Self-supervised 2D face presentation attack detection via temporal sequence sampling

Usman Muhammad, Zitong Yu, and Jukka Komulainen

Abstract—Conventional 2D face biometric systems are vulnerable to presentation attacks performed with different face artefacts, e.g. printouts, video-replays and wearable 3D masks. The research focus in face presentation attack detection (PAD) has been recently shifting towards end-to-end learning of deep representations directly from annotated data rather than designing hand-crafted (low-level) features. However, even the state-of-the-art deep learning based face PAD models have shown unsatisfying generalization performance when facing previously unseen attacks and acquisition conditions due to lack of representative training and tuning data available in the existing public benchmarks. To alleviate this issue, we propose a video pre-processing technique called Temporal Sequence Sampling (TSS) for 2D face PAD by removing the estimated inter-frame 2D affine motion in the view and encoding the appearance and dynamics of the resulting smoothed video sequence into a single RGB image. Furthermore, we leverage the features of a convolutional neural network (CNN) by introducing a self-supervised representation learning scheme, where the labels are automatically generated by the TSS method as the stabilized frames accumulated over video clips of different temporal lengths provide the supervision. The learnt feature representations are then fine-tuned for the downstream task using labeled face PAD data. Our extensive experiments on four public benchmarks, namely Replay-Attack, MSU-MFSD, CASIA-FASD and OULU-NPU, demonstrate that the proposed framework provides state-of-the-art generalization capability and encourage further study in this domain.

Index Terms—Face recognition, presentation attack detection, liveness detection, self-supervised learning, motion compensation.

I. INTRODUCTION

Face recognition (FR) has become an indispensable component in numerous real-world applications requiring reliable person verification and identification, such as device unlocking, banking, smart houses, airports, video surveillance and even law enforcement. One of the biggest concerns for FR is that of coping with the threat of presentation attacks (spoofing). Conventional FR techniques are vulnerable to direct sensor-level attacks, where an artificial biometric sample is presented to confuse the FR system using different presentation attack instruments (PAI), e.g. printouts, displays, paper masks and wearable 3D masks. Compared with other biometric traits, such as iris and fingerprint, face biometric samples of the targeted person are much easier to obtain. For instance, people are sharing their pictures on the internet using different social media platforms from which attackers can easily acquire

photographs to create face artefacts. Since conventional FR algorithms are not inherently capable of discriminating attacks from bona fide faces, dedicated presentation attack detection (PAD) methods are needed to mitigate the vulnerabilities to spoofing.

The accuracy of automatic FR is no longer a major concern in numerous real-world applications, thus the focus in FR research community has shifted towards mitigating the threat posed by presentation attacks. Traditionally, software-based face PAD techniques have been based on hand-crafted or low-level features describing e.g. liveness and motion cues, like eye blinking and lip movements [18], and facial texture [4, 22] and image quality [11, 34] properties between real and fake faces. However, low-level features rely heavily on human experience to extract detailed information and the designed feature spaces might not be able to distinguish subtle differences between bona fide samples and various face artefacts. In the past few years, end-to-end learning of deep features, e.g. deep Convolutional Neural Networks (CNNs) with different loss functions, have been successfully utilized to overcome some limitations of hand-crafted descriptors [15, 17, 21, 26, 27, 32, 33, 36, 38, 40]. A comprehensive overview of the recent advances in deep learning based face PAD can be found in [39].

Although promising results have been achieved, even the state-of-the-art deep learning based face PAD techniques have shown unsatisfying generalization performance when facing unknown operating conditions of unconstrained real-world applications. The generalization issue is largely due to the domain shift between source (train) and target (test) data as the existing public face PAD benchmarks suffer from severe bias across different covariates, including user demographics, PAIs, sensors, image/video resolution, illumination conditions and standoff distance between face and sensor. The domain generalization issues of software-based face PAD methods have been widely acknowledged and the recent trend in face PAD research has been increasingly on improving the performance in: 1) cross-database setup where a method is trained and tested with different datasets [9], and 2) specific intra-database cross-test evaluation protocols where pre-defined subsets of a dataset are used to introduce previously unseen test conditions, e.g. cameras, PAIs, illumination and environments [6].

Leveraging the potential of the state-of-the-art deep learning architectures and tuning well-generalizing face PAD models remain still a very difficult problem due to a large number of hyperparameters and the limited amount of representative training data available in the existing public datasets. The approaches proposed in the context of generalized face PAD

Usman Muhammad and Zitong Yu are with Center for Machine Vision and Signal Analysis, Faculty of Information Technology and Electrical Engineering, University of Oulu, Finland.

