phone sequence generated from a grapheme-to-phoneme model applied to textual inputs. Hidden Markov Models (HMMs) have traditionally played a prominent role in aligning phonetic and temporal information, including but not limited to Kaldi [13] and HTK [14]. Forced alignment systems using HMMs include Gentle¹ and ProsodyLab [15]. However, such models face limitations when attempting to predict both phones and phone boundaries simultaneously. The challenges for HMMs become apparent when confronted with long-range dependencies within a sequence. However, they struggle to effectively capture and understand the broader context of a word within a sentence. As a result, this hinders their ability to grasp nuanced linguistic relationships and context. Another reason why HMMs may not perform well is due to the incorrect phonetic transcriptions due to variations in conversational speech.

Recent developments have witnessed a shift towards deep learning models like Recurrent Neural Networks (RNNs) [16] and models trained with CTC loss [17], for phoneme recognition. This shift towards deep learning models allows us to efficiently predict both phones and phone boundaries simultaneously.

¹ <https://lowerquality.com/gentle/>

2.2. Phoneme Recognition

Deep learning models are known for their ability to capture complex patterns in data and are frequently employed for phoneme recognition. CTC loss is a popular training criterion for such models in this context. It allows the model to learn the alignment between the input speech signal and the corresponding phoneme sequence, even when the temporal correspondence is not provided during training. Wav2Vec2 [18,19] stands out as a self-supervised model that has been fine-tuned specifically for phoneme recognition using the CTC loss.

One of the advantages of the Wav2Vec2 approach is its ability to generate multi-lingual phonetic representations. By leveraging self-supervised learning during pre-training, the model learns to extract robust features from speech signals, capturing phonetic information that is broadly applicable across different languages. Furthermore, the fine-tuning process with CTC loss refines the model’s ability to map these learned representations to specific phoneme sequences.

2.3. Systems Predicting Phones and Boundaries

Charsiu [20] is recognized for its unique capability to predict both phones and phone boundaries. This integrated approach contributes to a more comprehensive understanding of the speech signal. While HMMs can perform similar tasks, they are often employed more prevalently for forced-alignment tasks, particularly in text-dependent scenarios. The emphasis on forced alignment implies that these systems are commonly utilized when the corresponding phoneme or phone sequence is available.

Therefore, in this paper, we conduct a comparative analysis between our proposed model and the text independent aspect of the charsiu model, recognized as the state-of-the-art for this task. Our methodology, detailed in the following section, combines self-supervised learning with phoneme recognition.

3. Methodology

3.1. Our Proposed Method

Our proposed method is an innovative approach, using wav2vec2, Principal Component Analysis (PCA) for dimensional reduction and frame-level phoneme classification, offers a robust text independent phone-to-audio alignment. The system is explained in details in subsection 3.3 and the system’s architecture is depicted in Figure 1. The model’s robustness is demonstrated through evaluation using synthetic native data, using the TIMIT [21] and SCRIBE dataset. Thanks to its versatility, we expect that our method will find applications in language learning and speech processing systems.

