A Summary of the ComParE COVID-19 Challenges

Harry Coppock[*]
Imperial College London, UK

Alican Akman
Imperial College London, UK

Christian Bergler
FAU Erlangen-Nürnberg, Germany

Maurice Gerczuk
Universität Augsburg, Germany

Chloë Brown
University of Cambridge, UK

Jagmohan Chauhan
University of Southampton, UK

Andreas Grammenos
University of Cambridge, UK

Apinan Hasthanasombat
University of Cambridge, UK

Dimitris Spathis
University of Cambridge, UK

Tong Xia
University of Cambridge, UK

Pietro Cicuta
University of Cambridge, UK

Jing Han
University of Cambridge, UK

Shahin Amiriparian
Universität Augsburg, Germany

Alice Baird
Universität Augsburg, Germany

Lukas Stappen
Universität Augsburg, Germany

Sandra Ottl
Universität Augsburg, Germany

Panagiotis Tzirakis
Imperial College London, UK

Anton Batliner
Universität Augsburg, Germany

Cecilia Mascolo
University of Cambridge, UK

Björn W. Schuller
Imperial College London, UK
and University of Augsburg, Germany

February 21, 2022

Abstract

The COVID-19 pandemic has caused massive humanitarian and economic damage. Teams of scientists from a broad range of disciplines have searched for methods to help governments and communities combat the disease. One avenue from the machine learning field which has been explored is the prospect of a digital mass test which can detect COVID-19 from infected individuals' respiratory sounds. We present a summary of the results from the INTERSPEECH 2021 Computational Paralinguistics Challenges: COVID-19 Cough, (CCS) and COVID-19 Speech, (CSS).

Keywords COVID-19 · Machine Learning · Mass Test · Challenge · Speech · Cough

1 Introduction

Significant work has been conducted exploring the possibility that COVID-19 yields unique audio biomarkers in infected individuals' respiratory signals [8, 41, 20, 36, 5, 28, 27, 4, 29, 31, 6, 11, 26, 30]. This has shown promising results although many still remain sceptical, suggesting that models could simply be relying on spurious bias signals in the datasets [12, 11]. These worries have been supported by findings that when sources of bias are controlled, the performance of the classifiers decreases [17, 13]. Along with this, cross dataset experiments have reported a marked

---

[*]harry.coppock@imperial.ac.uk

drop in performance when models trained on one dataset are then evaluated on another dataset, suggesting dataset specific bias [2].

Last summer, the machine learning community were called upon to address some of these challenges, and help answer the question whether a digital mass test was possible, through the creation of two COVID-19 challenges within the Interspeech Computational Paralinguistics challengE (ComParE) series: COVID-19 Cough, (CCS) and COVID-19 Speech, (CSS)[35]. Contestants were tasked to create the best performing COVID-19 classifier from user cough and speech recordings. We note that another COVID-19 detection from audio challenge was run at a similar time to ComParE, named DiCOVA [25], and point the inquisitive reader to their blog post [2] which details a summary of the results.

2   Challenge Methodology

Both COVID-19 cough and speech challenges were binary classification tasks. Given an audio signal of a user coughing or speaking, challenge participants were tasked with predicting whether the respiratory signal came from a COVID-19 positive or negative user. After signing up to the challenge, teams were sent the audio files along with the corresponding labels for both the training and development set. Teams were also sent the audio files from the test set without the corresponding labels. Teams were allowed to submit five predictions for the test set from which the best score was taken. The number of submissions was limited to avoid overfitting to the test set.

The datasets used in these challenges are two curated subsets of the crowd sourced Cambridge COVID-19 Sounds database [8, 42]. COVID-19 status was self reported and determined through either a PCR or rapid antigen test. The number of samples of both positive and negative cases for these selected subsets are detailed in Table 1. The submission date for both COVID-19 positive and negative case recordings are detailed in Figure 1a. Figure 1b shows the age distribution for both CSS and CCS challenges.

Table 1: ComParE COVID-19 Sub-Challenges dataset splits. Values specify the number of audio recordings. We note that disjoint participant train, development, and test splits were ensured.

	CCS*			CSS †		
#	Train	Dev	Test	Train	Dev	Test
COVID-19-positive	71	48	39	72	142	94
COVID-19-negative	215	183	169	243	153	189
Total	286	231	208	315	295	283

*CCS – COVID-19 Cough Sub-Challenge
† CSS – COVID-19 Speech Sub-Challenge

3   Overview of Methodologies Used in Accepted Papers at Interspeech 2021

Last year, 44 teams registered in both the ComParE COVID-19 Cough Sub-Challenge (CCS) and the COVID-19 Speech Sub-Challenge (CSS) of which 19 submitted test set predictions. Five of the 19 teams submitted papers to INTERSPEECH which were then accepted. Results for both CCS and CSS were reported in two of these papers, while two papers reported results exclusively for CCS and one paper exclusively for CSS. In this section, we provide a brief overview of methodologies used in these accepted works which included data augmentation techniques, feature types, classifier types, and ensemble model strategies. Teams that did not have their work accepted at INTERSPEECH 2021 will be named NN_X to preserve anonymity. NN refers to *nomen nescio* and X is the order in which they appear in Figure 2.The performance measured in Unweighted Average Recall (UAR) achieved by these methodologies is summarised in Table 2; UAR has been used as a standard measure in the Computational Paralinguistics Challenges at Interspeech since 2009 [34]. It is the mean of the diagonal of the confusion matrices in percent and by that, fair towards sparse classes.Note that UAR is sometimes called 'macro-average', see [24].

