
Deepfake Caricatures: Amplifying attention to artifacts increases deepfake detection by humans and machines

Camilo Fosco*
MIT
camilolu@mit.edu

Emilie Josephs*
MIT
ejosephs@mit.edu

Alex Andonian
MIT
andonian@mit.edu

Allen Lee
MIT
allenlee@mit.edu

Xi Wang
MIT
nicole.xiwang@gmail.com

Aude Oliva
MIT
oliva@mit.edu

Abstract

Deepfakes pose a serious threat to our digital society by fueling the spread of misinformation. It is essential to develop techniques that both detect them, and effectively alert the human user to their presence. Here, we introduce a novel deepfake detection framework that meets both of these needs. Our approach learns to generate attention maps of video artifacts, semi-supervised on human annotations. These maps make two contributions. First, they improve the accuracy and generalizability of a deepfake classifier, demonstrated across several deepfake detection datasets. Second, they allow us to generate an intuitive signal for the human user, in the form of "Deepfake Caricatures": transformations of the original deepfake video where attended artifacts are exacerbated to improve human recognition. Our approach, based on a mixture of human and artificial supervision, aims to further the development of countermeasures against fake visual content, and grants humans the ability to make their own judgment when presented with dubious visual media.

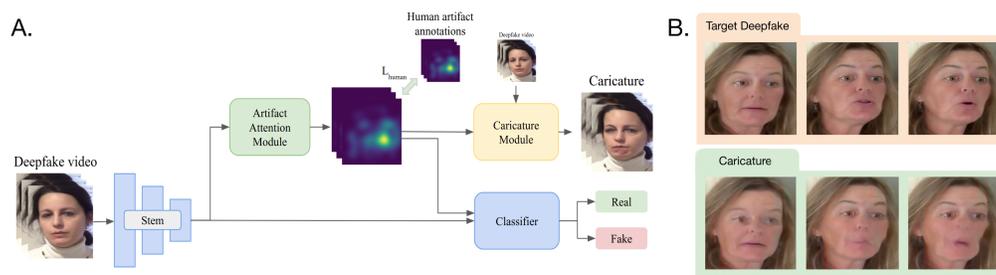


Figure 1: A. Overview of our framework, which uses an artifact attention module that highlights defects to improve classifier performance. Through supervision with human labels, we guide this module to learn artifacts visible to humans, then amplify them to generate Deepfake Caricatures: transformations of the original videos where artifacts are more visible. B. Example frames of a standard deepfake video (right) and a deepfake caricature (left). Our method distorts the fake video to exacerbate artifacts and unnatural motion produced by the deepfake process. The caricatures are best experienced in video form. See supp. material for caricature video clips.

1 Introduction

Fake or manipulated video media (“deepfakes”) pose a clear threat to safety and wellbeing in online spaces. Computer-generated fake videos are increasingly indistinguishable from genuine videos to the human eye [37, 16], pointing to a need for robust and effective detection systems

To date, research on deepfakes mitigation has largely focused on developing models that can identify fake videos. However, this approach does not take into account how human users engage with information. Simply being told that information is fake is not always enough to prevent a person from believing it. Prior exposure to plausible information raises the likelihood that a person will later endorse it as accurate, a cognitive bias known as the fluency heuristic [28]. This can happen even when they are alerted that the information might be false [40] [6]. Thus, preventing human users from absorbing fake information may require a more thoughtful intervention that simply alerting human users to the probability that a video is fake. Here, we explore for the first time the feasibility of exacerbating artifacts in deepfake videos to facilitate early detection by humans.

Overall, to make a societal impact, the development of deepfake detection models must proceed in lockstep with the development of innovative ways of alerting human users to a video’s truth status from the first glimpse. Here, we introduce a framework that achieve both of these goals, by capitalizing on distortions inherent in fake and manipulated videos. Our model, “CariNet”, explicitly attends to artifacts in deepfakes, using a novel *Artifact Attention Module* which leverages both human supervision and machine supervision to learn what distortions to attend to. CariNet also generates **deepfake caricatures**, distorted versions of deepfakes, using an *Caricature Generation Module* that magnifies unnatural movements in videos.

The success of our approach is evaluated on a tandem battery of machine and human detection tests. We show that our model can detect deepfakes, competitive with the highest performing models, and yields excellent generalization among deepfake generation methods. We also show in a controlled human study that our caricature module greatly increases human performance in fake video detection.

2 Related work

2.1 Deepfake detection systems

Deepfake detection is a young but active field, with many novel architectures proposed to detect videos or images where faces have been digitally manipulated [3] [5] [36] [42, 17][34].

