



a-DCF: an architecture agnostic metric with application to spoofing-robust speaker verification

Hye-jin Shim,^{1,*} Jee-weon Jung,^{1,*} Tomi Kinnunen,² Nicholas Evans,³
Jean-Francois Bonastre^{4,5} and Itshak Lapidot^{6,5,*†}

¹Carnegie Mellon University, USA – ²University of Eastern Finland, Finland

³EURECOM, France – ⁴Inria, France

⁵Avignon University, France – ⁶Afeka Tel-Aviv Academic College of Engineering, Israel

shimhz6.6@gmail.com, jeeveonj@ieee.org, tomi.kinnunen@uef.fi,

evans@eurecom.fr, jean-francois.bonastre@inria.fr, itshakl@afeka.ac.il

Abstract

Spoofing detection is today a mainstream research topic. Standard metrics can be applied to evaluate the performance of isolated spoofing detection solutions and others have been proposed to support their evaluation when they are combined with speaker detection. These either have well-known deficiencies or restrict the architectural approach to combine speaker and spoof detectors. In this paper, we propose an architecture-agnostic detection cost function (a-DCF). A generalisation of the original DCF used widely for the assessment of automatic speaker verification (ASV), the a-DCF is designed for the evaluation of spoofing-robust ASV. Like the DCF, the a-DCF reflects the cost of decisions in a Bayes risk sense, with explicitly defined class priors and detection cost model. We demonstrate the merit of the a-DCF through the benchmarking evaluation of architecturally-heterogeneous spoofing-robust ASV solutions.

1. Introduction

All biometric verification systems, including automatic speaker verification (ASV), have the single task of determining reliably whether or not a biometric sample corresponds to a claimed identity [1]. Early on, reliability was interpreted as classifiers which could discriminate between target trials and (zero-effort)¹ non-target trials [2]. In more recent years, and as a result of advances in text-to-speech synthesis and voice conversion, the threat of spoofing attacks, also known as presentation attacks [3], has come to the fore.

Despite the still-evolving consideration and study of spoofing, the single task of discriminating between target and non-target trials remains fundamentally unchanged; target trials should be accepted, while anything else should be rejected. The approach to address this challenging problem has nonetheless undergone rather more fundamental shifts. While there are examples of alternative approaches even in the early literature [4–7], almost all studies of spoofing-robust ASV adopt the use of so-called *tandem* architectures [8, 9]. These employ a pair of sub-systems, each also a binary classifier, one tasked with discriminating between target and non-target trials (the speaker detector), the other between bonafide and spoofed trials (the spoof detector). The tandem approach is

characteristic of the majority of related work, including studies involving other biometric traits [10, 11].

Standard metrics developed for the evaluation of speaker detectors can also be applied to the evaluation of spoof detectors, also known as countermeasures (CMs); they are both binary classifiers. Alternative metrics proposed in recent years also support the evaluation of speaker and spoof detectors when combined [12, 13]. While the combination of speaker and spoof detectors still constitutes a single, binary classifier with the very same original task of accepting bonafide target trials and rejecting anything else, the consideration of spoofing complicates evaluation. Despite still being a binary classifier, there are now three input classes (*target*, *non-target*, and *spoof*). The tandem detection cost function (t-DCF) [12] was hence developed to accommodate the evaluation of spoofing-robust ASV.

The t-DCF, however, is somewhat restrictive in terms of supported architectures. The speaker detector is used as a gate to the spoof detector,² while an AND decision logic is used to combine their respective classification decisions; given a target trial, the spoof detector should indicate that an input utterance is bonafide AND the speaker detector should indicate that the input utterance corresponds to the claimed identity. The t-DCF requires the computation of separate speaker and spoof detection scores and cannot be applied to the evaluation of alternative approaches which, for example, might produce only a single score. Such approaches have been reported in the past and continue to emerge [6, 14–19]. Their incompatibility with the t-DCF metric, has hence stimulated the development of alternative metrics, all forms of equal error rate (EER) estimates [14]. Different metrics are hence employed for the evaluation of different architectural approaches to the evaluation of spoofing-robust ASV solutions, despite them all sharing exactly the same goal. Use of different metrics then complicates the benchmarking of competing solutions. Furthermore, while an evaluation metric introduced in [20] can be adapted to incorporate the aforementioned condition, no existing research explicitly defines this problem to date.

We aim to provide a solution to this problem. We propose a derivative of the original DCF metric which is agnostic to the chosen architecture, be it a tandem system comprising separate, cascaded speaker and spoof detection sub-systems, alternatives whereby the roles of speaker and spoof detection are more closely integrated and perhaps jointly optimised, or indeed any other potential architecture. The only demand of the architecture-agnostic detection cost function (a-DCF) is that, whatever the architecture might be, it must produce a single score which provides an indication of whether or not an input

*Equal contribution. Similar ideas were devised independently from two different groups and later combined.

† Corresponding author

Code is available at https://github.com/shimhz/a_dcf.

¹*Zero-effort* implies purely casual impostors which make no concerted effort to deceive the system.

²There are plenty of other possible combination architectures.

utterance corresponds to the claimed identity and is also bonafide.

2. Related work

Spoof detectors were initially evaluated independently from speaker detectors, i.e. without consideration of impacts upon ASV behaviour. There were a number of advantages to this approach. First, as a fledgling initiative, it seemed sensible to simplify the task so that researchers without expertise in ASV, but perhaps with other relevant expertise, could still participate in the development of spoof detection solutions. Second, at the time the best architecture with which to combine speaker and spoof detectors was unclear. Third, without one having been identified, it seemed sensible to avoid the imposition of a particular combination architecture only for the sake of a common approach to evaluation. Spoof detection performance was then estimated straightforwardly using an EER metric, a measure of discrimination between the two classes (bonafide and spoof).

