

ASVspoof 2021: Automatic Speaker Verification Spoofing and Countermeasures Challenge Evaluation Plan*

Héctor Delgado, Nicholas Evans, Tomi Kinnunen, Kong Aik Lee,
Xuechen Liu, Andreas Nautsch, Jose Patino, Md Sahidullah,
Massimiliano Todisco, Xin Wang, Junichi Yamagishi

ASVspoof consortium
<http://www.asvspoof.org/>

July 16, 2021

TL;DR for new participants

- Your task is to develop a *bona fide* - *spoofed* classifier (spoofing countermeasure) for speech data.
- We provide labelled training/development data and baseline code. You apply your system to unlabelled test data (one score per audio file) and send your scores to us, along with a system description.
- We rank and analyse the results, and present a summary at an INTERSPEECH 2021 satellite workshop. You are encouraged to submit a paper.

... and for readers familiar with ASVspoof

- ASVspoof 2021 involves technically more challenging data to promote countermeasure generalisation:
 - A **logical access** (LA) task similar to those of ASVspoof 2015 and 2019, but involving the **coding and transmission** of text-to-speech (TTS) and voice conversion (VC) attacks;
 - A **physical access** (PA) task involving **real** bona fide and replayed samples similar to ASVspoof 2017, but with a **better controlled design** (similar to the ASVspoof 2019 PA task).
- A **new speech deepfake task** (DF), involving compressed audio similar to the LA task (but without ASV).
- **Protocols:** Training/development partitions are the same as for ASVspoof 2019. The evaluation partition will be new.
- **Metrics:** for the LA and PA tasks, a slightly revised t-DCF metric [1]. For the DF task, an equal error rate (EER) metric.

*Document Version 0.4 (July 16, 2021)

1 Introduction

The automatic speaker verification spoofing and countermeasures (ASVspoof) challenge series is a community-led initiative which aims to promote the consideration of spoofing and the development of countermeasures. ASVspoof 2021 is the 4th in a series of bi-annual, competitive challenges where the goal is to develop countermeasures capable of discriminating between bona fide and spoofed or deepfake speech.

ASVspoof 2021 comprises three different tasks:

- logical access (LA): bona fide and spoofed utterances generated using text-to-speech (TTS) and voice conversion (VC) algorithms are communicated across telephony and VoIP networks with various coding and transmission effects;
- physical access (PA): bona fide utterances are made in a real, physical space in which spoofing attacks are captured and then replayed within the same physical space using replay devices of varying quality;
- speech deepfake (DF): a fake audio detection task comprising bona fide and spoofed utterances generated using TTS and VC algorithms. Similar to the LA task (includes compressed data) but without speaker verification.

The LA and PA tasks maintain the focus of ASVspoof upon automatic speaker verification (ASV). Accordingly, the metric for these two tasks will be the minimum tandem decision cost function (min t-DCF) [1]. The new DF task has a fake media / fake audio / deepfake flavour in which there is no ASV system. The metric for the DF condition will revert to the equal error rate (EER).

This document provides a technical description of the ASVspoof 2021 challenge, including details of training, development and evaluation data, metrics, baselines, evaluation rules, submission procedures and the schedule. It is intentionally terse. Many aspects of the ASVspoof 2019 are identical to previous editions of ASVspoof. The similarities are not repeated. The focus in this document is instead upon the key details and differences to previous editions.

2 Technical objectives

While the promotion of *generalised* countermeasures remains, the specific objectives for ASVspoof 2021 are centred around the shift towards more practical, realistic scenarios. Moving beyond the technical focus of previous editions, the objectives for the 2021 challenge are to:

- promote the development of spoofing countermeasures that are robust to channel variability;
- improve the reliability of spoofing countermeasures in less simulated, more realistic conditions, i.e. to recordings (re-)captured in real physical spaces;
- evaluate impacts of data augmentation to replace in-domain training and development data;
- broaden the relevance of ASVspoof to non-ASV scenarios involving the detection of deepfakes.