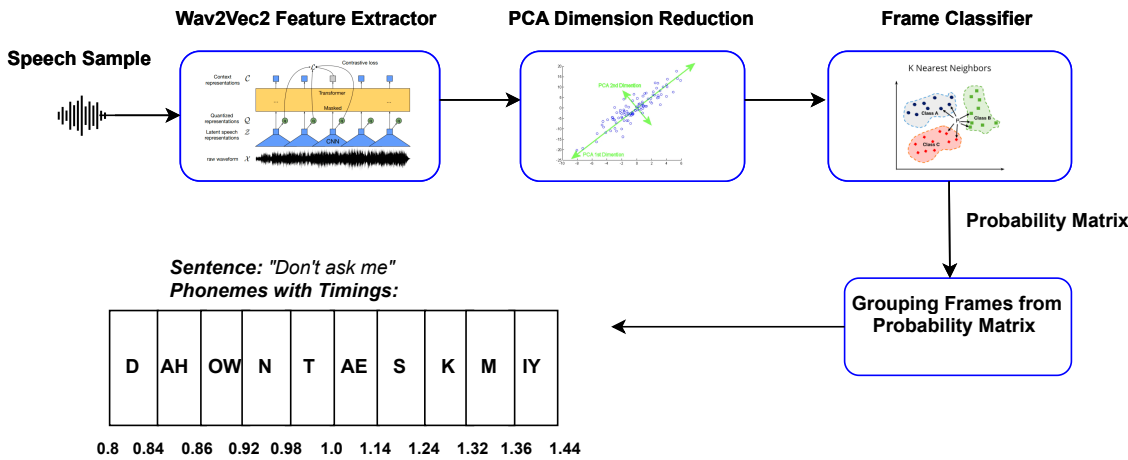


Figure 1. Model pipeline of the system

3.2. Pre-trained Self-Supervised Model

In this Section, we explain what open-source model we use as a basis² in order to leverage learned cross-lingual representations of speech adapted towards the phonetic space.

The basis of this model is a self-supervised model (wav2vec2) trained to extract latent representations from the audio data. The wav2vec2 model is trained on a large dataset of unlabeled speech data, using a contrastive predictive coding loss function to learn a representation that is capable of predicting future audio frames and can then be utilized for downstream tasks. The pre-training loss is defined as [22]:

$$L = L_m + \alpha L_d$$

Where, L_m is contrastive loss, L_d is diversity loss and α is the hypertuned parameter

$$L_m = -\log \frac{\exp(\text{sim}(\mathbf{c}_t, \mathbf{q}_t) / \kappa)}{\sum_{\tilde{\mathbf{q}} \sim \mathbf{Q}_t} \exp(\text{sim}(\mathbf{c}_t, \tilde{\mathbf{q}}) / \kappa)} \quad (1)$$

In the above equation, k represents the fixed temperature, and sim signifies the cosine similarity between context representations and quantized latent speech representations. The term L_m resembles the Softmax Function, but it employs cosine similarity instead of a score. For ease of optimization, the negative logarithm of the fraction is also applied. \mathbf{c}_t is the context network output centered over masked time step t and \mathbf{q}_t is the true quantized latent speech representation.

In this paper we use the version of the model pre-trained on cross-lingual speech data (53 different languages) [18] that was then fine-tuned [19] for the task of phoneme sequence prediction using a CTC loss, and the result is an open-source³.

The model itself thus predict sequences of phoneme labels, but no phone boundaries. What we propose in the following sections is an approach for leveraging the learned cross-lingual speech representations of this model that was already oriented to a phonetic space thanks to the CTC fine-tuning. The goal of this approach is to require limited amount of data and to be less biased towards American English accent compared to existing models.

3.3. Data Processing and Knowledge Transfer

Based on the representations that we can extract from the model explained in the previous section, we used PCA to reduce dimensionality while retaining most of the variance in the data. We use the last hidden layer before classification heads as our speech representations.

We aimed to preserve a substantial amount of variance of the data. Experimenting without PCA and with varying levels of variance retention (99%, 95%, and 90%), we determined that retaining 95% offers the most favorable compromise in terms of classification results. We hypothesized that this choice may serve as a means of noise filtration and provided a more manageable space for classifiers to process effectively. To train this dimension reduction model and the subsequent frame classifier model we used the MAILABS [23] dataset which is a large dataset of speech data, containing recordings from various languages and accents to train our reducer and classifier. We used the American and British English datasets from MAILABS.