Even if spoof detection solutions were initially evaluated independently, the potential for their interference with speaker detection is evident. Just like speaker detectors, spoof detectors make errors, either by accepting spoofs, or by rejecting bonafide trials; depending on the combination architecture, a speaker detector might not be able to recover from errors made by a spoof detector (or vice versa). The community was hence in need of a metric which could be used to evaluate the performance of *combined* speaker and spoof detection systems.

The tandem-detection cost function (t-DCF) [12] was developed for this purpose. Again, there were several advantages, not least because it provided an alternative to the EER metric. As a result of the assumed conditional independence of speaker and spoof detection decisions, the t-DCF also promotes the independent development of spoof detectors.³ Decisions produced by a spoof detector can be combined with those made by a speaker detector to provide a measure of spoofing-robust ASV performance. At the same time, use of a DCF-derived metric also helped bridge the gap between the relatively young spoof detection community and the mature speaker detection community which had embraced the DCF metric [21] decades earlier.

The well-known NIST-defined DCF [21], given by

$$\text{DCF}(t) := C_{\text{miss}} \pi_{\text{tar}} P_{\text{miss}}^{\text{asv}}(t) + C_{\text{fa}} \pi_{\text{non}} P_{\text{fa}}^{\text{asv}}(t), \quad (1)$$

reflects the cost of decisions in a Bayes risk sense [22,23] for an ASV system in the absence of spoofing attacks. In (1), π_{tar} and $\pi_{\text{non}} = 1 - \pi_{\text{tar}}$ are the class priors. C_{miss} and C_{fa} are non-negative costs assigned to missed targets (false rejections) and non-target false alarms (false acceptance), respectively. $P_{\text{miss}}^{\text{asv}}(t)$ and $P_{\text{fa}}^{\text{asv}}(t)$ are the two ASV detection error rates as a function of the detection threshold, t .

The t-DCF metric can be used for the evaluation of tandem systems when the non-target class is augmented to include the potential for spoofing attacks. The general form of the t-DCF is [12]

$$\begin{aligned} \text{t-DCF}(t_{\text{cm}}, t_{\text{asv}}) := & C_{\text{miss}} \cdot \pi_{\text{tar}} \cdot P_{\text{miss,tar}}^{\text{tdm}}(t_{\text{cm}}, t_{\text{asv}}) \\ & + C_{\text{fa,non}} \cdot \pi_{\text{non}} \cdot P_{\text{fa,non}}^{\text{tdm}}(t_{\text{cm}}, t_{\text{asv}}) \\ & + C_{\text{fa,spf}} \cdot \pi_{\text{spf}} \cdot P_{\text{fa,spoof}}^{\text{tdm}}(t_{\text{cm}}, t_{\text{asv}}), \end{aligned} \quad (2)$$

which now contains a third term related to false alarms stemming from spoofing attacks, and where the three *tandem* (tdm) detection error rates are now functions of *two* detection thresholds — one for the spoof detector or CM sub-system (t_{cm}), one for the speaker

³In principle, the conditional independence assumption of the two sub-systems even enables them to be developed using different datasets, yet they can still be evaluated in combination; in practice, this is not recommended, as it introduces an unnecessary data variation factor.

Table 1: Three EERs used in the SASV 2022 challenge. System is targeted to accept “+” trials and reject “-” trials.

	Target	Non-target	Spoof
SV-EER	+	-	
SPF-EER	+		-
SASV-EER	+	-	-

detector or ASV sub-system (t_{asv}). Each of the tandem detection error rates are computed under the assumption that the speaker and spoof detectors make (class-conditionally) independent decisions. This leads to an AND-gate rule for combining the two sub-system decisions; refer to [12, Section III.C] and [13, Section 4.3] for details.

Computation of the t-DCF in (2) requires two sets of detection scores, each corresponding to one of the two sub-systems. Nonetheless, one can choose to freeze either sub-system (including its operating point) to give a *constrained* t-DCF which then becomes a function of a single detection threshold only. Specifically, the ASV-constrained t-DCF [12, Eq. (10)] is given by

$$\text{t-DCF}(t_{\text{cm}}) := C_0 + C_1 P_{\text{miss}}^{\text{cm}}(t_{\text{cm}}) + C_2 P_{\text{fa}}^{\text{cm}}(t_{\text{cm}}), \quad (3)$$

where $P_{\text{miss}}^{\text{cm}}(t_{\text{cm}})$ and $P_{\text{fa}}^{\text{cm}}(t_{\text{cm}})$ are now the miss and false alarm rates for the spoof detector. The parameters C_0 , C_1 and C_2 [12, Eq. (11)] depend on pre-defined cost and prior parameters, in addition to the performance of a (frozen) ASV sub-system (speaker detector). The ASV-constrained t-DCF remains the cost of a complete (tandem) system — but where the machine learning engineer’s optimization efforts are constrained to modifying the spoof detector only, since both the ASV sub-system and the combination rule are already ‘written in stone’. For the ASVspoof challenges [24], this kind of *partial* optimisation strategy was adopted for the reasons noted above; the ASV system and t_{asv} in (2) were set by the organisers, while challenge participants could focus on improving the spoof detector, while having partial knowledge of speaker detector implementation.

Use of the t-DCF, whether (2) or (3), is nowadays unnecessarily restrictive. It assumes a specific, non-customisable combination architecture whereby speaker and spoof detectors are used in cascade, and whereby decisions produced by each are combined with an AND decision logic. The constrained t-DCF (3) further restricts the class of spoofing-robust ASV models. Many other architectural approaches are also possible, for instance score and embedding level combination architectures — or even systems comprising multiple spoof detection sub-systems. None of these systems can be evaluated using the t-DCF.