A number of successful approaches detect fake faces based on **image artifacts**, which emerge in fake media through the way they are created. For example, Durall et al. [15] distinguish real and fake video based on their amplitude spectra, and Li et al. [24] leverage image warping artifacts resulting from image reconstruction. Others rely on inconsistencies in color [21] or pose [52] which arise from differing camera views or lighting conditions between the source and target faces. Others target the boundary artifacts [23] that arise when a fake face is poorly blended with the background. Anomalous biological signals have also been useful for determining video authenticity: Ciftci et al. [12] show that real videos contain signals of blood volume changes which fakes cannot replicate, and Matern et al. [31] use imprecise rendering of face components, like eye color and specularities, teeth, and shadows to improve fake detection.

Other approaches employ **attention mechanisms** that highlight specific parts of the input. Tolosana et al. [48] propose to segment frames into face regions (nose, eyes, etc.) and process each region with a separate model. Nguyen et al.[35] and Li et al. [22] train models to both classify the video and estimate a binary map highlighting the manipulated regions. Li et al. [23] focus on face swapping, and estimate the binary mask corresponding to the blending boundary in their detection pipeline. Others, such as Dang et al. [45] and Zhao et al. [55], use attention mechanisms which learn to focus on manipulated regions, and weight the detector’s responses towards those regions. This is similar to our approach, however, given our joint goal of increasing human detection as well as machine detection, we additionally including human supervision, in order to target the learning toward artifacts that are informative to human users.

Overall, we expand and combine this previous work by proposing a novel network that both attends to artifacts to increase machine detection, and amplifies them to increase human detection.

2.2 Motion amplification

We amplify artifacts using a module that draws on motion magnification techniques. Motion magnification consists of applying transformations to a given video to expose small movements and perturbations. Liu et al. [26] were the first to introduce this idea: their Lagrangian method estimates small motions and warps frames to expose them. Wadhwa et al. [51] developed a phase-based method that reduced noise and artifacts compared to previous approaches. Oh et al. [39] introduced a fully learned approach to amplify motion, based on a neural network that disentangles a frame’s shape and texture before manipulation. To the best of our knowledge, we are the first to use a technique inspired from this work to amplify artifacts in fake videos.

2.3 Face caricatures

Humans are exceptionally sensitive to the proportions of faces. Psychology research has shown that faces are encoded in memory based on their deviations from a generic averaged face [8, 44], and that the proportions of facial features are at least as important as the particular shape of a facial component for distinguishing among faces [8]. This sensitivity is leveraged in the art style known as "caricature", where distinctive features of a face are exaggerated in a way that makes them easier to recognize and remember [8, 32, 44, 50], by drawing attention to the facial regions that differ most from the norm. Inspired by these caricatures, our method makes distortions of proportions in fake video more visible, increasing a viewer’s ability to recognize the video as fake.

3 Model specification

We present a framework that combines self attention and human-guided attention maps to both detect deepfakes, and expose them to the human eye. By learning to attend to the most diagnostic regions, our model can detect fake videos, and by magnifying motion in these regions, it can generate distorted videos (called *deepfake caricatures*) that are easier for humans to perceive as fake. Our main model, dubbed *CariNet*, contains three main modules:

- An **Artifact Attention Module** that outputs heatmaps indicating the probable location of artifacts in each input frame.
- A **Classifier Module**, consisting of a set of ResNet blocks which uses the output of the artifact attention module to weight self attention.
- A **Caricature Module** that amplifies artifacts and unnatural motions directly in the videos, with the guidance provided by the attention maps.

3.1 Artifact Attention module

This module aims to incorporate human perceptual behavior into the detection pipeline and to guide the module towards informative regions which may not be locally connected. We also incorporated human supervision since one of the goals of this framework was to increase artifact detectability to human users.

Human-informed ground truth. We collected heat maps corresponding to the locations of artifacts detected by humans. First, we used a simple selection interface to create a pool of videos which are difficult to spot as fake, as these are the most likely to yield non-trivial information about subtle artifacts. From the DFDCp dataset [14], we selected 500 videos which were difficult for humans to detect as fake (see supplemental materials), and 500 additional videos that were difficult for the XceptionNet [41] detection model. Next, we showed the selected deepfakes to a new set of participants, who annotated areas of the videos that appeared distorted or manipulated (Figure 2). Participants were shown blocks of 100 3-second clips, and used a paint-brush interface to paint on videos regions that appeared unnatural (see supp. for data analysis). This resulted in over 11K annotations across our 1000 videos. On average, each video had 22.6 sets of annotations, resulting in 4.1K sample points per video. For each video clip, we aggregate all annotation data and generate one 3D attention map. An anisotropic Gaussian kernel of size (20, 20, 6) in x , y and $time$ dimensions was first applied. We then normalize the attention map of each time frame to sum to one.