Jukka Komulainen is with Visidon Ltd., Oulu, Finland

can be roughly categorized into: 1) face PAD-specific feature learning to capture the intrinsic differences between real and fake faces [17], 2) data augmentation and synthesis [36], 3) auxiliary supervision [21, 33, 36, 38, 40], 4) domain adaptation [20, 23, 32] and generalization [15], and 5) continual detection and learning of novel attack types [26]. While face PAD has been traditionally treated as a “black box” binary classification problem, Jourabloo *et al.* [17] proposed a deep CNN architecture for explicitly extracting PAI dependent spoof noise, e.g. characteristic reflections, color distortions and moiré patterns, from facial images and then use spoof noise modeling for discriminating attacks from real samples. Yang *et al.* [36] introduced a data synthesis technique to simulate digital medium-based spoofing attacks and were able to significantly improve the PAD performance with their augmented training data. Liu *et al.* [21] proposed to increase the generalization of face PAD methods by exploiting spatial and temporal auxiliary supervision, where face depth can be considered as spatial information while remote photoplethysmography (rPPG) signals (i.e. pulse) are used as temporal cues. Several works have also approached the problem of domain-shift in face PAD from domain adaptation and domain generalization point of view. Domain adaptation [20, 23, 32] based approaches exploit some data from the target domain to match the feature distributions of source and target domains, whereas domain generalization [15] based techniques try to minimize the bias between diverse source domains without using any data from the target domain. Rostami *et al.* [26] proposed to tackle the problem of unknown attacks using continual detection and learning of novel attack types and developed a method to update a face PAD model with test samples that do not fit the training distribution in an embedding space.

Despite the generalization ability of the face PAD methods proposed in the literature has been gradually improving, the results have been still far from satisfying for real-world applications. For instance, the performance of methods using auxiliary supervision depends largely on the accuracy of the estimated auxiliary information. Monocular depth estimation from single face images or even short video sequences is rather difficult if active user interaction, e.g. challenge-response approach, is not utilized during liveness check. Also, reliable estimation of rPPG signals is hard when subject is moving or lighting conditions are challenging. A major problem with domain adaptation based approaches is that collecting data from the target domain is expensive or even impossible in some real-world use cases.

In this work, we propose to use spatiotemporal information for face PAD because we argue that both static and dynamic information provide important visual cues for discriminating fake faces from real ones. The successive frames in PAD videos are highly redundant, thus the videos might comprise of hundreds of frames repeating the similar patterns. This makes it difficult to extract meaningful liveness information even with the learning based approaches. Therefore, we propose a simple, yet effective pre-processing method called Temporal Sequence Sampling (TSS) to accumulate appearance and dynamic information of video sequences into single RGB images. This is achieved by splitting an input video sequence into non-

Figure 1. In the proposed TSS method, the input video is first divided into N equal video clips and 2D affine motion is estimated within each video clip based on the trajectories of SURF feature point matches. Then, the estimated inter-frame 2D affine motion is removed from the video frames and resulting clips are accumulated into a single RGB image.

Figure 2. First frame of a sample print attack video clip (left) and the mean of the corresponding stabilized video clip (middle). The result of simple frame averaging (right) is included for comparison to demonstrate the amount of inter-frame 2D affine motion in the original print attack video clip.

overlapping segments and then estimating the trajectories of extracted keypoints within each video clip. We focus on the problem of print and display attacks (consisting of both digital photos and video-replays) when the PAIs can be considered as planar 2D objects. We stabilize each video segment by removing the inter-frame 2D affine motion estimated based on the keypoint trajectories before accumulating the video frames into a single image. Fig. 1 provides an illustration of the steps described above. A comparison between straightforward frame aggregation and the output of the proposed TSS approach is shown in Fig. 2, which highlights the amount of inter-frame 2D affine motion in the original print attack video clip.

It is worth noting that the cumulative 2D affine transform estimated within a video clip is not directly used for face alignment but to enrich the spatiotemporal discrepancies between real 3D faces and flat 2D face artefacts in the observed view. The proposed approach can handle also print attacks where the 2D surface is warped as bending a photograph leads to a highly distorted cumulative 2D affine mapping that is not characteristic for real faces. The problem of video-replay attacks exhibiting also non-rigid facial motion is tackled by focusing on appearance information, which has shown promising generalization in detecting display attacks, e.g. in [7], due to evident screen bezels, video compression artefacts, display noise signatures, moiré effects, and luminance and color distortions, for instance.

Recently, self-supervised learning [16] has been receiving

Figure 3. An illustration of the proposed self-supervised and TSS training tasks. During the self-supervised learning phase, the CNN receives unlabeled TSS sampling outputs accumulated over different temporal lengths and the pretext task is to predict the length of the original video clip. The learnt feature representations are then fine-tuned on labeled TSS encoded video clips by performing PAD as the downstream task (stage 1). Finally, the BiLSTM sub-network is trained using the fine-tuned features (stage 2) to make the final PAD decision.

increasing attention as solving pretext tasks, like patch location, order and rotation prediction, in unsupervised manner has shown to be successful in learning meaningful and even more interpretable visual representations from the data itself, thus mitigating the need for human annotations for the downstream task. Inspired by the work on frame order prediction where 3D CNNs or optical flow information is utilized [19, 35], we propose a self-supervised learning scheme where the pretext task is to predict the length of the original video clip based on the TSS encoded data. To be more specific, the stabilized frames accumulated over video segments of different temporal lengths provide the supervision for training a 2D CNN with the aim of learning more meaningful representations from the videos aggregated into single RGB images. The learnt visual features are then fine-tuned for the downstream task of face PAD.