3 Logical Access (LA) task

Given the relevance to telephony scenarios, the focus of ASVspooF 2021 for the LA task is upon the development of spoofing countermeasures that are robust to codec and transmission channel variability. Bona fide and spoofed speech data generated with either text-to-speech (TTS), voice conversion (VC) or hybrid algorithms (voice conversion systems fed with synthetic speech) is transmitted across either a public switched telephone network (PSTN) or a voice over Internet Protocol (VoIP) network using some particular codec, prior to being processed by spoofing countermeasure and automatic speaker verification systems. Evaluation data will include trials transmitted across VoIP networks that use aLaw and G.722 codecs, though other codecs will also be used. Speech data may hence exhibit coding and transmission artefacts as well as differences in bandwidths and sampling frequencies, in addition to artefacts related to spoofing, but no additive noise. All data, no matter what the bandwidth, will be provided in a common 16 kHz, 16 bits per sample FLAC format. Neither new training nor new development data will be distributed; participants will receive new evaluation data only. Spoofing countermeasure systems should be trained and optimised using ASVspooF 2019 training and development data.

The TTS, VC and hybrid spoofing attacks will be the same as those for the ASVspooF 2019 LA evaluation partition. Codec meta data will not be provided for individual utterance. The task will hence be to design spoofing countermeasures that are capable of discriminating between bona fide and spoofed speech generated with known attacks when both contain additional nuisance variation relating to potential coding and transmission. The primary metric will be the tandem detection cost function (t-DCF).

4 Physical Access (PA) task

The objective of the PA task is to develop countermeasures for the detection of replay spoofing attacks. In keeping with the ASVspooF 2019 challenge, participants should use the ASVspooF 2019 PA training and development data which contains simulated replay speech for countermeasure training and optimisation. Accordingly, training and development protocols for the PA task are identical to those for ASVspooF 2019. As for the LA task, no new training or development data will be released.

New evaluation data will be distributed. This will contain predominantly *real* replayed speech, in addition to a smaller proportion of *simulated* replayed speech. Participants must therefore make use of simulated training and development replay data to train and optimise countermeasures which perform reliably in the case of real replay data. Real replay speech in the evaluation set is influenced by several environmental and attacker factors similar to, but more comprehensive to those of the ASVspooF 2019 challenge data.

The principal environmental factors include:

- ASV room size S_{asv} (m^2): $10 \leq S_{asv} \leq 65$
- Talker-to-ASV distance D_s (cm): $50 \leq D_s \leq 200$
- ASV microphone $Q_{asv,m}$: MEMS and condenser microphones

and the attacker factors include:

- Attacker room size S_a
- Attacker-to-talker distance D_a




Figure 1: An illustration of the ASVspooF 2021 physical access task. Rooms for replay acquisition (i.e., rooms in which the attacker captures recordings of the talker) and presentation (i.e., the room in which is situated the ASV system) may be different. In all cases the talker is simulated by a high quality loudspeaker.

- Attacker microphone $Q_{a,m}$,
- Replay device quality $Q_{a,s}$, and
- Replay-device to ASV distance D'_s

These factors are illustrated in Figure 1. The attacker factors are unknown. Neither meta information about the factors nor their labels will be provided. As a result of recordings being collected in real physical spaces (rooms), all data contains low levels of background noise. Audio files will be delivered in the same FLAC file format as for the LA task.

5 Speech deepfake (DF) task

Similar to the LA task, the DF task involves discrimination of bona fide speech from artificially generated or converted speech that may have undergone compression. In contrast to the LA task, however, (1) the DF task represents general audio compression (rather than telephony); and (2) there is no ASV system. While no specific use case is implied, the task is potentially relevant to social media or forensics applications and situations where the aim of an adversary is to fool a human listener rather than an ASV system.