We worked at the frame level for training the shallow dimension reduction and classifier models. As we observe from Figure 2, the frequency of phonemes in natural languages is not uniform. Hence, it is essential to perform data balancing on the data extracted from MAILABS for unbiased training. To perform data balancing, we randomly selected an equal number of latent frames in order to have a

² <https://huggingface.co/facebook/wav2vec2-xlsr-53-espeak-cv-ft>

³ <https://huggingface.co/facebook/wav2vec2-xlsr-53-espeak-cv-ft>



Our proposed approach, which integrates a reduction model with a classifier model can adeptly model intricate relationships stemming from the speech representation learning model. The task of mapping between these learned representations and phonetic units is efficiently managed by a compact and effective model. This ensures our ability to comprehend the complex relationships within the data while maintaining computational efficiency.

In evaluating the performance of our text independent phone-to-audio alignment model, we conducted a comprehensive comparison with the state-of-the-art model, charsiu [20]. In the charsiu model the authors compare two systems: text independent phone-to-audio alignment and text dependent phone-to-audio alignment using Wav2Vec2. The Wav2Vec2-FS is a semi-supervised model which learns alignment using contrastive learning and a forward sum loss. The second model, Wav2Vec2-FC, is a frame classification model trained on labels for forced alignment, capable of both forced alignment and

text-independent segmentation. The evaluation of both the systems for charsiu has been performed on the TIMIT dataset. It has been used for this evaluation because of the availability of human annotations especially phoneme level. We use a similar approach to evaluate the same metrics for our model. The assessment is based on statistical measures, namely precision, recall, F1 score, and r-value. In the referenced work, the authors present their best-performing model, W2V2-FC-32k-Libris, as can be observed in Table 1.

Table 1. Evaluation results of text independent alignment on TIMIT dataset

Model	P	R	F1	r-val
W2V2-CTC-10ms	0.31	0.29	0.30	0.42
W2V2-CTC-20ms	0.31	0.30	0.31	0.42
<i>Phone recognition + W2V2-FS</i>				
W2V2-FS-20ms	0.40	0.42	0.41	0.48
W2V2-FS-10ms	0.56	0.58	0.57	0.63
W2V2-FC-32k-Libris	0.57	0.57	0.57	0.64
<i>Direct inference</i>				
W2V2-FC-20ms-Libris	0.57	0.59	0.58	0.63
W2V2-FC-10ms-Libris	0.55	0.58	0.56	0.62
W2V2-FC-32k-Libris	0.60	0.63	0.61	0.66
<i>Our Proposed Model</i>				
W2V2-PCA-C	0.61	0.68	0.63	0.58

4.1. Text Independent Phone-to-Audio Alignment on TIMIT

In our case, for the task of phoneme recognition, we compare our model with the text independent W2V2-FC-10ms charsiu model. The statistical metrics for our proposed model outperform that of the charsiu model. We compare the TIMIT test dataset to assess our model and we observe from Table 1 that the r-value for our model deteriorates in performance. Apart from r-value, all the other metrics: Precision, Recall and F1 value have shown to perform well for our model.

The tests for our model are also performed on the TIMIT dataset but has the possibility to extend to other datasets with real speech and also other languages.

The r-value, which is known as the Pearson correlation coefficient measures the similarity between the ground truth phonemes and the predicted phonemes. Some of the reasons for low r-value could be alignment errors, variability in the pronunciation of speakers or changes in speaking style within the dataset. However, we need more experiments to confirm this hypothesis.

4.2. Text-Independent Phone-to-Audio Alignment on SCRIBE

A second part of the experiments in our proposed model is to evaluate it on British English.

The SCRIBE⁴ dataset is a small corpus of read speech and spontaneous speech specializing in British English. It consists of 200 'phonetically rich' sentences and 460 'phonetically compact' sentences. The 'phonetically rich' sentences are phonetically balanced. The 'phonetically compact' sentences are based on a British version of the MIT compact sentences (as in TIMIT). There are 45 files in the dataset of 30-50s containing 5-10 sentences.

The audio files and the phoneme annotations needed some processing before we could start using the dataset. The phonemes are in SAMPA form and needed to be converted to the English Arpabet. Even after the conversion from SAMPA to Arpabet, there were symbols that we could not retrieve so we filtered them out.