3.1   Data Augmentation

To combat the limited size and imbalance of the Cambridge COVID-19 Sounds database, the majority of the teams used data augmentation techniques in their implementation. Team *Casanova, et al.* exploited a noise addition method and SpecAugment to augment the challenge dataset [9]. Team *Illium, et al.* targeted spectrogram-level augmentations with

---

[2]https://dicova2021.github.io

(a)   (b)

Figure 1: Figure 1a is a cumulative plot detailing when COVID-19 positive and negative submission to both the CCS and CSS were made. Figure 1b details the age distribution of COVID-19 positive and negative participants for the CCS and CSS Sub-Challenges

Table 2: Summary of methodologies used in accepted papers at Interspeech 2021

Team Name	Data Aug.	Feature Type	Classifiers	Ensemble	Cough UAR [%]	Speech UAR [%]
Solera-Urena, et al. [37]	✗	TDNN-F, VGGish, PASE+	SVM	✓	69.3	–
Casanova, et al. [9]	✓	MFCC, mel-spectrogram	SpiraNet, CNN14, ResNet-38, MobileNet	✓	**75.9**	70.3
Klumpp, et al. [22]	✓	mel-spectrogram	CNN, LSTM, SVM, LR	✗	–	64.2
Illium, et al. [19]	✓	mel-spectrogram	Vision Transformer	✗	72.0	–
Baseline [35]	✗	openSMILE, openXBOX, DiFE, DeepSpectrum, auDeep	SVM, End2You	✓	73.9	**72.1**

temporal shifting, noise addition, SpecAugment and loudness adjustment [19]. Instead of using a data augmentation method to manipulate the challenge dataset, team *Klumpp, et al.* used three auxiliary datasets in different languages aiming their deep acoustic model to better learn the properties of healthy speech [22].

3.2   Feature Type

The teams chiefly used spectrogram-level features including mel-frequency cepstral coefficients (MFCC) and mel-spectrograms. For higher-level features, the teams used the common feature extraction toolkits openSMILE [15], openXBOX [33], DeepSpectrum [3], and auDeep [16], where a simple support vector machine (SVM) model was built on top of these features. Team *Solera-Urena, et al.* exploited transfer learning to extract feature embeddings by using pre-trained TDNN-F[39], VGGish[18], and PASE+[32] models with appropriate fine-tuning on the challenge dataset. Team *Klumpp, et al.* targeted to extract their own phonetic features by using an acoustic model consisting of convolutional neural network (CNN) and long short-term memory (LSTM) parts.

3.3   Classifier Type

Team *Solera-Urena, et al.* [37] and the challenge baseline [35] fitted a SVM model to high level audio embeddings extracted using TDNN-F[39], VGGish[18], and PASE+[32] models, and the openSMILE framework [15], respectively.

While the challenge baseline [35] searched for the complexity parameter of the SVM ranging from 10^{-5} to 1, team *Solera-Urena, et al.* [37] explored different kernels (linear, RBF), data normalisations (zero mean and unit variance, [0,1] range) and class balancing methods (majority class downsampling, class weighting). In addition to the SVM model, the baseline explored using the multimodel profiling toolkit End2You[38] to train a recurrent neural network using Gated Recurrent Units (GRUs) with hidden units of 64. Team *Casanova, et al.* [9] utilised the deep models: SpiraNet [10], CNN14 [23], ResNet-38 [23], and MobileNetv1 [23] where they explored kernel size, convolutional dilatation, dropout, number of fully connected layer neurons, learning rate, weight decay and optimizer. Team *Klumpp, et al.* [22] trained SVM and logistic regression (LR) models to perform COVID-19 classification on top of phonetic features extracted by their deep acoustic model. They explored the complexity parameter of the SVM ranging from 10^{-4} to 1. Team *Illium, et al.* [19] adapted a vision transformer [14] for mel-spectrogram representations of audio signals. Tree-structured Parzen Estimator-algorithm (TPE) [1] was exploited in [19] for hyperparameter search mainly exploring embedding size, learning rate, batch size, dropout, number of heads and head dimension. The teams *Solera-Urena, et al.*, *Casanova, et al.*, and the baseline also reported classification results by using the fusion of their best features and classifiers.


Figure 2: Team performance on the held out test set for the COVID-19 Cough Sub-Challenge.

3.4 Assessment of Performance Measures

Figure 4 visualises a two-sided significance test (based on a Z-test concerning two proportions, [21], section 5B) employing the CCS and CSS Test sets and the corresponding baseline systems [35]. Various levels of significance (α-values) were used for calculating an absolute deviation with respect to the Test set, being considered as significantly better or worse than the baseline systems. Due to the fact that a two-sided test is employed, the α-values must be halved to derive the respective Z-score used to calculate the p-value of a model fulfilling statistical significance for both sides [21]. Consequently, significantly outperforming the best CCS baseline system (73.9 % and 208 Test set samples) at a significance level of $\alpha = 0.01$ requires at least an absolute improvement of 6.7 %; for CSS (best baseline system with 72.1 % and 283 Test set samples), the improvement required is 6.0 %. Note that Null-Hypothesis-Testing with p-values as criterion has been criticised from its beginning; see the statement of the American Statistical Association in [40] and [7].Therefore, we provide this plot with p-values as a service for readers interested in this approach, not as a guideline for deciding between approaches.