The known compression algorithms include `mp3` and `m4a` while the evaluation data may also contain additional techniques. Here *compression* refers to compression-decompression; the audio

files will be delivered in the same **FLAC** file format as for other tasks, without meta data detailing the compression or processing applied. Since the aim is to promote the development of methods for the detection of speech deepfakes across multiple and unknown conditions, information neither about the compressor (or presence/absence of compression) nor their settings (such as bitrate) will be provided. Similar to the LA task, the data originates primarily from the ASVspoof 2019 LA data. Additionally, there will be new data for the purpose of evaluating robustness to domain mismatch.

6 Performance measures

Performance measures estimate the Bayes risk, and thus reflect the expected cost resulting from using a countermeasure (CM) in product and service operation. For legacy only, we refer to the equal-error rate (EER) of a CM—we acknowledge that EER reporting is deprecated by ISO/IEC standards on performance testing in biometrics and presentation attack detection (CMs in our case) [2], [3]. To estimate Bayes risk, one must quantify beliefs in class occurrence likelihoods (*i.e.*, priors) and beliefs in costs resulting from erroneous class predictions (*i.e.*, costs). The priors and costs shown in Table 1 are defined in the Appendix together with details of the min t-DCF computation. The ASV threshold is chosen to resemble the worst-case decision risk one can choose for the ASV subsystem, such that the optimisation task is specific to the CM subsystem (in an *ASV-constrained* setting).

Table 1: t-DCF cost function parameters assumed in ASVspoof 2021.

Evaluation condition	Priors			Costs		
	π_{tar}	π_{non}	π_{spoof}	C_{miss}	C_{fa}	$C_{\text{fa,spoof}}$
LA task PA task	0.9405	0.0095	0.0500	1	10	10
DF task	Equal error rate (EER)					

Note that, while these parameters are the same as for ASVspoof 2019, the ASV system and the conditions of the audio data are different. The above priors and costs, as well as ASV performance outline the ASV-constrained t-DCF computation. In consequence, the ASV-constrained t-DCF is computed according to three terms (C_0, C_1, C_2 – see Appendix) in addition to the CM error rates. C_0, C_1, C_2 are different to those used for ASVspoof 2019, and will vary between 2021 progress and evaluation sets (see Section 9.3).

7 Common ASV system and baseline countermeasures

7.1 Common ASV system

The common ASV system is based on deep speaker embeddings followed by a PLDA-based scoring back-end. The speaker embedding extractor is trained on VoxCeleb data. The PLDA-based scoring back-end is trained on part of the same dataset and then adapted to each data condition. No ASV scores will be made available during the evaluation. Full details of the ASV system will be released at post-evaluation phase.

7.2 Baseline countermeasures

A number of baseline CM solutions will be made available. They include: a high resolution LFCC-GMM [4], [5] system; a CQCC-GMM [6], [7] system; an LFCC-LCNN [8] system; a RawNet2 [9] system. Source codes for all systems are available from the following git repository:

<https://github.com/asvspoof-challenge/2021>

Pre-trained models can be downloaded using scripts provided in the repository.

8 Evaluation rules

The evaluation rules are common to each of the three tasks. They are as follows.