An ensemble of three different metrics, all EER estimates, was developed subsequently to support the evaluation of more flexible spoofing-robust ASV architectures [14]. Each one corresponds to an evaluation protocol comprising a different mix of trial types illustrated in Table 1. The traditional speaker verification EER (SV-EER) is used to measure the discrimination between targets and non-targets. The spoofing EER (SPF-EER) measures the discrimination between bonafide targets and spoofs. Finally, the spoofing-aware speaker verification EER (SASV-EER) measures discrimination between bonafide targets (which should be accepted) and everything else (bonafide non-targets and spoofing attacks, which should be rejected).

The ensemble of metrics provides insights into classifier performance under different trial combinations, as well as the impacts of spoofing and spoof detection upon ASV (just like the t-DCF). Even so, there are *three* metrics (instead of one), and they are all EER estimates (instead of DCF estimates). The shortcomings of *any* EER-based metric are well-known [3, 13, 25]. For ASV applications,

one generally prioritises *either* a low miss rate *or* a low false alarm rate, which conflicts with the notion of *equal* with the EER. Unlike DCF-based metrics, the EER is not customizable or an optimization target for different applications. The SASV-EER metric which pools non-targets and spoofing attacks among the negative class has even *further* issues. The classic EER, when used to measure discrimination between any pair of *two* classes is not dependent on the empirical class priors. As soon as the third class (spoofing attacks) is pooled with non-targets, however, this property no longer holds [13, Section 4.4]; the SASV-EER metric becomes a function of the empirical class priors. Hence, the metric itself depends on the evaluation data.

Note the subtle, but critically important, difference in the role of the class priors for the DCF, t-DCF and the proposed a-DCF. For all these metrics, the class priors are *not* a property of any dataset; they are *not* estimated or computed from empirical data, but instead reflect the evaluator’s assumed (uncertain) class priors for the given application. Different to the *implicit* (hidden) weighting of the non-target and spoof false alarm rates for the SASV-EER, DCF-based metrics make the class priors (and consequences of classification errors) *explicit*. For these reasons, **we strongly discourage any further use of the SASV-EER metric**. If one absolutely must report EERs, our recommendation is to limit such reporting to the SV-EER and SPF-EER. For the interested reader, we point to the recently-introduced *concurrent tandem* EER (t-EER) [13] metric which is not dependent on the empirical class priors. Nonetheless, similar to t-DCF, it can only be used to evaluate cascaded systems; and similar to the other EERs, it cannot be customised to different applications.

It is time to reconsider the architectures for spoofing-robust ASV. New metrics are needed for their comparative evaluation. They should avoid the use of EER-based metrics, and especially use of the SASV-EER, and inherit the favourable properties of the DCF and t-DCF metrics with an explicit detection cost model. Last, they should be agnostic to the classifier architecture.

3. Architecture-agnostic DCF

In this section we outline the theoretical basis for the proposed architecture-agnostic detection cost function (a-DCF). We start with a general formulation (Section 3.1), and show how it relates to the NIST-defined DCF in (1), before presenting the specific formulation adopted for experiments reported later in the paper (Section 3.3). We then present the a-DCF as a generalisation of the traditional two-class DCF (Section 3.4) before arguing how the a-DCF can be applied to the evaluation of other spoofing-robust ASV architectures.

3.1. General form

Let us assume a multi-class classification problem. Let $\mathcal{A} = \{a_1, a_2, \dots, a_K\}$ be a set of K ground-truth class labels, let $T \in \mathcal{A}$ a true class label for a given trial and $E \in \mathcal{A}$ is an estimated/predicted class label (classifier output). Let us define also the following:

- $K \times 1$ column vector of class priors $\pi_k = \Pr(T = a_k)$:

$$\Pi = \begin{bmatrix} \pi_1 & \pi_2 & \dots & \pi_K \end{bmatrix}^\top,$$

such that, $\forall k, \pi_k \geq 0, \sum_{k=1}^K \pi_k = 1$.

- $K \times K$ matrix of classifier conditional probabilities:

$$\mathcal{P} = [P_{qk} = \Pr(E = a_q | T = a_k)] \\ \forall q, k \in \{1, \dots, K\}, P_{qk} \geq 0, \quad \forall k, \sum_{q=1}^K P_{qk} = 1,$$

where P_{qk} is the conditional probability that the classifier outputs decision a_q given the ground-truth class a_k .

- $K \times K$ matrix of conditional costs:

$$\mathcal{C} = [c_{qk}] \\ \forall q, k \in \{1, \dots, K\}, c_{qk} \geq 0, \quad \forall k, c_{kk} = 0,$$

where c_{qk} is the cost of the classifier outputting decision a_q given the ground-truth class a_k . Without loss of generality, costs in \mathcal{C} can be set to zero in the case of correct decisions ($c_{kk} = 0$), whereas incorrect decisions can be assigned non-negative, real values.

The *total* cost of making decisions can then be expressed in compact form by:

$$C_T = \mathbf{1}_{1 \times K} \cdot (\mathcal{C} \circ \mathcal{P}) \cdot \Pi, \quad (4)$$

where $\mathbf{1}_{1 \times K}$ is a vector of ones which acts to sum the columns of the matrix resulting from the terms to the right, and where \circ is the Hadamard or entry-wise product. Though expressed using different formalism, (4) coincides with [26, Eq. (4)].

In practice, classifier conditional probabilities are defined for some operating point t for which entries in the matrix of classifier conditional probabilities, now $\mathcal{P}(t)$, can be approximated by $p_{qk}(t) \approx N_{qk}(t)/N_k$. N_k is the number of trials belonging to class a_k whereas $N_{qk}(t)$ is the number of trials among them that are classified as belonging to class a_q .