Another way of assessing performance measures as for their 'uncertainty' is computing confidence intervals (CIs). [35] employed two different CIs: first, 1000x bootstrapping for Test (random selection with replacement) and UARs based on the same model that was trained with Train and Dev; in the following, the CIs for these UARs are given first.
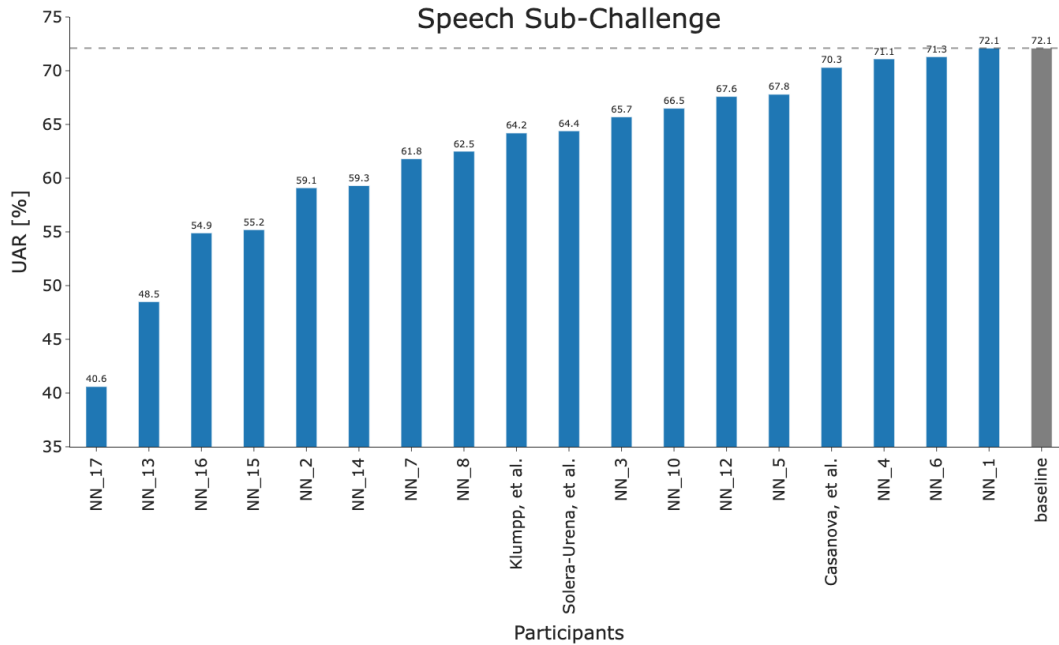
Figure 3: Team performance on the held out test set for the COVID-19 Speech Sub-Challenge.
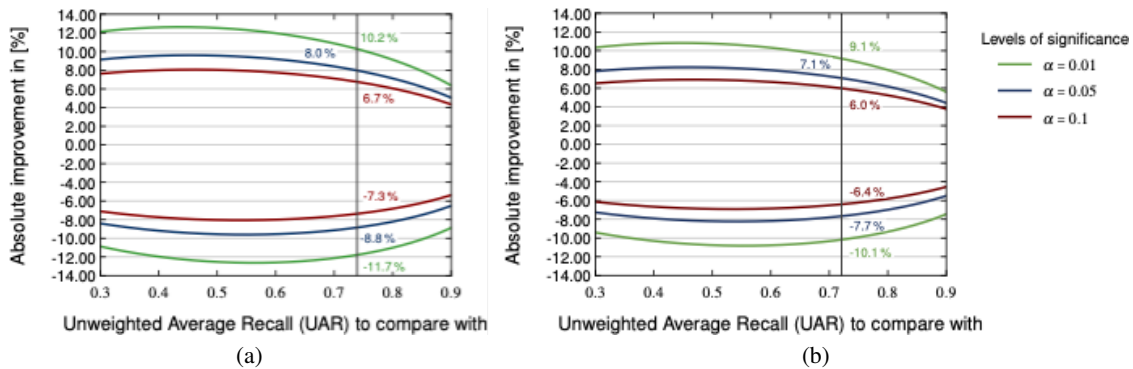

| (a) | (b) |

Figure 4: Two-sided significance test on the COVID-19 Cough (4a) and Speech (4b) Test sets with various levels of significance according to a two-sided Z-test.

Then, 100x bootstrapping for the corresponding combination of Train and Dev; the different models obtained from these combinations were employed to get UARs for Test and subsequently, CIs; these results are given in second place. Note that for this type of CI, the Test results are often above the CI, sometimes within and in a few cases below, as can be seen in [35]; obviously, reducing the variability of the samples in the training phase with bootstrapping results on average in somehow lower performance. For CCS with a UAR of 73.9 %, the first CI was 66.0 %-82.6 %; the second one could not be computed because this UAR is based on a fusion of different classifiers. For CSS with a UAR of 72.1 %, the CIs were 66.0 %-77.8 % and 70.2 %-71.1 %, respectively. Both Figure 4 and the spread of the CIs reported demonstrate the uncertainty of the results, caused by the relatively low number of data points in the test set.

4 Results and Discussion

Figures 2 and 3 detail the rankings for the 19 teams which submitted predictions for the test set. We congratulate [9] for winning the COVID-19 Cough Sub-Challenge with an UAR of 75.9 % on the held out test set. We note that for

the COVID-19 Speech Sub-Challenge, no team exceeded the performance of the baseline which scored 72.1 % UAR on the held out test set. To significantly outperform the baseline system for the cough modality, with a significance level of $\alpha = 0.1$, as detailed in Figure 4, would require an improvement of 6.7 %, an improvement which the winning submission fell short of by 4.7 %.

For both Sub-Challenges, teams struggled to outperform the baseline. Postulating why this could be the case one could suggest one, or a combination, of the following: COVID-19 detection from audio is a particularly hard task, the baseline score – being already a fusion of several state-of-the-art systems for CCS – represents a performance ceiling and that higher classification scores are not possible for this dataset, or, as a result of the limited size of the dataset, the task lends itself to less data hungry algorithms, such as the openSMILE-SVM baseline models for CSS.

It is important to analyse the level of agreement of COVID-19 detection between participant submissions. This is shown schematically in Figures 5 and 6. From these figures, we can see that there are clearly COVID-19 positive cases which teams across the board are able to correctly predict, but there are also positive COVID-19 cases which all teams have missed. These findings are reflected in the minimal performance increase of 0.3 % and 0.8 % for cough and speech tasks, respectively, obtained when fusing n best submission predictions through majority voting schemes. The results from fusing n best models using majority voting are detailed in Figures 8 and 9. This suggests that models from all teams are depending on similar audio features when predicting COVID-19 positive cases.

Figures 5 and 6 b and c detail the level of agreement across submissions for curated subset of the test set, where participants were selected if they were displaying at least one symptom (b) and when they were displaying no symptoms (c). These figures can be paired with Figure 7 which details the recall scores for positive cases across these same curated test sets. From this analysis, it does not appear that there was a trend across teams to perform favourably on cases where symptoms were being displayed or visa versa. While this does not disprove worries that these algorithms are simply cough or symptom identifiers, it does not add evidence in support of this claim.

4.1 Limitations

While this challenge was an important step in exploring the possibilities of a digital mass test for COVID-19, it has a number of limitations. A clear limiting factor of the challenge was the small size of the dataset. While many participants addressed this through data augmentation and regularisation techniques, it restricted the extent to which conclusions could be taken from the results, particularly investigating teams' performance on carefully controlled subsets of the data. We look forward to the newly released COVID-19 sounds dataset [42] which represents a vastly greater source of COVID-19 samples.