- Participants commit to using data from the ASVspoof 2019 countermeasures *training* and *development* partitions **only** to construct their countermeasures. The use of *any* other external *speech* dataset is strictly forbidden; this includes, but is not limited to, the use of any other public or in-house corpora, found speech or spoofed speech samples, externally trained models, feature vectors (or other statistical descriptors) extracted from external data, or externally trained speech activity detectors. Use of ASVspoof 2019 *evaluation* data as well as *ASV* data is also forbidden.
- The use of external *non-speech* resources (e.g. noise samples and impulse responses), speech codecs and audio compression tools for training or data augmentation *is* permitted as long as their use is detailed in the system description accompanying any submission.
- Participants are allowed to re-partition the ASVspoof 2019 training and development data as they wish (e.g. by pooling the two to create a larger training set, or by excluding some training files).
- Each test sample must be scored independently; techniques such as domain adaption using the evaluation segments, or any use of evaluation data for score normalisation is strictly forbidden.
- In contrast to ASVspoof 2019, participants may submit scores for only **one** system and we do not require the submission of scores for development data. The challenge will be implemented using CodaLab. Participants will be permitted to update their scores multiple times during the progress phase. During the evaluation phase, only a single submission will be permitted. See Section 9.3 for further details.
- Teams are required to submit a detailed system description relating to their final submission (1 system description per team). A system description template in the form of a high-level annotation will be provided and will serve as an indicator of the level of detail required. Compliance with the template and high-level annotation is mandatory and is in everyone's interest. System descriptions in such a common format will be essential to meaningful post-evaluation analysis and will help everyone to put different approaches into context with one another (to promote open science and a better, common understanding).
- Participants should not make direct, public comparisons to the results or rankings of competing teams. This rule applies also to the disclosure of a team's own *ranking* which is, by definition, a comparison to the results and rankings of competing teams. A team can, however, publish their own *results* in terms of t-DCF or EER and the like. A summary of the challenge results

will also be prepared by the organisers. Any team wishing to retain anonymity should provide an anonymous team identifier with their registration. The organisers reserve the right to adjust this policy prior to the scoring platform closing.

The organisers reserve the right to exclude systems from ranking and to exclude teams from future participation in ASVspooF in the event that the above rules are not adhered to.

9 Registration and submission of results

9.1 General mailing list

All participants and team members are encouraged to subscribe to the general mailing list. Subscribe by sending an email to:

sympa@lists.asvspoof.org

with ‘subscribe ASVspooF2021’ as the subject line. Successful subscriptions are confirmed by return email. To post messages to the mailing list itself, emails should be addressed to:

asvspoof2021@lists.asvspoof.org

Subscribers can unsubscribe at any time by following the instructions in the subscription confirmation email.

9.2 Registration and database download

Participants/teams are requested to register for the evaluation. Registration should be performed once only for each participating entity and by accessing the following registration form:

<https://www.spsc-sig.org/form/asvspoof-2021-registration-form>

The organisers aim to validate all registrations within 48 hours. Validated registrations will be confirmed by email.

There are three databases, one for each of the three tasks, all available from the Zenodo website:

Logical Access: <https://zenodo.org/record/4837263>

Physical Access: <https://zenodo.org/record/4834716>

Speech Deepfake: <https://zenodo.org/record/4835108>

Each dataset contains a list of evaluation data and corresponding audio files. All audio files are distributed in FLAC format. All audio data is sampled at a rate of 16 kHz and stored in 16-bit.

9.3 Submission of results

The ASVspooF 2021 challenge will be hosted on CodaLab:

Logical Access: <https://competitions.codalab.org/competitions/32343>

Physical Access: <https://competitions.codalab.org/competitions/32347>

Speech Deepfake: <https://competitions.codalab.org/competitions/32345>

In order to access each site, participants will be required to create an account on the CodaLab website. This will enable you to request access to each submission site. Teams should make only one request per task and requests should normally be made in the name of the team’s registered contact person. Requests will be validated manually by ASVspoof 2021 organisers. The leaderboard is anonymous; usernames will not be disclosed to other participants.

The challenge will run in 2 phases, a *progress* phase and the main *evaluation* phase. During the progress phase, each team may make up to 3 submissions per day. Results determined from a subset of trials will be made available for each submission. During the evaluation phase, participants will be allowed only a single submission for which results will be determined from the remaining evaluation trials. The schedule is shown in Section 11.

Results for 4 reproducible baselines (LFCC-GMM, CQCC-GMM, LFCC-LCNN, RawNet2) are displayed on the Codalab site for each task.