The main contributions of this work can be summarized as follows:

- 1) In order to reduce temporal redundancy and remove inter-frame 2D affine motion in videos, Temporal Sequence Sampling (TSS) is introduced to encode video clips into a compact representation in form of a single RGB image.
- 2) The need for annotated data in face PAD is mitigated using self-supervised learning.
- 3) The effectiveness of the proposed approach is demonstrated in several widely used cross-database settings and using the official cross-test evaluation protocols of the OULU-NPU database [6] as promising generalization ability with new state-of-the-art results is achieved.

In addition, the source code for reproducing our results will

be shared to the research community.

II. PROPOSED METHOD

The backbone of the proposed face PAD approach is the TSS method which removes inter-frame 2D affine motion within a video segment and accumulates frames of the resulting motion-compensated video clip into a single RGB image (see Fig. 1 and Fig. 2). The main architecture of our face PAD framework is illustrated in Fig. 3. During the self-supervised learning phase, the CNN receives unlabeled TSS encoded images accumulated over video segments of different temporal lengths and the pretext task is to predict the length of the original video clip. The learnt feature representations are then fine-tuned on labeled TSS encoded video clips by performing PAD as the downstream task (stage 1 in Fig. 3). Finally, a bidirectional long short-term memory (BiLSTM) [28] sub-network is trained using the fine-tuned CNN features to make the final face PAD decision (stage 2 in Fig. 3). A more detailed description of the proposed TSS method and self-supervised learning scheme is provided in the following, while the implementation details and the training process are discussed later in Section III-C.

A. Temporal Sequence Sampling (TSS)

The steps of the proposed Temporal Sequence Sampling method are illustrated in Fig 1. First, the input video is equally partitioned into S non-overlapping segments (clips), where each video clip contains the same number of frames, e.g. 45 frames. We estimate the 2D affine motion between all adjacent frames of a video clip based on sparse point correspondences. We use first Speeded Up Robust Features (SURF) descriptor

[1] to detect keypoints from both the face and background regions in each video frame and then find the corresponding points between all adjacent video frames using the Hamming distance.

The M-estimator SAmple Consensus (MSAC) algorithm [31] is utilized to mitigate the impact of incorrect point correspondences and to get a robust estimates of the 2D affine transforms between between the adjacent frames. MSAC is an improved version of RANdom SAmple Consensus (RANSAC) where M-estimator is introduced to set outlier point correspondences a constant weight while inliers are weighted based on how well they fit the estimated transformation.

The resulting 2D affine transform between adjacent frames is a 3×3 matrix:

$$\begin{bmatrix} a_1 & a_2 & t_x \\ a_3 & a_4 & t_y \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

where a_n represents scale, rotation, and shearing transforms and t_x and t_y correspond to translation.

The 2D affine transform between two frames F_i and F_{i+1} is denoted with A_i when the cumulative 2D affine transform of a frame F_i with respect to the first frame F_0 (reference) of the video segment corresponds to cascaded inter-frame transforms:

$$A_{0,i} = \prod_{j=0}^{i-1} A_j \quad (2)$$

The estimated cumulative transformation $A_{0,i}$ is used to remove inter-frame 2D affine motion by warping the each frame F_i relative to the first frame F_0 of the video clip. However, we convert the 2D affine transformation described above into a simpler and more stable four parameter transform to produce the final motion-compensated video clip:

$$\begin{bmatrix} s \cos \theta & -s \sin \theta & t_x \\ s \sin \theta & s \cos \theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \quad (3)$$

where s is a scale factor and θ rotation angle. It is worth highlighting that the cumulative transform aims at removing the inter-frame 2D motion but the frame F_i is not necessarily aligned with reference frame F_0 due to the cumulative errors in estimating the motion between adjacent frames and changes in the observed view within the video clip.

Finally, we take the temporal average of the motion-compensated frames to encode the whole video clip into a single RGB image. An example of the TSS output is shown in Fig. 2.

B. Self-supervised representation learning

The amount and nature of spatiotemporal variations within video segments, and, consequently, the TSS outputs and the resulting CNN feature representations, depend largely on the duration of the input sequence. Therefore, the key idea of our self-supervision scheme is to generate sets of TSS encoded video clips with different temporal lengths L and then learn the spatiotemporal variations and context across these different

length settings. To be more specific, we first generate the T classes of TSS outputs with e.g. $L = \{5, 15, 30, 45, 60\}$ and then use these labels to train a deep CNN with softmax loss to predict the length of a given TSS encoded video segment. After the self-supervised spatiotemporal context adaptation step, the learnt 2D visual features are then further fine-tuned for the actual downstream task of face PAD and finally the BiLSTM sub-network is trained using the resulting CNN features.