We evaluated the charsiu model and our proposed model using SCRIBE and achieved the metrics that are depicted in Table 2.

⁴ <https://www.phon.ucl.ac.uk/resource/scribe/>

Table 2. Evaluation results of text-independent alignment on SCRIBE dataset

Model	P	R	F1	r-val
charsiu Model				
W2V2-FC-10ms-Libris	0.93	0.71	0.80	0.79
Our Proposed Model				
W2V2-PCA-C	0.89	0.85	0.87	0.88

Upon close examination of the results presented in Table 2, it becomes evident that the metrics associated with our proposed model exhibit a greater uniformity in values when compared to charsiu. Furthermore, these metrics generally surpass those of charsiu, albeit with the exception of precision. The uniformity observed in our model’s metrics can be attributed to the utilization of a more balanced reduced frame dataset during training, which serves to mitigate biases and yield more consistent outcomes. Moreover, the superior quality of audio files within the SCRIBE dataset, resembling professional studio recordings with minimal noise, likely contributes to the heightened metrics observed for both charsiu and our proposed model when contrasted with the TIMIT dataset.

Consequently, our system demonstrates enhanced generalization across various accents, laying the foundation for potential expansion to other languages. The key lies in leveraging the underlying Wav2vec2 XLSR framework, and the methodology employed with MAILABS can seamlessly be replicated with different datasets, opening avenues for utilizing non-native English accents and broader linguistic applications.

5. Discussion

In this paper, we introduced an innovative text independent phone alignment system designed to be language-independent. Our method harnesses a self-supervised model (wav2vec2) fine-tuned for phoneme recognition through a CTC loss, alongside a dimensional reduction model (PCA) and a frame-level phoneme classifier. Through experiments using synthetic native data, we assessed our model’s performance on both American and British English, benchmarking it against the state-of-the-art charsiu model. Encouragingly, our results demonstrate robustness and applicability to diverse accents, such as British English.

However, certain limitations are acknowledged. Firstly, the reducer and classifier have been trained on a restricted amount of data specific to native English, posing a constraint. Secondly, our model relies on forced-aligned datasets of native speech for effective learning, introducing another limitation. Thirdly, the SCRIBE dataset that we used is a smaller dataset of British English as compared to the TIMTI dataset. Lack of well-annotated datasets, especially on phoneme level is one of our biggest challenges.

Future research directions could explore the incorporation of datasets containing non-native English data. This involves re-training the shallow reducer and classifier models using non-native speech data. Furthermore, extending our approach to languages beyond English and evaluating its performance with real-world data from language learners presents an intriguing avenue for exploration. Our work lays the foundation for further experiments in the realm of text independent phone-to-audio alignment, especially in the context of non-native English.

Funding: Part of this work was done during the project REDCALL that is partially funded by a FIRST Enterprise Docteur program from SPW Recherche⁵.
Part of this work was done during the project DEEPCALL that is partially funded by a Win4Doc program from SPW Recherche.