A further limitation of this challenge is the unforeseen correlation between low sample rate recordings, below 12 kHz, and COVID-19 status. In fact all low sample rate recordings in the challenge for both CCS and CSS were COVID-19 positive. For CCS and CSS there were 30 and 37 low sample rate cases, respectively. The reason for this is that at the start of the study the label in the survey for COVID-19 negative was unclear, and could have been interpreted as either 'not tested' or 'tested negative'. For this reason no negative samples from the time period were used. This can be seen in 1a. This early version of data collection also correlated with the study allowing for lower sample rate recordings, a feature which later was changed to restrict submissions to higher sample rates. This resulted in all the low sample rate recordings being COVID-19 positive. As can be seen in Figures 10, 11, 12 and 13, teams' trained models were able to pick up on the sample rate bias, with most teams correctly predicting all the low sample rate cases as COVID-19 positive. When this is controlled for and low sample rate recordings are removed from the test set, as shown in Figures 12 and 13, teams' performances drop significantly. For the challenge baselines this too was the case, with the fusion of baseline models for CCS falling from 73.8% to 68.6% UAR and the opensmile-SVM baseline for CSS dropping from 72.1% to 70.9% UAR. This is a great example of the effect of overlooked bias which expresses itself as an identifiable audio feature, leading to inflated classification scores. We regret that this was not found earlier.

As with most machine learning methods, it still remains unclear how to interpret the decision making process at inference time. This results in it being tricky to determine which acoustic features the model is correlating with COVID-19. Whether that be true, acoustic features caused by the COVID-19 infection or other acoustic bias is still an unanswered question [12]. We also note that this is a binary classification task, in that models only had to decide between COVID-19 positive or negative. This 'closed word fallacy' [7] leads to inflated performance as models are not tasked with discerning between confounding symptoms such as heavy cold or asthma. Tasking models to predict COVID-19 out of a wide range of possible conditions/symptoms would be a harder task.

In this challenge, participants were provided with the test set recordings (without the corresponding labels). In future challenges, test set instances should be kept private, requiring participants to submit trained models along with pipeline scripts for inference. Teams' test set predictions can then be run automatically by the challenge organisers. This will

help in reducing the possibility of overfitting and foul play. We note that there was no evidence of foul play, e. g., training in an unsupervised manner on the test set, in this challenge.

Another limitation of this challenge was the lack of meta data that organisers could provide to participants. This tied teams' hands to some extent in evaluating for themselves the level of bias in the dataset and so their opportunity to implement methods to combat it. This was not a desired feature. However, we now point teams towards the newly open sourced COVID-19 Sounds database [42] which also provides collected meta data. It is this dataset from which a subset of samples was taken for this challenge.