III. EXPERIMENTAL SETUP

In the following, we introduce briefly the public benchmark face PAD datasets and describe the evaluation metrics and protocols used in our experimental analysis. Finally, the implementation details of the proposed approach are also provided.

A. Experimental data

To assess the generalization of the proposed face PAD approach, we considered four widely used publicly available databases consisting of bona fide and 2D face presentation attack videos, namely Idiap Replay-Attack Database [8] (denoted as I), CASIA Face Anti-Spoofing Database [41] (denoted as C), MSU Mobile Face Spoofing Database [34] (denoted as M), and OULU-NPU Database [6] (denoted as O).

Idiap Replay-Attack Database [8] consists of bona fide and attack videos of 50 subjects captured under two different lighting conditions. Five different attacks are launched with iPhone 3GS (digital photo and video-replay), 1st generation iPad (digital photo and video-replay) and hard copies. All videos are recorded with a built-in webcam of a Macbook Air laptop. Altogether the database contains 1,200 videos, which are divided into three subject-disjoint subsets for training, development and testing (15, 15 and 20 subjects, respectively).

CASIA Face Anti-Spoofing Database (CASIA-FASD) [41] contains bona fide and attack videos of 50 subjects recorded with three different of imaging qualities (low, normal and high) and considers three kinds of attack presentations (warped photo, cut-photo and video-replay). Consequently, each subject has three kinds of bona fide videos and altogether nine different attack presentations. Altogether the database contains 600 videos, which are divided into two subject-disjoint subsets for training and testing (20 and 30 subjects, respectively).

MSU Mobile Face Spoofing Database (MSU-MFSD) [34] includes bona fide and attack videos of 35 subjects recorded with two mobile devices (a Google Nexus 5 smartphone and a MacBook Air laptop). Three kinds of attack presentations are considered, including two video-replay attacks of different quality (iPhone 5S and iPad Air) and a print attack. Consequently, each subject has two kinds of real videos and six kinds of attack presentations. Altogether the database contains 280 videos, which are divided into two subject-disjoint subsets for the training and testing (15 and 20 subjects, respectively).

OULU-NPU Database [6] is one of the most recent commonly used face PAD datasets. It contains bonafide and attack videos of 55 subjects recorded in several different acquisition conditions (six high-resolution smartphone front cameras and

three sessions) and considers two kinds of print attacks and two kinds of video-replay attacks. Four cross-test protocols are used to evaluate the generalization performance of a face PAD method across different covariates. Protocols 1, 2, and 3 introduce a single previously unseen test condition, namely illumination, PAI and sensor, respectively, while the fourth and most challenging protocol evaluates the generalization performance simultaneously across unknown sensors, attacks and illumination conditions. Altogether the database contains 5,940 videos, which are divided into three subject-disjoint subsets for training, development and testing (20, 15 and 20, respectively).

B. Evaluation metrics and protocols

All four datasets are used in our cross-database experiments, while only the OULU-NPU database is utilized also for intra-database experiments following its official cross-test protocols that assess generalization across different covariates.

For our cross-database experiments, we follow the widely used evaluation metrics and protocols introduced in [9]. The results are reported using Half Total Error Rate (HTER), which denotes the mean of the False Acceptance Rate (FAR) and False Rejection Rate (FRR). The HTER is computed on the test set of the target domain using the threshold τ corresponding to the equal error rate (EER) operating point on the development set of the source domain. In the case of the CASIA-FASD and MSU-MFSD datasets, the threshold τ is computed on the training set because they lack pre-defined validation sets (see Section III-A).

The intra-database results on the official cross-test protocols of the OULU-NPU database are reported in terms of Average Classification Error Rate (ACER), which denotes the mean of Attack Presentation Classification Error Rate (APCER) and Bona Fide Presentation Classification Error Rate (BPCER). APCER and BPCER essentially correspond to FAR and FRR, respectively, but APCER is computed separately for each PAI, e.g. print and video-replay, and the final PAD performance corresponds to the attack with the highest APCER, i.e. the most successful PAI. Similarly to the HTER, the ACER is computed on the test set using the threshold τ corresponding to the EER operating point on the development set.

C. Implementation details

We utilize both the face and background regions for PAD, thus no face cropping is applied. The TSS method processes the frames of the input video clips at their native resolution and the TSS generated accumulated output frames are resized to 224×224 according to the input image size of the pre-trained CNN (ResNet-101 [14]). The video segment length for TSS methods was set to 45 and the number of TSS encoded segments depends on the total length of an input video sequence. For instance, a video of 270 frames provides six TSS encoded video clips. No data augmentation is applied during training.

To evaluate the generalization performance of the proposed TSS method, the pre-trained CNN is fine-tuned using stochastic gradient descent (SGD) with mini-batch size of

32 and validation frequency of 30, and by shuffling every epoch. We do not use fixed epochs because an early stopping function is utilized to automatically stop the model training when over-fitting is observed [25]. The learning rate is fixed to 0.0001 in our cross-database experiments, while we adjust the learning rate to 0.001 on all four intra-database test protocols of the OULU-NPU dataset.