9.4 System descriptions

Participants must submit a system description for each task they participate via the following web forms:

Logical Access: <https://www.spsc-sig.org/form/asvspoof2021-la-systems>

Physical Access: <https://www.spsc-sig.org/form/asvspoof2021-pa-systems>

Speech Deepfake: <https://www.spsc-sig.org/form/asvspoof2021-df-systems>

10 ASVspoof 2021 Workshop

The [ASVspoof 2021 Workshop](#) (an official [Interspeech 2021 satellite event](#)) is planned for September 16, 2021. Due to uncertainty related to the global pandemic, the workshop will be organised fully online (details TBA). Participation will be free of charge and open to anyone (i.e. not limited to challenge participants). Challenge participants are encouraged to submit papers which describe their systems and results. Besides challenge papers, we welcome research contributions related (but not limited) to speaker recognition, anti-spoofing, text-to-speech, voice conversion and speech deepfakes. Submitted papers will be subject to blind peer review. Accepted papers will be published in the ISCA archive¹ and a DOI will be assigned to each paper. While the usual requirements regarding novelty are more relaxed for challenge system papers, reviewing criteria are independent to challenge results and ranking. The usual requirements will apply to regular research contributions.

11 Schedule

A tentative schedule is as follows:

Challenge

- | | |
|---------------------------------|----------------|
| • Initial release of eval plan: | April 19, 2021 |
| • Registration opens: | April 19, 2021 |
| • Evaluation data available: | May 28, 2021 |

¹<https://www.isca-speech.org/iscaweb/index.php/archive/online-archive>

- Registration deadline: June 23, 2021
- Scoring platform closes: July 14, 2021 (AoE)
- System descriptions: July 16, 2021

Workshop

- Paper submission deadline: July 16, 2021
- Notifications: August 16, 2021
- Camera-ready papers: August 30, 2021
- ASVspooof 2021 Workshop (virtual): September 16, 2021

12 Glossary

Generally, the terminologies of automatic speaker verification are consistent with that in the NIST speaker recognition evaluation. Terminologies more specific to spoofing and countermeasure assessment are listed as follows:

Spoofing attack: An adversary, also named impostor, attempts to deceive an automatic speaker verification system by impersonating another enrolled user in order to manipulate speaker verification results.

Anti-Spoofing: Also known as countermeasure. It is a technique to countering spoofing attacks to secure automatic speaker verification.

Bona fide trial: A trial in which the speech signal is recorded from a live human being without any modification.

Spoof trial: In the case of the physical access, a spoofing trial means a trial in which an authentic human speech signal is first played back through a digital-to-analog conversion process and then re-recorded again through analog-to-digital channel; an example would be using smartphone *A* to replay an authentic target speaker recording through the loudspeaker of *A* to the microphone of smartphone *B* that acts as the end-user terminal of an ASV system. In the case of the logical access, a spoofing trial means a trial in which the original, genuine speech signal is modified automatically in order to manipulate ASV.

References

- [1] T. Kinnunen, H. Delgado, N. Evans, K. A. Lee, V. Vestman, A. Nautsch, M. Todisco, X. Wang, M. Sahidullah, J. Yamagishi, and D. A. Reynolds, “Tandem assessment of spoofing countermeasures and automatic speaker verification: Fundamentals,” *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 2195–2210, 2020.
- [2] “ISO/IEC FDIS 19795-1:2021. Information Technology – Biometric performance testing and reporting – Part 1: Principles and framework,” International Organization for Standardization, Geneva, Switzerland, Standard, 2021.
- [3] “ISO/IEC 30107. Information Technology – Biometric presentation attack detection,” International Organization for Standardization, Geneva, Switzerland, Standard, 2016.
- [4] M. Sahidullah, T. Kinnunen, and C. Hanilçi, “A comparison of features for synthetic speech detection,” in *Proc. Interspeech*, 2015, pp. 2087–2091.
- [5] H. Tak, J. Patino, A. Nautsch, N. Evans, and M. Todisco, “Spoofing Attack Detection Using the Non-Linear Fusion of Sub-Band Classifiers,” in *Proc. Interspeech*, 2020, pp. 1106–1110.
- [6] M. Todisco, H. Delgado, and N. Evans, “A new feature for automatic speaker verification anti-spoofing: Constant Q cepstral coefficients,” in *Proc. Odyssey*, 2016, pp. 283–290.
- [7] —, “Constant Q cepstral coefficients: A spoofing countermeasure for automatic speaker verification,” *Computer Speech & Language*, vol. 45, pp. 516–535, 2017.
- [8] X. Wang and J. Yamagishi, “A comparative study on recent neural spoofing countermeasures for synthetic speech detection,” 2021. arXiv: [2103.11326](https://arxiv.org/abs/2103.11326) [eess.AS].
- [9] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, “End-to-end anti-spoofing with RawNet2,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6369–6373.
- [10] NIST, *NIST 2020 CTS Speaker Recognition Challenge Evaluation Plan*, 2020.
- [11] T. Kinnunen, K. Lee, H. Delgado, N. Evans, M. Todisco, M. Sahidullah, J. Yamagishi, and D. A. Reynolds, “t-DCF: A detection cost function for the tandem assessment of spoofing countermeasures and automatic speaker verification,” in *Proc. Odyssey*, Les Sables d’Olonne, France, Jun. 2018.