3.2. Spoofing-robust speaker verification

For a standard ASV task, classifier decisions result in the labeling of an input trial as either a target or non-target, corresponding respectively to either accept or reject decisions. In this case, there are two possible input classes and two possible classifier predictions ($K = 2$). The formulation in (4) can then written as:

$$C_T(t) = C_{\text{non,tar}} \pi_{\text{tar}} P_{\text{non,tar}}(t) + C_{\text{tar,non}} \pi_{\text{non}} P_{\text{tar,non}}(t), \quad (5)$$

where $c_{1,2} = C_{\text{non,tar}}$ and $c_{2,1} = C_{\text{tar,non}}$ are the costs of classifying targets as non-targets and non-targets as targets respectively, π_{tar} and $\pi_{\text{non}} = 1 - \pi_{\text{tar}}$ are the class priors, and where $p_{\text{non,tar}}(t)$ and $p_{\text{tar,non}}(t)$ are the classifier conditional probabilities of each decision error at the set classifier threshold t . (5) is identical to the familiar detection cost function (DCF) in (1) in which ‘miss’ signifies targets mistaken for non-targets and ‘fa’ (false alarm) signifies non-targets mistaken for targets.

When subjected to spoofing attacks, the negative input class becomes a union of bonafide non-target (zero-effort impostor) trials and spoofed target trials. For brevity, we refer to them simply as non-target and spoofed trials respectively. Different priors and costs can be assigned to each. Even so, the decision is still binary, with inputs still being labeled as either targets (positive class) or non-targets (negative class). Identical to the traditional ASV scenario (without spoofing attacks), there hence remains a single cost for missed target detections. However, since there are two different negative class priors, one for non-targets and one for spoofs,⁴ there are likewise two (possibly different)⁵ false alarm costs. Classifier conditional probabilities can also be estimated in the usual way for a specified operating point or threshold t .

⁴Note that positive and negative class assignments are arbitrary [12]. In keeping with the literature, here we assume the assignment of bonafide target trials to the positive class and everything else to the negative class.

⁵From an applications perspective, the mistaking of non-targets for targets when caused by an unintentional impostor may not have the same ramifications (cost) of when a spoof mistaken for a target is caused by a fraudster. In the spirit of ‘an error is an error, no matter the cause’, for all experiments reported later in the paper, we use the *same* costs for each type of false alarm.

Table 2: Comparison on the characteristics of existing EER and DCF metrics used in the ASV and CM field. SASV refers to spoofing-robust ASV where ASV and CM sub-systems are combined to produce a single score output. In the first row, ‘Explicit detection cost’, ‘Independent of class prior’ refers to whether the detection cost is explicitly used or not (difference between DCF and EER families) and whether it is affected by the number of different trial classes. ‘# Supported classes’ and ‘Scores required’ indicate the number of class types considered and input score(s) needed for the calculation, respectively.

	Explicit detection cost	Independent of class prior	# Supported classes	# scores required
DCF [21]	✓	✓	2	1
t-DCF [12]	✓	✓	3	2 (ASV and CM)
a-DCF (proposed)	✓	✓	3	1 (SASV)
EER	✗	✓	2	1 (ASV or CM)
SASV-EER [14]	✗	✗	3	1 (SASV)
t-EER [13]	✗	✓	3	2 (ASV and CM)

The priors, the matrix of classifier conditional probabilities and the matrix of conditional costs then have the form:

$$\Pi = \begin{bmatrix} \pi_{\text{tar}} \\ \pi_{\text{non}} \\ \pi_{\text{spf}} \end{bmatrix}$$

$$\mathcal{P}(t) = \begin{bmatrix} P_{\text{tar,tar}}(t) & P_{\text{tar,non}}(t) & P_{\text{tar,spf}}(t) \\ P_{\text{non,tar}}(t) & P_{\text{non,non}}(t) & P_{\text{non,spf}}(t) \\ P_{\text{spf,tar}}(t) & P_{\text{spf,non}}(t) & P_{\text{spf,spf}}(t) \end{bmatrix}$$

$$\mathcal{C} = \begin{bmatrix} 0 & C_{\text{tar,non}} & C_{\text{tar,spf}} \\ C_{\text{non,tar}} & 0 & 0 \\ C_{\text{spf,tar}} & 0 & 0 \end{bmatrix}$$

The off-diagonal zero elements in the cost matrix correspond to the cost of mistaking non-target trials for spoofed trials and vice versa; the binary classifier distinguishes target trials from anything else and hence cannot distinguish between non-targets and spoofs. The total cost is then given again from (4) by:

$$C_T(t) := C_{\text{non,tar}}\pi_{\text{tar}}P_{\text{non,tar}}(t) + C_{\text{spf,tar}}\pi_{\text{tar}}P_{\text{spf,tar}}(t) + C_{\text{tar,non}}\pi_{\text{non}}P_{\text{tar,non}}(t) + C_{\text{tar,spf}}\pi_{\text{spf}}P_{\text{tar,spf}}(t). \quad (6)$$

Remembering that the negative class is the union of non-target and spoofed trials, and like the correspondence between (5) and (1), we obtain the architecture-agnostic detection cost function (a-DCF):

$$a\text{-DCF}(t) := C_{\text{miss}}\pi_{\text{tar}}P_{\text{miss}}(t) + C_{\text{fa,non}}\pi_{\text{non}}P_{\text{fa,non}}(t) + C_{\text{fa,spf}}\pi_{\text{spf}}P_{\text{fa,spf}}(t), \quad (7)$$

in which the first two lines of (6) are reduced to a single term which encapsulates the miss-classification of target trials as non-targets. C_{miss} , $C_{\text{fa,non}}$ and $C_{\text{fa,spf}}$ are, respectively, the costs of missing (falsely rejecting) a target speaker, falsely accepting a non-target speaker, and falsely accepting a spoof. π_{tar} , π_{non} and π_{spf} are the asserted class priors.⁶ Finally, $P_{\text{miss}}(t)$ (the miss rate), $P_{\text{fa,non}}(t)$ (the non-target false alarm rate) and $P_{\text{fa,spf}}(t)$ (the spoof false alarm rate) are the respective empirical detection error rates at detection threshold t .⁷

⁶As a probability mass function, $\pi_{\text{tar}} + \pi_{\text{non}} + \pi_{\text{spf}} = 1$ and therefore any two of these priors are sufficient to characterize the class priors.