The fine-tuned feature vectors are extracted from the output of the last pooling layer with size of 2048. The BiLSTM sub-network is trained using the cross-entropy loss and Adam optimizer is utilized with fixed learning rate of 0.0001. The number of hidden units is fixed to 100 for the cross-database experiments, while the number of hidden units is increased to 500 on all four intra-database test protocols of the OULU-NPU dataset. We set the recurrent weights with He initializer [13] which performs the best in all scenarios of our experiments.

During the self-supervised training stage, we first fine-tune the CNN by keeping the same parameters as mentioned above on unlabeled data, i.e. sets of TSS outputs with different temporal length combinations, for pretext task learning. Then, the model is fine-tuned on the binary labels (bona fide and attack) of the downstream task by replacing the new fully connected layer with the output size of 2. Finally, BiLSTM takes the input of the last average pooling layer of the fine-tuned CNN and gives the final binary PAD decision.

Figure 4. Cross-database performance with different video segment lengths.

IV. EXPERIMENTAL RESULTS

Our experimental analysis focuses on assessing the generalization performance of the proposed approach in two settings: 1) different cross-database configurations, and 2) the official intra-database cross-test protocols of the OULU-NPU dataset. In the following, we first investigate the effect of the input video sequence length on the TSS method and then study the effectiveness of the proposed self-supervised learning scheme in cross-database setup between the CASIA-FASD and Replay-Attack datasets. Finally, we compare the performance of the proposed approach against the state of the art on several widely used cross-database scenarios and the four official evaluation protocols of the OULU-NPU database.

Table I
CROSS-DATABASE PERFORMANCE OF THE PROPOSED SELF-SUPERVISED
LEARNING SCHEME IN TERMS OF HTER (%).

Number of classes	Segment lengths	Train C / Test I	Train I / Test C
2	30, 45	8.2	18.4
3	15, 30, 45	5.9	15.2
5	5, 15, 30, 45, 60	24.8	22.1

Table II
CROSS-DATABASE PERFORMANCE IN TERMS OF HTER (%) ON THE
REPLAY-ATTACK AND CASIA-FASD DATABASES. COMPARATIVE
RESULTS ARE OBTAINED FROM [37].

Method	Train CASIA-FASD Test Replay-Attack	Train Replay-Attack Test CASIA-FASD
LBP [22]	47.0	39.6
LBP-TOP [10]	49.7	60.6
Color-LBP [3]	30.3	37.7
Motion-Mag [2]	50.1	47.0
Spectral cubes [24]	34.4	50.0
Auxiliary [21]	27.6	28.4
FaceDs [17]	28.5	41.1
STASN [36]	31.5	30.9
DSGTD [33]	17.0	22.8
CDCN [37]	6.5	29.8
TSS with ResNet	30.4	39.9
TSS with ResNet-BiLSTM	9.3	18.1
Self-supervised learning	5.9	15.2

A. The effect of video segment length

We begin our experiments by exploring how the performance of the proposed TSS method depends on the length of the video segment. We examine the generalization performance by varying the length of video segments L from 5 to 60 frames. The cross-database results with Replay-Attack and CASIA-FASD databases shown in Fig. 4 depict that the HTER decreases as the number of frames per video segment increases. However, when we further increase the temporal length of frames (more than 45 frames), the face PAD performance on CASIA-FASD and Replay-Attack datasets starts decreasing. Therefore, we set $L = 45$ where the best performance is achieved, i.e. HTER of 9.3% on Replay-Attack and 18.1% on CASIA-FASD database.

B. Effectiveness of self-supervised learning

For assessing the performance of the proposed self-supervised learning scheme, we construct three sets of TSS encoded video segments with varying number of temporal lengths $L = \{5, 15, 30, 45, 60\}$ for the pretext task. It can be observed from Table I that the use of pretext task with three temporal lengths ($L = \{15, 30, 45\}$) leads to best cross-database performance improving the HTER with 3.5% and 2.9% for Replay-Attack and CASIA-FASD databases, respectively.

C. Comparison against the state of the art

The results presented in Table II depict that our TSS method achieves astonishing cross-database performance between the Replay-Attack and CASIA-FASD datasets and that the generalization ability can be further improved when the proposed self-supervised learning stage is included in training. The TSS method combined with self-supervised learning obtains the

Table III
COMBINED CROSS-DATABASE EVALUATION USING MSU-MFSD (M),
IDIAP REPLAY-ATTACK (I), CASIA-FASD (C) AND OULU-NPU (O)
DATABASES. COMPARATIVE RESULTS ARE OBTAINED FROM [15].