Appendix: evaluation metric details

The t-DCF [1] is based on statistical detection theory and involves specification of an envisioned *application*. A key feature of t-DCF is the assessment of a tandem system while keeping the two subsystems (CM and ASV) isolated from each other; the participant will develop standalone CM and submit their detection scores that the organizers will combine with the ASV scores. Thus, the participant does not need to develop, optimize or execute any ASV system at any stage. To enable rapid kick-off, the organizers provide t-DCF scoring code and provide pointers to existing CM implementations. Besides t-DCF, the organizers will evaluate equal error rate (EER) of each submitted countermeasure as a secondary metric. This section details the evaluation metrics.

The ASVspoof 2021 challenge is to develop a two-class *bona fide* vs. *spoof* classifier. The positive (bona fide) class corresponds to speech segments of natural and genuine persons (combination of target and non-target classes) and the negative (spoof) class corresponds to attacks. The challenge participants are provided labeled audio files for development purposes (same as in 2019) while the evaluation set samples will be unlabeled.

Both the t-DCF and EER metrics are computed from a set of *detection scores*, one score corresponding to a single audio file (evaluation trial). Participants should assign a real-valued, finite numerical value to each trial which reflects the support for two competing hypotheses, namely that the trial is bona fide or spoofed audio. Similar to the former two edition of ASVspoof, **high detection score should indicate bona fide and low score should indicate spoofing attack**. The scores might be, for instance, log-likelihood ratios (LLRs) between bona fide vs. spoof hypothesis, bonafide posterior probabilities or inner product scores, to give examples.

The primary metric of the LA and PA tasks is based on *ASV-constrained* t-DCF [1], defined as

$$\text{t-DCF}(\tau_{\text{cm}}) = C_0 + C_1 P_{\text{miss}}^{\text{cm}}(\tau_{\text{cm}}) + C_2 P_{\text{fa}}^{\text{cm}}(\tau_{\text{cm}}), \quad (1)$$

where $P_{\text{miss}}^{\text{cm}}(\tau_{\text{cm}})$ and $P_{\text{fa}}^{\text{cm}}(\tau_{\text{cm}})$ are, respectively, the empirical miss and false alarm rates of the CM at threshold τ_{cm} , computed by

$$\begin{aligned} P_{\text{miss}}^{\text{cm}}(\tau_{\text{cm}}) &= \frac{\#\{\text{bona fide trials with CM score} \leq \tau_{\text{cm}}\}}{\#\{\text{Total bona fide trials}\}} \\ P_{\text{fa}}^{\text{cm}}(\tau_{\text{cm}}) &= \frac{\#\{\text{spoof trials with CM score} > \tau_{\text{cm}}\}}{\#\{\text{Total spoof trials}\}}, \end{aligned}$$

and where C_0 , C_1 , and C_2 depend both on the t-DCF parameters and the ASV error rates. They are given by:

$$\begin{aligned} C_0 &= \pi_{\text{tar}} C_{\text{miss}} P_{\text{miss}}^{\text{asv}} + \pi_{\text{non}} C_{\text{fa}} P_{\text{fa}}^{\text{asv}} \\ C_1 &= \pi_{\text{tar}} C_{\text{miss}} - (\pi_{\text{tar}} C_{\text{miss}} P_{\text{miss}}^{\text{asv}} + \pi_{\text{non}} C_{\text{fa}} P_{\text{fa}}^{\text{asv}}) \\ C_2 &= \pi_{\text{spoof}} C_{\text{fa,spoof}} P_{\text{fa,spoof}}^{\text{asv}}, \end{aligned} \quad (2)$$

The meaning of the terms in Eq. (2) are

- $\pi_{\text{tar}}, \pi_{\text{non}}, \pi_{\text{spoof}}$ — **a priori probabilities** of target user, non-target user, and a spoofing attack, respectively (nonnegative and sum = 1).
- $C_{\text{miss}}, C_{\text{fa}}, C_{\text{fa,spoof}}$ — nonnegative **detection costs** assigned to missed target, falsely accepted nontarget and falsely accepted spoofing attack, respectively.

- $P_{\text{miss}}^{\text{asv}}, P_{\text{fa}}^{\text{asv}}, P_{\text{fa,spoofer}}^{\text{asv}}$ — **ASV system error rates** at a pre-specified ASV threshold — miss rate of targets, false accept rate of nontargets, and false accept rate of spoofing attacks, respectively.

The priors/costs shown in Table 1 do not to reflect some real-world case, which is subjective to and thus different for each participant, but—to reflect security-centred settings in which submitted systems are better to distinguish from one another.²

The raw t-DCF value can be difficult to interpret. Following the standard convention adopted in the NIST speaker recognition evaluation campaigns, it is useful to normalize the cost with that of a non-informative detector. In the realm of (1) this refers to a spoofing countermeasure which either accepts or rejects every input, whichever of these two yields lower t-DCF. A normalised version of the ASV-constrained t-DCF is therefore defined by

$$\text{t-DCF}(\tau_{\text{cm}})' = \frac{\text{t-DCF}(\tau_{\text{cm}})}{C_0 + \min\{C_1, C_2\}}, \quad (3)$$

where the denominator is obtained by setting either $P_{\text{miss}}^{\text{cm}}(\tau_{\text{cm}})$ or $P_{\text{fa}}^{\text{cm}}(\tau_{\text{cm}})$ to 1 (in which case the other error rate is 0) and taking the minimum. Finally, as in 2019, we consider the ideally-calibrated countermeasure by adjusting the threshold τ_{cm} to the value whichever gives lowest $\text{t-DCF}(\tau_{\text{cm}})'$. That is, the primary metric is

$$\text{t-DCF}'_{\min} = \min_{\tau_{\text{cm}}} \text{t-DCF}(\tau_{\text{cm}})'. \quad (4)$$

This metric ranges in between $[\text{ASV floor}, 1]$, where the value of 1 corresponds to a useless countermeasure (meaning the system would have worked better without it) while the lower bound

$$\text{Normalized ASV floor} \equiv \frac{C_0}{C_0 + \min\{C_1, C_2\}} \quad (5)$$

corresponds to a perfect (error-free) CM. We call it *ASV floor* to indicate that the performance is floored by the ASV errors.

The reader familiar with the DCF evaluation framework used in NIST speaker recognition evaluations (SREs) recognizes immediately the similarities³. Nonetheless, unlike in the ASV tasks in the NIST SRE campaigns where the weightage of the miss and false alarms is *known a priori*, note that the coefficients C_0 , C_1 and C_2 depend not only on prior and cost terms but *also* on the errors of the ASV system. The reader who participated ASVspoofer 2019 recognizes that in ASVspoofer 2021 we retain the term C_0 in the metric (it was discarded in 2019). Note also that the cost formulation itself has been simplified from the original work [11]; the reader is pointed to [1] for further explanations.