⁵ <https://recherche.wallonie.be/>

References

1. Chapelle, C.A. The relationship between second language acquisition theory and computer-assisted language learning. *The modern language journal* **2009**, *93*, 741–753.
2. Golonka, E.M.; Bowles, A.R.; Frank, V.M.; Richardson, D.L.; Freynik, S. Technologies for foreign language learning: A review of technology types and their effectiveness. *Computer assisted language learning* **2014**, *27*, 70–105.
3. Tits, N.; Broisson, Z. Flowchase: a Mobile Application for Pronunciation Training. Proc. 9th Workshop on Speech and Language Technology in Education (SLaTE), 2023, pp. 93–94.
4. Tan, C.; Sun, F.; Kong, T.; Zhang, W.; Yang, C.; Liu, C. A survey on deep transfer learning. International conference on artificial neural networks. Springer, 2018, pp. 270–279.
5. Jaiswal, A.; Babu, A.R.; Zadeh, M.Z.; Banerjee, D.; Makedon, F. A survey on contrastive self-supervised learning. *Technologies* **2020**, *9*, 2.
6. Zoph, B.; Yuret, D.; May, J.; Knight, K. Transfer learning for low-resource neural machine translation. *arXiv preprint arXiv:1604.02201* **2016**.
7. Tits, N.; El Haddad, K.; Dutoit, T. Exploring Transfer Learning for Low Resource Emotional TTS. Intelligent Systems and Applications; Bi, Y.; Bhatia, R.; Kapoor, S., Eds.; Springer International Publishing: Cham, 2020; pp. 52–60.
8. Tits, N.; Wang, F.; Haddad, K.E.; Pagel, V.; Dutoit, T. Visualization and Interpretation of Latent Spaces for Controlling Expressive Speech Synthesis through Audio Analysis. Proc. Interspeech 2019, 2019, pp. 4475–4479. doi:10.21437/Interspeech.2019-1426.
9. Tits, N.; El Haddad, K.; Dutoit, T. Analysis and assessment of controllability of an expressive deep learning-based tts system. Informatics. MDPI, 2021, Vol. 8, p. 84.
10. Tits, N.; El Haddad, K.; Dutoit, T. ASR-based Features for Emotion Recognition: A Transfer Learning Approach. Proceedings of Grand Challenge and Workshop on Human Multimodal Language (Challenge-HML). Association for Computational Linguistics, 2018, pp. 48–52.
11. Zhou, K.; Sisman, B.; Liu, R.; Li, H. Emotional voice conversion: Theory, databases and ESD. *Speech Communication* **2022**, *137*, 1–18.
12. Hu, W.; Qian, Y.; Soong, F.K.; Wang, Y. Improved mispronunciation detection with deep neural network trained acoustic models and transfer learning based logistic regression classifiers. *Speech Communication* **2015**, *67*, 154–166.
13. Povey, D.; Ghoshal, A.; Boulianne, G.; Burget, L.; Glembek, O.; Goel, N.; Hannemann, M.; Motlicek, P.; Qian, Y.; Schwarz, P.; others. The Kaldi speech recognition toolkit. IEEE 2011 workshop on automatic speech recognition and understanding. IEEE Signal Processing Society, 2011, number CONF.
14. Young, S.; Evermann, G.; Gales, M.; Hain, T.; Kershaw, D.; Liu, X.; Moore, G.; Odell, J.; Ollason, D.; Povey, D.; others. The HTK book. *Cambridge university engineering department* **2002**, *3*, 12.
15. Gorman, K.; Howell, J.; Wagner, M. Prosodylab-aligner: A tool for forced alignment of laboratory speech. *Canadian Acoustics* **2011**, *39*, 192–193.
16. Koizumi, T.; Mori, M.; Taniguchi, S.; Maruya, M. Recurrent neural networks for phoneme recognition. Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP'96. IEEE, 1996, Vol. 1, pp. 326–329.
17. Müller, M.; Stüker, S.; Waibel, A. Phonemic and graphemic multilingual ctc based speech recognition. *arXiv preprint arXiv:1711.04564* **2017**.
18. Conneau, A.; Baevski, A.; Collobert, R.; Mohamed, A.; Auli, M. Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979* **2020**.
19. Xu, Q.; Baevski, A.; Auli, M. Simple and effective zero-shot cross-lingual phoneme recognition. *arXiv preprint arXiv:2109.11680* **2021**.
20. Zhu, J.; Zhang, C.; Jurgens, D. Phone-to-audio alignment without text: A semi-supervised approach. ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2022, pp. 8167–8171.
21. Garofolo, J.S. Timit acoustic phonetic continuous speech corpus. *Linguistic Data Consortium*, 1993 **1993**.

22. Baevski, A.; Zhou, Y.; Mohamed, A.; Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems* **2020**, *33*, 12449–12460.
23. Solak, I. The M-AILABS speech dataset, 2019.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.