Method	M&I to C		M&I to O	
	HTER(%)	AUC(%)	HTER(%)	AUC(%)
MS-LBP [22]	51.16	52.09	43.63	58.07
LBP-TOP [10]	45.27	54.88	47.26	50.21
Color-LBP [4]	55.17	46.89	53.31	45.16
IDA [34]	45.16	58.80	54.52	42.17
MADDG [30]	41.02	64.33	39.35	65.10
SSDG-M [15]	31.89	71.29	36.01	66.88
TSS with ResNet	29.58	79.44	37.62	70.26
TSS with ResNet-BiLSTM	28.66	83.73	30.12	79.06

state of the art in this widely used cross-database scenario by achieving an HTER improvement from 9.3% to 5.9% on the Replay-Attack dataset and from 18.1% to 15.2% on the CASIA-FASD dataset, respectively. Thus, the proposed self-supervised learning scheme indeed helps in fine-tuning a 2D CNN to learn more meaningful representations from the TSS encoded video segments.

Table III presents the generalization performance of the proposed TSS approach in another cross-database setting combining MSU-MFSD and Replay-Attack databases for training and CASIA-FASD and OULU-NPU databases for testing. The proposed TSS method with CNN-BiLSTM framework achieves the best results and significant improvement with respect the state of the art in HTER from 31.89% to 28.66% on the CASIA-FASD and 36.01% to 30.12% on the OULU-NPU dataset.

The results of the intra-database experiments following the official evaluation protocols of the OULU-NPU database are presented in Table IV. From these results it can be seen that the proposed TSS method with CNN-BiLSTM framework ranks first on the protocols 1 and 2 of the OULU-NPU database obtaining ACER of 0.1% and 0.6%, respectively, and achieves very competitive performance of 1.5% and 7.1% in terms of ACER on the protocols 3 and 4, respectively.

We have also included the performance of the proposed TSS method in the cross-database experiments and the intra-database tests on the OULU-NPU without the BiLSTM sub-network in order to demonstrate the importance of the BiLSTM component on the final face PAD performance.

D. Network visualization and analysis

In this section, we generate the class activation maps by using the Grad-CAM [29] to get an intuition about the decisions that the network makes when producing a particular prediction. Sample images of a real face, a video-replay attack and a print attack are presented for further analysis in Fig. 5. The first row represents sample images of the real faces where one can see that the network gives clear focus on the actual facial region due to e.g. head motion, non-rigid facial movements, eye blinking and skin texture, while background region does not provide liveness cues. In contrast, the sample images of video-replay attacks and print attacks in the second and third rows, respectively, depict that the discriminative visual cues are display device (PAI) related and attention more dispersed

Table IV
INTRA-DATABASE EVALUATION ON THE FOUR OFFICIAL PROTOCOLS OF
THE OULU-NPU DATABASE. COMPARATIVE RESULTS ARE OBTAINED
FROM [37].

Protocol	Method	APCER(%)	BPCER(%)	ACER(%)
1	DeepPixBIS [12]	0.8	0.0	0.4
	GRADIENT [5]	1.3	12.5	6.9
	Auxiliary [21]	1.6	1.6	1.6
	FaceDs [17]	1.2	1.7	1.5
	STASN [36]	1.2	2.5	1.9
	DSGTD [33]	2.0	0.0	1.0
	CDCN [37]	0.4	0.0	0.2
	BiFPN [27]	3.1	0.8	2.0
	TSS with ResNet	0.6	10.3	5.5
2	TSS with ResNet-BiLSTM	0.0	0.2	0.1
	DeepPixBIS [12]	11.4	0.6	6.0
	GRADIENT [5]	3.1	1.9	2.5
	Auxiliary [21]	2.7	2.7	2.7
	FaceDs [17]	4.2	4.4	4.3
	STASN [36]	4.2	0.3	2.2
	DSGTD [33]	2.5	1.3	1.9
	CDCN [37]	1.8	0.8	1.3
	BiFPN [27]	1.7	1.1	1.4
3	TSS with ResNet	2.0	2.1	2.1
	TSS with ResNet-BiLSTM	0.4	0.8	0.6
	DeepPixBIS [12]	11.7±19.6	10.6±14.1	11.1±9.4
	GRADIENT [5]	2.6±3.9	5.0±5.3	3.8±2.4
	Auxiliary [21]	2.7±1.3	3.1±1.7	2.9±1.5
	FaceDs [17]	4.0±1.8	3.8±1.2	3.6±1.6
	STASN [36]	4.7±3.9	0.9±1.2	2.8±1.6
	DSGTD [33]	3.2±2.0	2.2±1.0	2.7±0.6
	CDCN [37]	1.7±1.5	2.0±1.2	1.8±0.7
4	BiFPN [27]	0.7±0.7	0.3±0.7	0.5±0.6
	TSS with ResNet	7.2±8.3	3.9±3.4	5.5±3.0
	TSS with ResNet-BiLSTM	2.5±1.8	0.5±0.6	1.5±0.8
	DeepPixBIS [12]	36.7±29.7	13.3±14.1	25.0±12.7
	GRADIENT [5]	5.0±4.5	15.0±7.1	10.0±5.0
	Auxiliary [21]	9.3±5.6	10.4±6.0	9.5±6.0
	FaceDs [17]	1.2±6.3	6.1±5.1	5.6±5.7
	STASN [36]	6.7±10.6	8.3±8.4	7.5±4.7
	DSGTD [33]	6.7±7.5	3.3±4.1	5.0±2.2
	CDCN [37]	4.2±3.4	5.8±4.9	5.0±2.9
	BiFPN [27]	2.5±3.2	3.3±4.1	2.9±3.4
	TSS with ResNet	5.7±4.5	16±13	10.8±5.3
	TSS with ResNet-BiLSTM	4.7±10.5	9.2±10.4	7.1±5.3

Figure 5. Sample class activation maps for TSS encoded videos corresponding to a real face (first row), print attack (second row) and video-replay attack (third row).

and focusing also on background regions, i.e. non-face related information.