5   Conclusion

This challenge demonstrated that there is a signal in crowdsourced COVID-19 respiratory sounds that allows for machine learning algorithms to fit a classifier which achieves moderate detection rates of COVID-19 in infected individuals' respiratory sounds. Exactly what this signal is, however, still remains unclear. Whether these signals are truly audio biomarkers in respiratory sounds of infected individuals uniquely caused by COVID-19 or rather identifiable bias in the datasets, such as confounding flu like symptoms, is still an open question to be answered next.
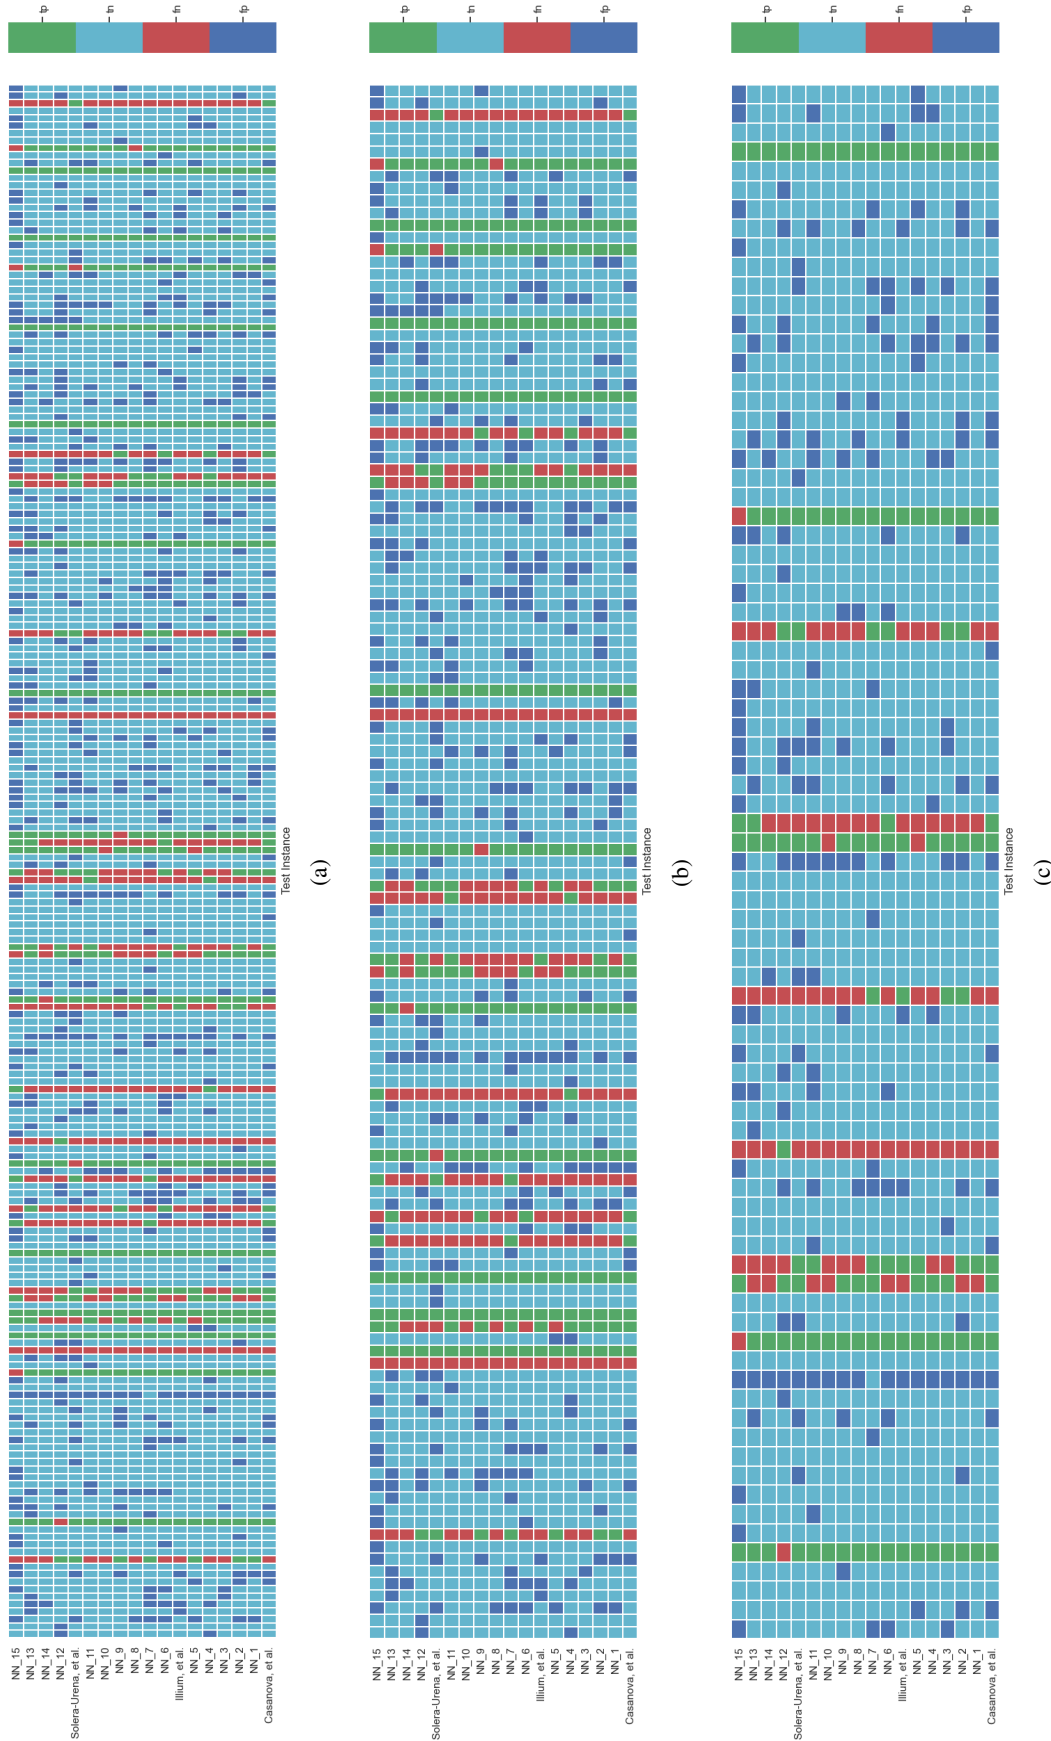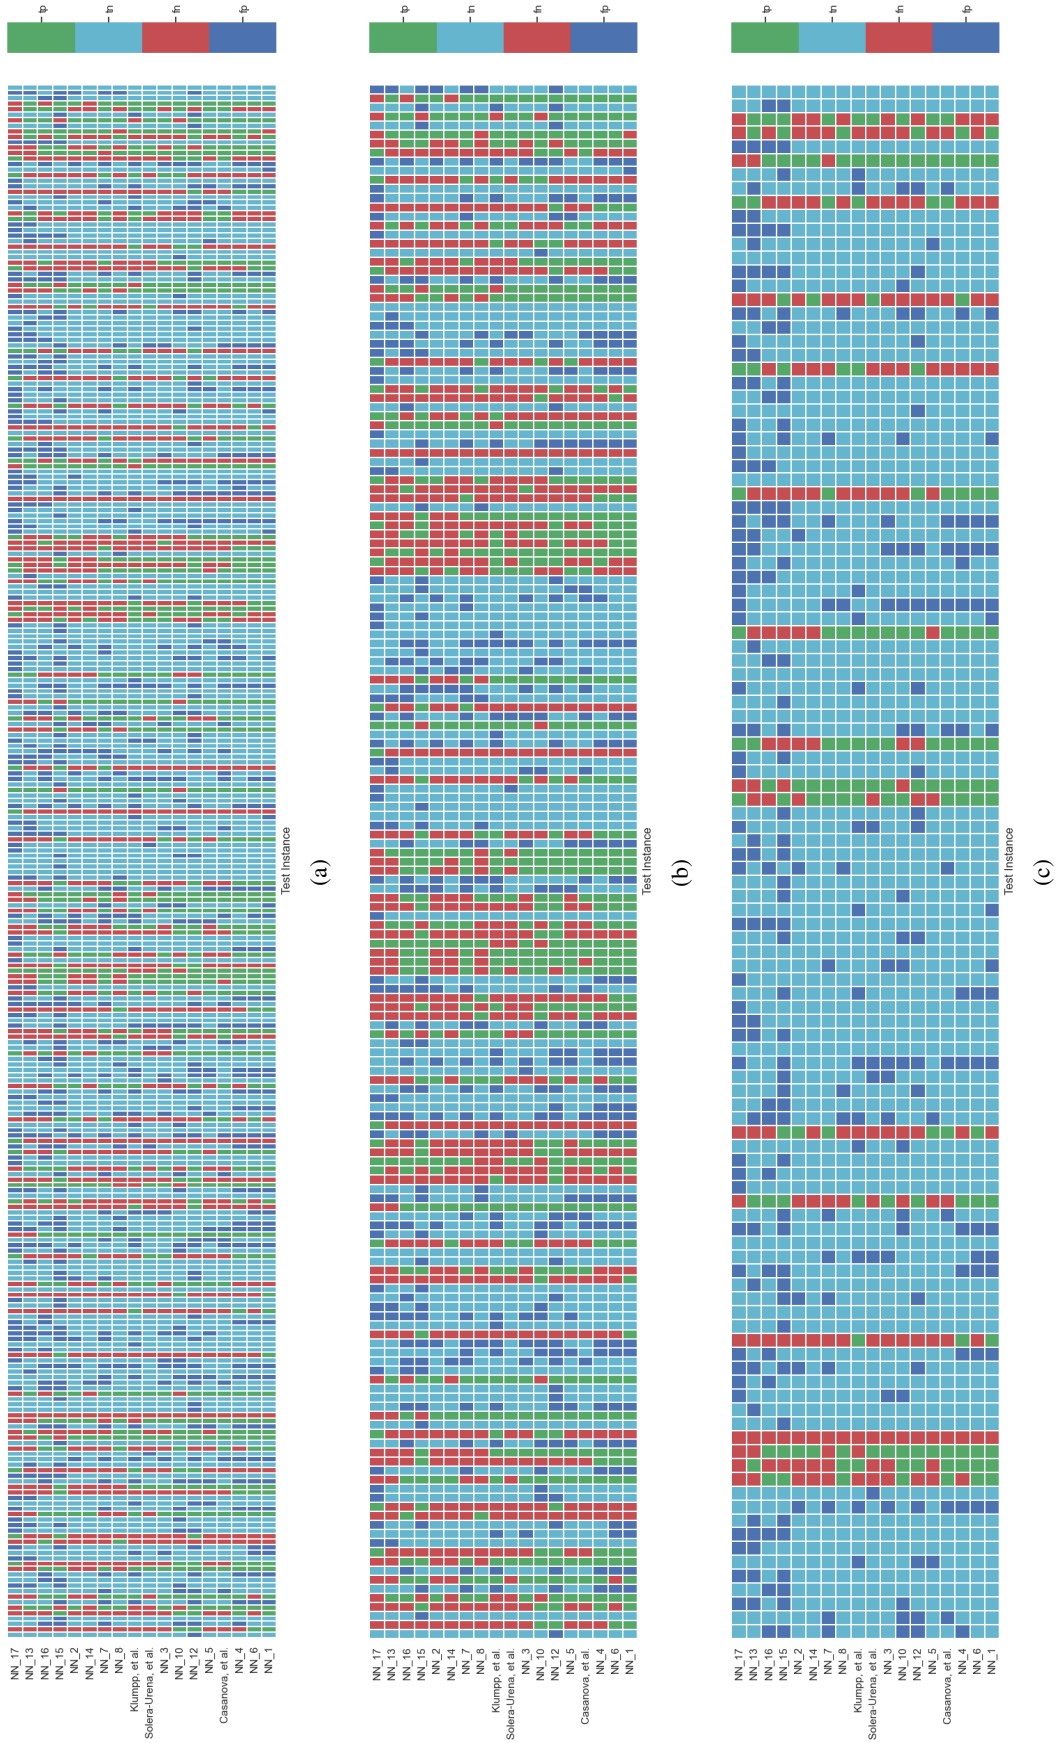
Figure 5: Schematic detailing the level of agreement between teams for each test instance for the **COVID-19 Cough Sub-Challenge.** Each row represents a team's submission results. The teams have been ordered by Unweighted Average Recall, from the bottom up (team *Casanova, et al.*'s predictions represent the highest scoring submission). Each column represents all teams predictions, across the competition, for one test instance. The test instances appear in the order in which they are in the test set. 5a details all the test instances, 5b details only the test instances which were experiencing symptoms at the time of recording, and 5c details only the test instances which were experiencing no symptoms at the time of recording.