⁷While, as argued above, costs $C_{\text{fa,non}}$ and $C_{\text{fa,spf}}$ can be set to different or identical values, empirical detection error rates $P_{\text{fa,non}}(t)$ and $P_{\text{fa,spf}}(t)$ are almost certainly different, with the latter normally being greater than the former.

3.3. Normalisation

Identical to the t-DCF, the value of a-DCF is not bounded and can be difficult to interpret. Therefore, similar to normalisation of the NIST DCF [21] as well as the t-DCF [12], we further scale the a-DCF using the cost of a *default system* [27, 28] which is configured to either accept ($t \rightarrow -\infty$) or reject ($t \rightarrow +\infty$) every trial:

$$a\text{-DCF}_{\text{def}} := \min \left\{ C_{\text{miss,tar}}\pi_{\text{tar}}, C_{\text{fa,non}}\pi_{\text{non}} + C_{\text{fa,spf}}\pi_{\text{spf}} \right\}$$

where the minimum costs of the two edge cases are considered (these expressions follow directly from the limit behavior of the miss and false alarm rates in (7)). The *normalized* a-DCF is then given by:

$$a\text{-DCF}(t)_{\text{norm}} = \frac{a\text{-DCF}(t)}{a\text{-DCF}_{\text{def}}}. \quad (8)$$

While in real operational settings the detection threshold t must be set *before* observing test data, for analysis purposes it is informative to report the lowest possible cost when t is allowed to vary. The *minimum* a-DCF is defined simply by:

$$a\text{-DCF}_{\text{min}} := \min_{t \in \mathbb{R}} a\text{-DCF}_{\text{norm}}(t). \quad (9)$$

While the normalised a-DCF function in (8) can exceed 1, which implies that it performs worse than a non-informative system which either accepts or rejects every trial, the minimum is always bounded between 0 and 1. All detection costs reported later in the paper are normalised minimums computed according to (9). For simplicity, they are denoted simply as a-DCF.

3.4. Relation to NIST DCF and t-DCF

It would be helpful to discuss the relation between a-DCF and other DCFs reported in [12, 20, 21]. The comparison is shown in Table 2. The original DCF [21] in (1), as endorsed by NIST through the speaker recognition evaluation (SRE) series, is designed for the benchmarking of binary classifiers which produce a single detection score — such as standalone speaker or spoof detection systems. The NIST DCF can be seen as a degenerate case of the a-DCF with $\pi_{\text{spf}} = 0$ (no spoofing attacks). Interestingly, the NIST 2012 SRE campaign [20] used a 3-term DCF similar to (7), though the two FA rate terms are those of *known/unknown nontargets* (as opposed to nontargets and spoofs as in our work). Though the 2012 DCF and a-DCF are in the same function class of evaluation metrics (i.e. differ only by the choice of costs and priors), they are motivated (and used) from entirely different perspectives.

The t-DCF [12] in (2) and (3) is closely related to the a-DCF. The difference in terms of structure is that the t-DCF requires *two*

Table 3: Two a-DCF priors and costs accounting for different scenarios.

	π_{spf}	π_{non}	π_{tar}	C_{miss}	$C_{\text{fa,non}}$	$C_{\text{fa,spf}}$
a-DCF1	0.05	0.01	0.94	1	10	10
a-DCF2	0.01	0.01	0.98	1	10	10

detection scores, as opposed to a single score in the case of the NIST DCF and the a-DCF. Nonetheless, by assigning a peculiar ‘dummy’ countermeasure, we can view the a-DCF as a special case of the t-DCF. To be specific, by letting the threshold in [12, Eq. (7)] approach $-\infty$, the spoof detection miss and false alarm rates approach 0 and 1, respectively. This yields precisely the a-DCF expression in (7). Nonetheless, this is *not* our point, as we discuss next.

3.5. a-DCF: beyond tandem systems

While the above special cases reveal formal connections between cost functions, degenerate spoof prior and dummy spoof detectors are not in the *typical scope* of conventional (spoofing-robust) ASV and tandem recognizers, respectively. The former aims to improve speaker discrimination, without any consideration for spoofing attacks, hence neither the benchmarking datasets, nor the evaluation metrics or even the systems need be concerned with the potential for spoofing attacks. Likewise, the key feature of tandem systems is that they consist of two non-dummy sub-systems, the decisions of which result from the application of *finite* thresholds. A tandem system with either one of the two systems being a dummy system is not in the typical scope.

Consider now a slightly *revised* tandem detection system which still uses separate thresholds for the speaker and spoof detector, but which outputs a single soft decision as long as the spoof detector score is high enough. Namely,

$$s = \begin{cases} s_{\text{asv}}, & \text{if } s_{\text{cm}} \geq t_{\text{cm}} \\ -\infty, & \text{if } s_{\text{cm}} < t_{\text{cm}}, \end{cases} \quad (10)$$

where $s_{\text{asv}}, s_{\text{cm}} \in \mathbb{R}$ are the detection scores produced by the speaker detector (the ASV sub-system) and the spoof detector (the CM sub-system), respectively, and where t_{cm} is a pre-set CM threshold.

4. Experimental setup

While the key contribution of this work is the new a-DCF metric, we report also an example application to the evaluation of spoofing-robust ASV solutions. All experiments were performed using an open source, publicly-available implementation⁸ which can be used to reproduce our results.