V. CONCLUSIONS

In this paper, we addressed the generalization issue of 2D face presentation attack detection (PAD). We proposed a Temporal Sequence Sampling (TSS) method that removes the inter-frame 2D affine motion within short video clips estimated using keypoint matching and encodes the appearance and dynamics of the resulting frames into a single colour image. We also introduced a self-supervised learning scheme where the stabilized video frames accumulated over sequences of different temporal lengths provide the supervision to train a 2D convolutional neural network. We conducted extensive experimental analysis using the official intra-test protocols of

the OULU-NPU database and cross-database scenarios on four public face PAD databases to demonstrate the robustness of the proposed framework.

The proposed approach needs capturing and processing of relatively long input sequences, i.e. approximately two seconds of video, in order to achieve robust face PAD performance, thus it cannot be used in authentication applications requiring real-time response or biometric systems operating on single facial images. A drawback of encoding registered video clips into a compact representation in form of a single colour image is that the subtle inter-frame motion (direction) information is lost due to frame aggregation. Therefore, we plan to extend our work by developing methods that explicitly model the geometrical differences in the feature or facial landmark based trajectories between stabilized real and attack videos. In this work, we focused only on detecting attacks launched with 2D PAI, i.e. prints and displays, thus it is yet unknown whether the proposed approach generalizes well on unseen or other types of facial artefacts, including paper and 3D masks. In the future, we will explore the robustness of our method on new emerging face PAD datasets and evaluation protocols.

VI. ACKNOWLEDGMENTS

The financial support of the Tauno Tönning Foundation is gratefully acknowledged.

REFERENCES

- [1] H. Bay, T. Tuytelaars, and L. Van Gool. SURF: Speeded up robust features. In A. Leonardis, H. Bischof, and A. Pinz, editors, *European Conference on Computer Vision (ECCV)*, pages 404–417. Springer, 2006.
- [2] S. Bharadwaj, T. I. Dhamecha, M. Vatsa, and R. Singh. Computationally efficient face spoofing detection with motion magnification. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 105–110, 2013.
- [3] Z. Boulkenafet, J. Komulainen, and A. Hadid. Face anti-spoofing based on color texture analysis. In *IEEE International Conference on Image Processing (ICIP)*, pages 2636–2640, 2015.
- [4] Z. Boulkenafet, J. Komulainen, and A. Hadid. Face spoofing detection using colour texture analysis. *IEEE Transactions on Information Forensics and Security*, 11(8):1818–1830, 2016.
- [5] Z. Boulkenafet, J. Komulainen, Z. Akhtar, A. Benlamoudi, D. Samai, S. E. Bekhouche, A. Ouafi, F. Dornaika, A. Taleb-Ahmed, L. Qin, et al. A competition on generalized software-based face presentation attack detection in mobile scenarios. In *International Joint Conference on Biometrics (IJCB)*, pages 688–696, 2017.
- [6] Z. Boulkenafet, J. Komulainen, L. Li, X. Feng, and A. Hadid. OULU-NPU: A mobile face presentation attack database with real-world variations. In *IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 612–618, 2017.
- [7] Z. Boulkenafet, J. Komulainen, and A. Hadid. On the generalization of color texture-based face anti-spoofing. *Image and Vision Computing*, 77:1–9, 2018.