Figure 6: Schematic detailing the level of agreement between teams for each test instance for the **COVID-19 Speech Sub-Challenge**. Each row represents a team's submission results. The teams have been ordered by Unweighted Average Recall (UAR), from the bottom up (team *yoshiharuyamamoto*'s predictions represent the highest scoring submission). Each column represents all teams' predictions, across the competition, for one test instance. The test instances appear in the order which they are in the test set. note: There are more test cases in the COVID-19 Speech Sub-Challenge than in the COVID-19 Cough Sub-Challenge. 6a details all the test instances, 6b details only the test instances which were experiencing symptoms at the time of recording, and 6c details only the test instances which were experiencing no symptoms at the time of recording.

6   Acknowledgements

References

[1] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. Optuna: A next-generation hyperparameter optimization framework. *CoRR*, abs/1907.10902, 2019. URL `http://arxiv.org/abs/1907.10902`.

[2] A. Akman, H. Coppock, A. Gaskell, P. Tzirakis, L. Jones, and B. W. Schuller. Evaluating the COVID-19 Identification ResNet (CIdeR) on the INTERSPEECH COVID-19 from Audio Challenges, 2021.

[3] S. Amiriparian, M. Gerczuk, S. Ottl, N. Cummins, S. Pugachevskiy, and B. Schuller. Bag-of-deep-features: Noise-robust deep feature representations for audio analysis. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7, 2018. doi:10.1109/IJCNN.2018.8489416.

[4] J. Andreu-Perez, H. Perez-Espinosa, E. Timonet, M. Kiani, M. I. Giron-Perez, A. B. Benitez-Trinidad, D. Jarchi, A. Rosales, N. Gkatzoulis, O. F. Reyes-Galaviz, A. Torres, C. Alberto Reyes-Garcia, Z. Ali, and F. Rivas. A Generic Deep Learning Based Cough Analysis System from Clinically Validated Samples for Point-of-Need COVID-19 Test and Severity Levels. *Transactions on Services Computing*, pages 1–1, 2021. doi:10.1109/TSC.2021.3061402.

[5] P. Bagad, A. Dalmia, J. Doshi, A. Nagrani, P. Bhamare, A. Mahale, S. Rane, N. Agarwal, and R. Panicker. Cough Against COVID: Evidence of COVID-19 Signature in Cough Sounds, 2020.

[6] K. D. Bartl-Pokorny, F. B. Pokorny, A. Batliner, S. Amiriparian, A. Semertzidou, F. Eyben, E. Kramer, F. Schmidt, R. Schönweiler, M. Wehler, and B. W. Schuller. The voice of COVID-19: Acoustic correlates of infection in sustained vowels. *The Journal of the Acoustical Society of America*, page 4377–4383, 2021. doi:10.1121/10.0005194.

[7] A. Batliner, S. Hantke, and B. Schuller. Ethics and Good Practice in Computational Paralinguistics. *Transactions on Affective Computing*, pages 1–1, 2020.