4.1. Dataset

We employ the ASVspoof 2019 logical access (LA) corpus [31] for all experiments. The corpus is organised into three subsets: training, development, and evaluation. The training and development sets include genuine and spoofed speech from 20 speakers (8 male, 12 female). We employ the evaluation protocol⁹ used for the SASV 2022 challenge [14]. The protocol was proposed initially for calculating the coefficients of ASV-constrained t-DCF (Eq. 10 in [12]) in the ASVspoof 2019 challenge. It consists of three trial

⁸https://github.com/shimhz/a_dcf

⁹https://github.com/sasv-challenge/SASVC2022_Baseline/blob/main/protocols/ASVspoof2019.LA.asv.eval.gi.trl.txt

types, *target*, *non-target*, and *spoof*, given an enrolment and test utterances. Hence it can be used for evaluating different types of ASV systems, including spoofing-robust ASV systems.

4.2. Systems

We consider three types of systems: cascade, jointly optimised, and single-model, to demonstrate the architecture-agnostic of the a-DCF metric.¹⁰ For the cascade (decision-level tandem) systems, RawNet3 [29] is used as an ASV model. As a CM system, AASIST [30], AASIST-L [30] are used and several cascade systems are composed. In addition, we consider systems submitted to the SASV 2022 challenge [14, 32–34] which explored diverse methodologies including score-level and embedding-level combinations. For the single-model systems, we employ various versions from [35], comprising MFA-Conformer [36] and SKA-TDNN [37] models, where the training data varies throughout the devised multi-stage training.¹¹ To avoid confusion regarding system details and to focus on comparing metric measurements, all systems are denoted simply with the alphanumeric identifiers in Table 5.

5. Results

The discussion below is oriented toward the comparison of different evaluation metrics, highlighting the merits and agnosticity of each, rather than the performance of different systems (which is not the objective). Results are expressed in terms of the four different EERs, three from Table 1,¹² the t-EER [13], the min t-DCF [12], where possible, and the a-DCF.

As illustrated in the second column of Table 4 which shows results only for cascade systems, results are shown for two different configurations in each case where either the speaker detector (an ASV sub-system), or the spoof detector (a CM sub-system) comes first. The detection threshold for the first system is set to either 0.05 or 0.5. t-DCF results are computed in the usual way. The a-DCF results are computed according to (9) in which the roles of the speaker and spoof detector can be interchanged. In both cases, the ASV sub-system is the RawNet3 model described in [29]. The CM sub-system is either the AASIST or AASIST-L model, a lightweight variant, both described in [30]. t-DCF1 and t-DCF2 results are computed with the same priors and costs as a-DCF1 and a-DCF2 respectively (see Table 3). Though different formulations are readily derived, computation of the t-DCF according to (3) is specific to the evaluation of cascade systems comprising a spoof detector and a frozen speaker detector.¹³ Finally, the t-EER [13] is not dependent on the order of the two sub-systems, hence there is only one value for each combination of ASV and CM sub-systems. SASV-EER, SV-EER and SPF-EER results can all be computed for all setups, but they remain EERs and suffer from the detractions described earlier. Use of the a-DCF avoids these issues with results also being computable for each setup, with a common formulation

¹⁰Note that although we categorise systems in Table 5 into three categories, it does not mean that all systems should fall into one of the three categories. ‘Jointly optimised’ can be defined differently under diverse circumstances, where in our case, we use the term to call systems that jointly optimise pre-trained ASV and CM systems to derive the final score for each input.

¹¹S1, S2, S3, and S4 represent MFA-Conformer index 7 and SKA-TDNN indices 3, 7, and 11, respectively, from the model indices shown in https://github.com/sasv-challenge/SASV2_Baseline.

¹²For the sake of completeness, we include the SASV-EER despite our previously stated reservations concerning its use.

¹³The unconstrained t-DCF formulation in (2) is more flexible and can be applied to the evaluation of other architectures, such as systems where the CM (rather than the ASV) subsystem is frozen. Nonetheless, this leads to different cost functions that are not comparable with (3), hence the omission of some results in Table 4.

Table 4: Results of cascade (decision-level tandem with an AND gate) systems. The ASV system is a pre-trained RawNet3 model [29] whereas the CM is the AASIST or AASIST-L (a lightweight variant) model described in [30]. The SASV, SV and SPF EERs are those defined in the context of the SASV 2022 challenge [14]. t-EER is the tandem EER [13]. Results also reported for a-DCF and t-DCF metrics.

	Cascade system (threshold t)	EERs				min a-DCFs		min t-DCFs	
		SASV-EER	SV-EER	SPF-EER	t-EER	a-DCF1	a-DCF2	t-DCF1	t-DCF2
RawNet3 + AASIST	ASV→CM ($t_{asv}=0.05$)	21.51	44.45	0.65	1.36	0.1492	0.4316	0.0226	0.0644
	ASV→CM ($t_{asv}=0.5$)	0.89	0.74	0.91		0.1492	0.0506	0.0358	0.0918
	CM→ASV ($t_{cm}=0.05$)	1.10	0.69	1.49		0.0240	0.0405	N/A	N/A
	CM→ASV ($t_{cm}=0.5$)	0.76	0.71	0.80		0.0180	0.0382	N/A	N/A
RawNet3 + AASIST-L	ASV→CM ($t_{asv}=0.05$)	21.68	44.71	0.82	1.23	0.1551	0.4405	0.0264	0.0633
	ASV→CM ($t_{asv}=0.5$)	1.02	0.74	1.04		0.0222	0.0560	0.0379	0.0883
	CM→ASV ($t_{cm}=0.05$)	1.58	0.69	2.24		0.0322	0.0453	N/A	N/A
	CM→ASV ($t_{cm}=0.5$)	1.02	0.78	1.30		0.0241	0.0448	N/A	N/A

Table 5: EER, a-DCF, and t-DCF results for diverse spoofing-robust ASV systems. Results shown for four different single-model systems from [35] (S1-S4) and three SASV 2022 challenge submissions (C1-C3). C4 is the same RawNet3 + AASIST ($t_{asv}=0.5$) system for which results are also presented in Table 4. The t-EER and t-DCF can be calculated only for system C4. Tandem metrics require two separate scores, while EER and a-DCF metrics can be calculated for all systems.