- [8] I. Chingovska, A. Anjos, and S. Marcel. On the effectiveness of local binary patterns in face anti-spoofing. In *International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–7, 2012.
- [9] T. de Freitas Pereira, A. Anjos, J. M. De Martino, and S. Marcel. Can face anti-spoofing countermeasures work in a real world scenario? In *International Conference on Biometrics (ICB)*, 2013.
- [10] T. de Freitas Pereira, J. Komulainen, A. Anjos, J. M. De Martino, A. Hadid, M. Pietikäinen, and S. Marcel. Face liveness detection using dynamic texture. *EURASIP Journal on Image and Video Processing*, 2014(1):1–15, 2014.
- [11] J. Galbally and S. Marcel. Face anti-spoofing based on general image quality assessment. In *International Conference on Pattern Recognition (ICPR)*, pages 1173–1178, 2014.
- [12] A. George and S. Marcel. Deep pixel-wise binary supervision for face presentation attack detection. In *International Conference on Biometrics (ICB)*, pages 1–8, 2019.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *International Conference on Computer Vision (ICCV)*, pages 1026–1034, 2015.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [15] Y. Jia, J. Zhang, S. Shan, and X. Chen. Single-side domain generalization for face anti-spoofing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8484–8493, 2020.
- [16] L. Jing and Y. Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [17] A. Jourabloo, Y. Liu, and X. Liu. Face de-spoofing: Anti-spoofing via noise modeling. In *European Conference on Computer Vision (ECCV)*, pages 290–306, 2018.
- [18] K. Kollreider, H. Fronthaler, M. I. Faraj, and J. Bigun. Real-time face detection and motion analysis with application in “liveness” assessment. *IEEE Transactions on Information Forensics and Security*, 2(3):548–558, 2007.
- [19] H.-Y. Lee, J.-B. Huang, M. Singh, and M.-H. Yang. Unsupervised representation learning by sorting sequences. In *International Conference on Computer Vision (ICCV)*, pages 667–676, 2017.
- [20] H. Li, W. Li, H. Cao, S. Wang, F. Huang, and A. C. Kot. Unsupervised domain adaptation for face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 13(7):1794–1809, 2018.
- [21] Y. Liu, A. Jourabloo, and X. Liu. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 389–398, 2018.
- [22] J. Määttä, A. Hadid, and M. Pietikäinen. Face spoofing detection from single images using micro-texture analysis. In *International Joint Conference on Biometrics (IJCB)*, pages 1–7. IEEE, 2011.
- [23] A. Mohammadi, S. Bhattacharjee, and S. Marcel. Domain adaptation for generalization of face presentation attack detection in mobile settings with minimal information. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [24] A. Pinto, H. Pedrini, W. R. Schwartz, and A. Rocha. Face spoofing detection through visual codebooks of spectral temporal cubes. *IEEE Transactions on Image Processing*, 24(12):4726–4740, 2015.
- [25] L. Prechelt. Early stopping-but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer, 1998.
- [26] M. Rostami, L. Spinoulas, M. Hussein, J. Mathai, and W. Abd-Almageed. Detection and continual learning of novel face presentation attacks. In *International Conference on Computer Vision (ICCV)*, 2021.
- [27] K. Roy, M. Hasan, L. Rupty, M. S. Hossain, S. Sengupta, S. N. Taus, and N. Mohammed. Bi-FPNFAS: Bi-directional feature pyramid network for pixel-wise face anti-spoofing by leveraging fourier spectra. *Sensors*, 21(8), 2021.
- [28] M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [29] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *International Conference on Computer Vision (ICCV)*, pages 618–626, 2017.
- [30] R. Shao, X. Lan, J. Li, and P. C. Yuen. Multi-adversarial discriminative deep domain generalization for face presentation attack detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10023–10031, 2019.
- [31] P. Torr and A. Zisserman. Robust parametrization and computation of the trifocal tensor. *Image and Vision Computing*, 15(8):591–605, 1997.
- [32] G. Wang, H. Han, S. Shan, and X. Chen. Unsupervised adversarial domain adaptation for cross-domain face presentation attack detection. *IEEE Transactions on Information Forensics and Security*, 16:56–69, 2020.
- [33] Z. Wang, Z. Yu, C. Zhao, X. Zhu, Y. Qin, Q. Zhou, F. Zhou, and Z. Lei. Deep spatial gradient and temporal depth learning for face anti-spoofing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [34] D. Wen, H. Han, and A. K. Jain. Face spoof detection with image distortion analysis. *IEEE Transactions on Information Forensics and Security*, 10(4):746–761, 2015.
- [35] D. Xu, J. Xiao, Z. Zhao, J. Shao, D. Xie, and Y. Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10334–10343, 2019.
- [36] X. Yang, W. Luo, L. Bao, Y. Gao, D. Gong, S. Zheng, Z. Li, and W. Liu. Face anti-spoofing: Model matters,

- so does data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3507–3516, 2019.
- [37] Z. Yu, C. Zhao, Z. Wang, Y. Qin, Z. Su, X. Li, F. Zhou, and G. Zhao. Searching central difference convolutional networks for face anti-spoofing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [38] Z. Yu, X. Li, J. Shi, Z. Xia, and G. Zhao. Revisiting pixel-wise supervision for face anti-spoofing. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(3):285–295, 2021.
- [39] Z. Yu, Y. Qin, X. Li, C. Zhao, Z. Lei, and G. Zhao. Deep learning for face anti-spoofing: A survey. *arXiv preprint arXiv:2106.14948*, 2021.
- [40] Z. Yu, J. Wan, Y. Qin, X. Li, S. Z. Li, and G. Zhao. NAS-FAS: Static-dynamic central difference network search for face anti-spoofing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):3005–3023, 2021.
- [41] Z. Zhang, J. Yan, S. Liu, Z. Lei, D. Yi, and S. Z. Li. A face antispoofing database with diverse attacks. In *International Conference on Biometrics (ICB)*, pages 26–31, 2012.