[8] C. Brown, J. Chauhan, A. Grammenos, J. Han, A. Hasthanasombat, D. Spathis, T. Xia, P. Cicuta, and C. Mascolo. Exploring Automatic Diagnosis of COVID-19 from Crowdsourced Respiratory Sound Data. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 3474–3484, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3394486.3412865.

[9] E. Casanova, A. Candido Jr., R. C. Fernandes Jr., M. Finger, L. R. S. Gris, M. A. Ponti, and D. P. Pinto da Silva. Transfer Learning and Data Augmentation Techniques to the COVID-19 Identification Tasks in ComParE 2021. In *Proc. Interspeech 2021*, pages 446–450, 2021. doi:10.21437/Interspeech.2021-1798.

[10] E. Casanova, L. Gris, A. Camargo, D. da Silva, M. Gazzola, E. Sabino, A. Levin, A. Candido Jr, S. Aluisio, and M. Finger. Deep learning against COVID-19: Respiratory insufficiency detection in Brazilian Portuguese speech. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 625–633. Association for Computational Linguistics, 2021. doi:10.18653/v1/2021.findings-acl.55.

[11] H. Coppock, A. Gaskell, P. Tzirakis, A. Baird, L. Jones, and B. W. Schuller. End-to-end convolutional neural network enables COVID-19 detection from breath and cough audio: a pilot study. *BMJ Innovations*, pages 356–362, 2021. doi:10.1136/bmjinnov-2021-000668.

[12] H. Coppock, L. Jones, I. Kiskin, and B. W. Schuller. COVID-19 Detection from Audio: Seven Grains of Salt. *The Lancet Digital Health*, 2021. doi:10.1016/S2589-7500(21)00141-2.

[13] H. Coppock, L. Jones, I. Kiskin, and B. W. Schuller. Bias and privacy in AI's cough-based COVID-19 recognition – Authors' reply. *The Lancet Digital Health*, 2021. doi:10.1016/S2589-7500(21)00233-8.

[14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.

[15] F. Eyben, F. Weninger, F. Gross, and B. Schuller. Recent developments in opensmile, the munich open-source multimedia feature extractor. In *Proceedings of the 21st ACM International Conference on Multimedia*, MM '13, page 835–838, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450324045. doi:10.1145/2502081.2502224. URL `https://doi.org/10.1145/2502081.2502224`.

[16] M. Freitag, S. Amiriparian, S. Pugachevskiy, N. Cummins, and B. Schuller. audeep: Unsupervised learning of representations from audio with deep recurrent neural networks. *Journal of Machine Learning Research*, 2018.

[17] J. Han, T. Xia, D. Spathis, E. Bondareva, C. Brown, J. Chauhan, T. Dang, A. Grammenos, A. Hasthanasombat, A. Floto, P. Cicuta, and C. Mascolo. Sounds of COVID-19: exploring realistic performance of audio-based digital testing, 2021.

[18] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson. Cnn architectures for large-scale audio classification, 2017.

[19] S. Illium, R. Müller, A. Sedlmeier, and C.-L. Popien. Visual Transformers for Primates Classification and Covid Detection. In *Proc. Interspeech 2021*, pages 451–455, 2021. doi:10.21437/Interspeech.2021-273.

[20] A. Imran, I. Posokhova, H. N. Qureshi, U. Masood, M. S. Riaz, K. Ali, C. N. John, M. I. Hussain, and M. Nabeel. AI4COVID-19: AI enabled preliminary diagnosis for COVID-19 from cough samples via an app. *Informatics in Medicine Unlocked*, page 100378, 2020. doi:10.1016/j.imu.2020.100378.

[21] E. Isaac. Test of Hypothesis - Concise Formula Summary, 2015.

[22] P. Klumpp, T. Bocklet, T. Arias-Vergara, J. Vásquez-Correa, P. Pérez-Toro, S. Bayerl, J. Orozco-Arroyave, and E. Nöth. The Phonetic Footprint of Covid-19? In *Proc. Interspeech 2021*, pages 441–445, 2021. doi:10.21437/Interspeech.2021-1488.

[23] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley. Panns: Large-scale pretrained audio neural networks for audio pattern recognition, 2020.

[24] C. D. Manning, P. Raghavan, and H. Schütze. *An Introduction to Information Retrieval*. Cambridge University Press, Cambridge, 2009.

[25] A. Muguli, L. Pinto, N. R., N. Sharma, P. Krishnan, P. K. Ghosh, R. Kumar, S. Ramoji, S. Bhat, S. R. Chetupalli, S. Ganapathy, and V. Nanda. DiCOVA Challenge: Dataset, task, and baseline system for COVID-19 diagnosis using acoustics, 2021.

[26] M. A. Nessiem, M. M. Mohamed, H. Coppock, A. Gaskell, and B. W. Schuller. Detecting COVID-19 from Breathing and Coughing Sounds using Deep Neural Networks. In *International Symposium on Computer-Based Medical Systems (CBMS)*, pages 183–188, Aveiro, Portugal, 2021. IEEE. doi:10.1109/CBMS52027.2021.00069.

[27] L. Orlandic, T. Teijeiro, and D. Atienza. The COUGHVID crowdsourcing dataset, a corpus for the study of large-scale cough analysis algorithms. *Scientific Data*, page 156, 2021. doi:10.1038/s41597-021-00937-4.

[28] G. Pinkas, Y. Karny, A. Malachi, G. Barkai, G. Bachar, and V. Aharonson. SARS-CoV-2 Detection From Voice. *Open Journal of Engineering in Medicine and Biology*, pages 268–274, 2020. doi:10.1109/OJEMB.2020.3026468.

[29] D. T. Pizzo, S. Santiago Esteban, and M. de los Ángeles Scetta. IATos: AI-powered pre-screening tool for COVID-19 from cough audio samples, 2021.

[30] A. Ponomarchuk, I. Burenko, E. Malkin, I. Nazarov, V. Kokh, M. Avetisian, and L. Zhukov. Project achoo: A practical model and application for covid-19 detection from recordings of breath, voice, and cough. *IEEE Journal of Selected Topics in Signal Processing*, pages 1–1, 2022. doi:10.1109/JSTSP.2022.3142514.