	System type	EERs				min a-DCFs		min t-DCFs	
		SASV-EER	SV-EER	SPF-EER	t-EER	a-DCF1	a-DCF2	t-DCF1	t-DCF2
S1	Single-model	1.19	1.82	0.58	N/A	0.0222	0.0578	N/A	N/A
S2	Single-model	1.25	1.27	1.23	N/A	0.0268	0.0417	N/A	N/A
S3	Single-model	1.82	2.51	1.16	N/A	0.0366	0.0853	N/A	N/A
S4	Single-model	2.48	3.32	1.56	N/A	0.0485	0.1068	N/A	N/A
C1	Jointly optimized (<i>score</i>)	0.13	0.11	0.17	N/A	0.0032	0.0060	N/A	N/A
C2	Jointly optimized (<i>embedding</i>)	0.28	0.28	0.28	N/A	0.0067	0.0147	N/A	N/A
C3	Jointly optimized (<i>score</i>)	0.37	0.45	0.26	N/A	0.0080	0.0219	N/A	N/A
C4	Cascade	0.89	0.74	0.91	1.36	0.1492	0.0506	0.0357	0.0358

and for any given set of priors and costs.

Shown in Table 5 are results for single-model, jointly optimised, and cascade systems. Four single models from [35] are simply denoted as ‘S’. Three jointly optimised systems, namely the top-3 performing submissions [32–34] to the SASV 2022 challenge [14], and one cascade system from Table 4 are denoted as ‘C’. EERs, though similarly computable for all systems, are still EERs with the same shortcomings. According to (3), t-DCF values can be computed only for the cascade system and not for either single or jointly optimised systems. Once again, the a-DCF can be applied to the evaluation of all 8 systems when configured to produce a single score, and with explicitly defined priors and costs.

6. Conclusions

We propose an architecture-agnostic detection cost function (a-DCF) designed for the evaluation of spoofing-robust automatic speaker verification (ASV) systems. The a-DCF extends the time-tested DCF adopted by the ASV community decades ago. An alternative to the derided equal error rate metric, the a-DCF reflects the cost of decisions in a Bayes risk sense, with explicitly defined class priors and detection cost model. The a-DCF is also more flexible than its previously proposed t-DCF cousin in terms of supported classifier architectures, so long as they can be configured to produce a single output score. The a-DCF is also task-agnostic and could also be applied to studies involving other biometric traits, or indeed entirely different problems.

Experiments serve as a demonstration of the a-DCF in benchmarking a broad variety of competing spoofing-robust ASV solutions. Even if developing the means to compare the performance of such different approaches was an objective of this work, these

experiments are not, and were not intended to be, sufficient on their own to help identify the most promising. This requires far more extensive experimental analysis which is left for future work.

7. Acknowledgement

We appreciate Dr. Shinji Watanabe for the valuable discussions. The work has been partially supported by the Academy of Finland (Decision No. 349605, project “SPEECHFAKES”) and by the Agence Nationale de la Recherche (ANR) in France (BRUEL project, No. ANR-22-CE39-0009). Experiments of this work used the Bridges2 system at PSC and Delta system at NCSA through allocations CIS210014 and IRI120008P from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

8. References

- [1] A. K. Jain, A. Ross and S. Prabhakar, “An introduction to biometric recognition,” *IEEE Transactions on circuits and systems for video technology*, vol. 14, no. 1, pp. 4–20, 2004.
- [2] J. H. Hansen and T. Hasan, “Speaker recognition by machines and humans: A tutorial review,” *IEEE Signal processing magazine*, vol. 32, no. 6, pp. 74–99, 2015.
- [3] ISO/IEC 30107-1:2016, “Information technology — biometric presentation attack detection — part 1: Framework,” <https://www.iso.org/obp/ui/#iso:std:iso-iec:30107:-1:ed-1:v1:en>, 2016.
- [4] A. Sizov, E. Khoury, T. Kinnunen et al., “Joint speaker verification and antispoofing in the i -vector space,” *IEEE*