As in 2019, the challenge participant cannot interact with the ASV system and will have limited information about it. They *will* be provided labeled ASV scores of the development set to demonstrate t-DCF computation, but will *not* be provided any evaluation set ASV scores. The organizer’s ASV system will be scored exactly on the same way on development and evaluation data but the evaluation set error rates may be different from those of the development set. What is common to all the ASV systems is that their detection threshold will be set at the equal error (EER) point,

²To reflect priors and costs of, e.g., a specific business case for, particular services and products requires owners of these to quantify their own beliefs. This would be necessary in any case, even when CMs are compared by means of EER. When referring to EER, risk assessment is implied indirectly (Which EER is good? Why even equal-errors?): it is outsourced to owners of products and services by making lots of hidden assumptions. Here, we are fully transparent on our assumptions and provide formalism and reference implementations to Bayes risk reporting. When choosing more general/specific use cases for operation, priors/costs need to be adapted accordingly; expectations are contextual.

³In fact, the NIST DCF [10] is a special case of the tandem t-DCF that assume zero spoof prior and which uses ‘accept all’ countermeasure

at which $P_{\text{miss}}^{\text{asv}} = P_{\text{fa}}^{\text{asv}}$ (see [1, Section V] for further discussion). Further information on ASV performance in one eval-set will be made available after the challenge.

Similar to the past two editions of ASVspoof, performance will be computed from detection scores pooled across all trials within a given challenge task (LA, PA or DF). That is, within one of these two tasks, all the bona fide trials in the evaluation set will be used to obtain $P_{\text{miss}}^{\text{cm}}(\tau_{\text{cm}})$ and all the spoof trials will be used to obtain $P_{\text{fa}}^{\text{cm}}(\tau_{\text{cm}})$ —ASV system error rates are determined by the challenge organisers. This strategy is intended to encourage development of countermeasures that calibrate well across different variants of attacks and conditions. That is, to develop countermeasures whose scores are mutually compatible across a wide range of conditions. The different challenge tasks, however, will be scored independent of each other. The three tasks should be considered independent tasks and a participant interested in all of them is welcome to use even completely different technologies to address each task.

The t-DCF parameters used in ASVspoof 2019 are provided in Table 1. The selected values, shared across the LA and PA tasks, are arbitrary but intended to be representative of *authentication* applications. For instance, a banking domain may involve a very large number of authentication requests, most of which represent genuine target users; therefore, we have asserted high target prior. Priors of both ‘undesired’ trials (nontargets and spoofing attacks) are set to low values for the same reason. The costs, in turn, reflect the idea that false acceptances are more detrimental than false rejections. Note that the priors do not necessarily reflect the empirical proportions of different classes in the evaluation data.

Secondary metric: equal error rate (EER): The secondary evaluation metric is *equal error rate* (EER) which was used in the past two editions of ASVspoof. EER corresponds to CM threshold τ_{EER} at which the miss and false alarm rates equal each other⁴, i.e. $\text{EER} = P_{\text{fa}}^{\text{cm}}(\tau_{\text{EER}}) = P_{\text{miss}}^{\text{cm}}(\tau_{\text{EER}})$. Similar to t-DCF, EER will be computed from pooled evaluation trials within the given task.

⁴One may not find such threshold exactly as $P_{\text{fa}}(\tau_{\text{EER}})$ and $P_{\text{miss}}(\tau_{\text{EER}})$ change in discrete steps. You may use $\theta_{\text{EER}} = \arg \min_{\theta} |P_{\text{fa}}^{\text{cm}}(\tau_{\text{cm}}) - P_{\text{miss}}^{\text{cm}}(\tau_{\text{cm}})|$ or more advanced methods such as EER based on convex hull (ROCCH-EER) implemented in the Bosaris toolkit, <https://sites.google.com/site/bosaristoolkit/>. In ASVspoof2019, we use a simplified nearest-neighbor type computation demonstrated in the scoring package.