[31] K. Qian, M. Schmitt, H. Zheng, T. Koike, J. Han, J. Liu, W. Ji, J. Duan, M. Song, Z. Yang, Z. Ren, S. Liu, Z. Zhang, Y. Yamamoto, and B. W. Schuller. Computer Audition for Fighting the SARS-CoV-2 Corona Crisis – Introducing the Multi-task Speech Corpus for COVID-19. *Internet of Things Journal*, pages 1–1, 2021. doi:10.1109/JIOT.2021.3067605.

[32] M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio. Multi-task self-supervised learning for robust speech recognition, 2020.

[33] M. Schmitt and B. Schuller. openxbow – introducing the passau open-source crossmodal bag-of-words toolkit. *Journal of Machine Learning Research*, 2017.

[34] B. Schuller, S. Steidl, and A. Batliner. The INTERSPEECH 2009 Emotion Challenge. In *Proc. of INTERSPEECH*, pages 312–315, Brighton, 2009.

[35] B. W. Schuller, A. Batliner, C. Bergler, C. Mascolo, J. Han, I. Lefter, H. Kaya, S. Amiriparian, A. Baird, L. Stappen, S. Ottl, M. Gerczuk, P. Tzirakis, C. Brown, J. Chauhan, A. Grammenos, A. Hasthanasombat, D. Spathis, T. Xia, P. Cicuta, L. J. M. Rothkrantz, J. Zwerts, J. Treep, and C. Kaandorp. The interspeech 2021 computational paralinguistics challenge: Covid-19 cough, covid-19 speech, escalation & primates. In *INTERSPEECH 2021, 22nd Annual Conference of the International Speech Communication Association*, Brno, Czechia, 2021.

[36] N. Sharma, P. Krishnan, R. Kumar, S. Ramoji, S. R. Chetupalli, N. R., P. K. Ghosh, and S. Ganapathy. Coswara — A Database of Breathing, Cough, and Voice Sounds for COVID-19 Diagnosis. In *Proceedings INTERSPEECH 2020*, pages 4811–4815, Shanghai, China, 2020. ISCA. doi:10.21437/Interspeech.2020-2768.

[37] R. Solera-Ureña, C. Botelho, F. Teixeira, T. Rolland, A. Abad, and I. Trancoso. Transfer Learning-Based Cough Representations for Automatic Detection of COVID-19. In *Proc. Interspeech 2021*, pages 436–440, 2021. doi:10.21437/Interspeech.2021-1702.

[38] P. Tzirakis, S. Zafeiriou, and B. W. Schuller. End2you – the imperial toolkit for multimodal profiling by end-to-end learning, 2018.

[39] J. Villalba, N. Chen, D. Snyder, D. Garcia-Romero, A. McCree, G. Sell, J. Borgstrom, L. P. García-Perera, F. Richardson, R. Dehak, P. A. Torres-Carrasquillo, and N. Dehak. State-of-the-art speaker recognition with neural network embeddings in nist sre18 and speakers in the wild evaluations. *Computer Speech & Language*, 60:101026, 2020. ISSN 0885-2308. doi:https://doi.org/10.1016/j.csl.2019.101026. URL `https://www.sciencedirect.com/science/article/pii/S0885230819302700`.

[40] R. L. Wasserstein and N. A. Lazar. The ASA's Statement on p-values: Context, Process, and Purpose. *The American Statistician*, pages 129–133, 2016.

[41] T. Xia, J. Han, L. Qendro, T. Dang, and C. Mascolo. Uncertainty-Aware COVID-19 Detection from Imbalanced Sound Data, 2021.

[42] T. Xia, D. Spathis, C. Brown, A. G. J Ch, J. Han, A. Hasthanasombat, E. Bondareva, T. Dang, A. Floto, P. Cicuta, and C. Mascolo. Covid-19 sounds: A large-scale audio dataset for digital covid-19 detection. *NeurIPS 2021 Track Datasets and Benchmarks Round2 Submission*, 2021. Open Review *https://openreview.net/forum?id=9KArJb4r5ZQ*.

7 Appendix

Here we present some results from ablation studies of teams' performances through evaluating performance on curated subsets of the test set. Figure 7 details the effect of controlling for symptom cofounders on teams' performance. Figures 12 and 13 repeats this analysis however controlling for sample rate. Figures 10 and 11 details the level of agreement between teams for the low 10a, 11a and high 10b 11b sample rate test cases. Figures 8 and 9 detail the classification performance of a fusion of teams' predictions on the test set.

(a)

(b)

Figure 7: Team performance on the full test set (NoControl) and two curated test sets featuring only test instances where the participants either had at least one symptom (AnySymptoms) or were displaying no symptoms at all (NoSymptoms). The metric reported is recall for positive cases. 95 % confidence intervals are shown, calculated via the normal approximation method. 7a corresponds to the COVID-19 Cough Sub-Challenge, CCS, and 7a the COVID-19 Speech Sub-Challenge, CSS.

Figure 8: The performance of the fusion model of n-best models for the COVID-19 Cough Sub-Challenge using majority voting.


Figure 9: The performance of the fusion model of n-best models for the COVID-19 Speech Sub-Challenge using majority voting..

Figure 10: Schematic detailing the level of agreement as in 5 with test instances with either a low sample rate (below 12 kHz) (10a) or high sample rate (above 12 kHz) (10b).

Figure 11: Schematic detailing the level of agreement as in 6 with test instances with either a low sample rate (below 12 kHz) (11a) or high sample rate (above 12 kHz) (11b).

(a)

(b)

Figure 12: Team performance on two curated test sets from the COVID-19 Cough Sub-Challenge. 12a controls for test samples with a sample rate of greater than 12 kHz and 12b controls for test samples with a sample rate of 12 kHz and below. The metric reported is recall for positive cases. 95 % confidence intervals are shown, calculated via the normal approximation method.

(a)

(b)

Figure 13: Team performance on two curated test sets from the COVID-19 Speech Sub-Challenge. 13a controls for test samples with a sample rate of greater than 12 kHz and 13b controls for test samples with a sample rate of 12 kHz and below. The metric reported is recall for positive cases. 95 % confidence intervals are shown, calculated via the normal approximation method.