- Transactions on Information Forensics and Security*, vol. 10, no. 4, pp. 821–832, 2015.
- [5] M. Todisco, H. Delgado, K. A. Lee et al., “Integrated presentation attack detection and automatic speaker verification: common features and Gaussian back-end fusion,” in *Proc. Interspeech*, 2018.
 - [6] A. Gomez-Alanis, J. A. Gonzalez-Lopez, S. P. Dubagunta et al., “On joint optimization of automatic speaker verification and anti-spoofing in the embedding space,” *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 1579–1593, 2020.
 - [7] J. Li, M. Sun, X. Zhang and Y. Wang, “Joint decision of anti-spoofing and automatic speaker verification by multi-task learning with contrastive loss,” *IEEE Access*, vol. 8, pp. 7907–7915, 2020.
 - [8] M. Sahidullah, H. Delgado, M. Todisco et al., “Integrated spoofing countermeasures and automatic speaker verification: an evaluation on asvspoof 2015,” in *Proc. Interspeech*, 2016.
 - [9] H.-J. Shim, H. Tak, X. Liu et al., “Baseline systems for the first spoofing-aware speaker verification challenge: Score and embedding fusion,” in *Proc. Speaker Odyssey*, 2022.
 - [10] P. Porwik, R. Doroz and K. Wrobel, “An ensemble learning approach to lip-based biometric verification, with a dynamic selection of classifiers,” *Expert Systems with Applications*, vol. 115, pp. 673–683, 2019.
 - [11] S. Ergin, A. K. Uysal, E. S. Gunal et al., “ECG based biometric authentication using ensemble of features,” in *Proc. Iberian Conference on Information Systems and Technologies (CISTI)*. IEEE, 2014.
 - [12] T. Kinnunen, H. Delgado, N. Evans et al., “Tandem Assessment of Spoofing Countermeasures and Automatic Speaker Verification: Fundamentals,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2195–2210, 2020.
 - [13] T. H. Kinnunen, K. A. Lee, H. Tak et al., “t-EER: Parameter-free tandem evaluation of countermeasures and biometric comparators,” *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 01, pp. 1–16, 2023.
 - [14] J.-w. Jung, H. Tak, H.-j. Shim et al., “SASV 2022: The first spoofing-aware speaker verification challenge,” in *Proc. Interspeech*, 2022.
 - [15] X. Liu, M. Sahidullah, K. A. Lee and T. Kinnunen, “Generalizing speaker verification for spoof awareness in the embedding space,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
 - [16] J. Li, M. Sun and X. Zhang, “Multi-task learning of deep neural networks for joint automatic speaker verification and spoofing detection,” in *Proc. APSIPA*, 2019.
 - [17] H.-j. Shim, J.-w. Jung, J.-h. Kim and H.-j. Yu, “Integrated replay spoofing-aware text-independent speaker verification,” *Applied Sciences*, vol. 10, no. 18, pp. 6292, 2020.
 - [18] L. Zhang, Y. Li, H. Zhao and L. Xie, “Explore backend ensemble of speaker verification and spoofing countermeasure,” https://sasv-challenge.github.io/pdfs/2022_descriptions/FlySpeech.pdf, [Online; accessed 5-Feb-2024].
 - [19] J. Heo, J.-h. Kim and H.-s. Shin, “Two methods for spoofing-aware speaker verification: Multi-layer perceptron score fusion model and integrated embedding projector,” in *Proc. Interspeech*, 2022.
 - [20] C. S. Greenberg, “The NIST year 2012 speaker recognition evaluation plan,” *NIST, Technical report*, 2012.
 - [21] G. R. Doddington, M. A. Przybocki, A. F. Martin and D. A. Reynolds, “The NIST speaker recognition evaluation – overview, methodology, systems, results, perspective,” *Speech Communication*, vol. 31, no. 2, pp. 225–254, 2000.
 - [22] E. T. Jaynes, *Probability theory: The logic of science*, Cambridge University Press, Cambridge, 2003.
 - [23] R. O. Duda, P. E. Hart and D. G. Stork, *Pattern Classification*, Wiley, 2 edition, 2001.
 - [24] X. Liu, X. Wang, M. Sahidullah et al., “ASVspoof 2021: Towards spoofed and deepfake speech detection in the wild,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2507–2522, 2023.
 - [25] ISO/ IEC FDIS 19795–1:2021, “Information technology–biometric performance testing and reporting – part 1: Principles and framework,” <https://www.iso.org/obp/ui/#iso:std:iso-iec:19795:-1:ed-1:vl:en>, 2021.
 - [26] L. Ferrer, “Analysis and comparison of classification metrics,” *arXiv preprint arXiv:2209.05355*, 2022.
 - [27] N. Brümmer and E. De Villiers, “The Bosaris toolkit: Theory, algorithms and code for surviving the new dcf,” *Proc. NIST SRE’11 Analysis Workshop*, 2011.
 - [28] A. Nautsch, *Speaker recognition in Unconstrained Environments*, Ph.D. thesis, Technische Universität Darmstadt, 2019.
 - [29] J.-w. Jung, Y. J. Kim, H.-S. Heo et al., “Pushing the limits of raw waveform speaker recognition,” in *Proc. Interspeech*, 2022.
 - [30] J.-w. Jung, H.-S. Heo, H. Tak et al., “AASIST: Audio anti-spoofing using integrated spectro-temporal graph attention networks,” in *Proc. ICASSP*, 2022.
 - [31] X. Wang, J. Yamagishi, M. Todisco et al., “ASVspoof 2019: A large-scale public database of synthesized, converted and replayed speech,” *Computer Speech & Language*, vol. 64, pp. 101114, 2020.
 - [32] A. Alenin, N. Torgashov, A. Okhotnikov et al., “ID R&D team submission description for sasv challenge 2022,” https://sasv-challenge.github.io/pdfs/2022_descriptions/IDVoice.pdf, [Online; accessed 5-Feb-2024].
 - [33] X. Wang, X. Qin, Y. Wang et al., “The DKU-OPPO system for the 2022 spoofing-aware speaker verification challenge,” in *Proc. Interspeech*, 2022.
 - [34] J.-H. Choi, J.-Y. Yang, Y.-R. Jeoung and J.-H. Chang, “HYU submission for the SASV challenge 2022: Reforming speaker embeddings with spoofing-aware conditioning,” in *Proc. Interspeech*, 2022.
 - [35] S. H. Mun, H.-j. Shim, H. Tak et al., “Towards single integrated spoofing-aware speaker verification embeddings,” in *Proc. Interspeech*, 2023.
 - [36] Y. Zhang, Z. Lv, H. Wu et al., “MFA-conformer: Multi-scale feature aggregation conformer for automatic speaker verification,” in *Proc. Interspeech*, 2022.
 - [37] S. H. Mun, J.-w. Jung, M. H. Han and N. S. Kim, “Frequency and multi-scale selective kernel attention for speaker verification,” in *Proc. SLT*, 2023.