Module details and training. Our module is based on the encoder-decoder architectural paradigm, and consists of an Xception-based encoder [11] and a 6-block resnet-based decoder, where upsampling and convolutions are layered between each block (Figure 4). The encoder produces codes e that retain spatial location information, and the decoder utilizes those compressed feature maps to generate output heatmaps. We use the human data to supervise the generation of these heatmaps directly. This happens through three losses: the Pearson Correlation Coefficient, KL-Divergence and L-1 loss.

3.2 Classifier module

Our Classifier module (Figure 3) attempts to detect if the input is real or fake, and draws from the information provided by the Artifact Attention Module to attend to key parts of the input. It is composed of a convolutional stem and a main trunk of Attention Blocks: we define these blocks as a Residual Block followed by a modified self-attention operation that is modulated by the artifact attention outputs. Importantly, this module receives feature maps generated by a global convolutional stem of three 3×3 convolution layers with strides 2, 1 and 1. The layers use the Mish [33] activation function and Batch Normalization layers. A 3×3 max pooling layer with stride 2 follows the stem, and a sequence of L Attention Blocks followed by a global average pooling and a sigmoid activation composes the rest of the network. We consider two types of Residual blocks in our architecture: the common Residual block of two 3×3 convolution layers, and the bottleneck block of three convolution layers with kernel size of 1×1 , 3×3 and 1×1 . We analyzed Attention block sequence sizes of $L = 18$ and $L = 34$: each alternative yields a different model version, referred to as CariNet18 and CariNet34. Cross entropy is used as the loss function. The output of the binary logit function is assigned to each frame as a detection score; we take the averaged score of the whole sequence as the prediction for the video clip. Our modules are trained jointly, and the full loss is the combination of this cross-entropy loss with the heatmap losses from the Artifact Attention Module.

Self-attention with artifact heatmaps. We define our self-attention layers in a similar manner to prior self-attention work [53, 7], but extend the traditional construction to incorporate modulation from the artifact attention heatmaps. Specifically, our self attention layer computes an affinity matrix between keys and queries where the keys are re-weighted by the artifact attention map. The rationale here is that if human attention maps point to valuable areas in the frame, we want to maximize query affinity in those areas and inhibit it in others. Given a feature map \mathbf{x}_i over one frame, and an artifact attention map A over that frame, the module learns to generate an affinity matrix \mathbf{a}_i :

$$\mathbf{a}_i = \text{softmax}((\mathbf{W}_Q \mathbf{x}_i)^T (\mathbf{W}_K \mathbf{x}_i \odot \mathbf{A})). \quad (1)$$

The softmaxed key-query affinity tensor is then matrix-multiplied with the values $V = W_V x_i$ to generate the output residual r . That residual is then scaled by γ and added to input x_i to yield the output feature map y_i :

$$y_i = \gamma \mathbf{a}_i^T (\mathbf{W}_V \mathbf{x}_i) + \mathbf{x}_i. \quad (2)$$

$\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$ are learned weight matrices of shape $\mathbf{R}^{\bar{C} \times C}$, $\mathbf{R}^{\bar{C} \times C}$ and $\mathbf{R}^{C \times C}$ respectively, with $\bar{C} = C/4$. γ is a learnable scalar which controls the impact of the learned attention feature vis-a-vis the original feature maps x_i .

3.3 Caricature Generation Module

Our caricature module is inspired by the motion magnification network proposed in [39]. The encodings e generated by the artifact attention module are used to amplify differences between the representations of consecutive frames. The artifact attention maps guide this module: element-wise multiplication with the artifact attention map of frame x_i is utilized to re-weight the difference in codes for each pair of frames x_i and x_{i+1} . The distorted code \hat{e}_i of frame x_i is computed as

$$\hat{e}_i + \mathbf{1} = \mathbf{e}_i + \alpha(\mathbf{e}_i - \mathbf{e}_{i+1}) \odot \mathbf{A}, \quad (3)$$

where α is a user-defined distortion factor that controls the strength of the resulting caricature.

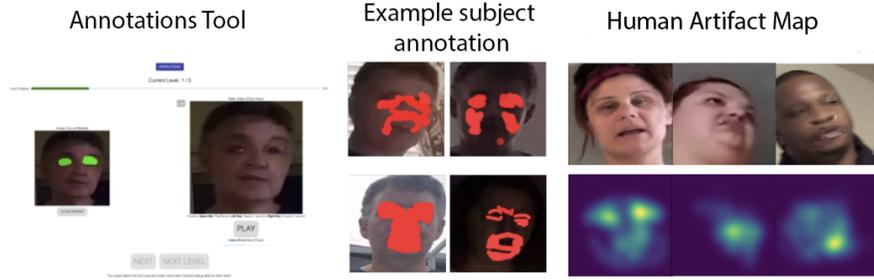


Figure 2: Left Panel: Our annotation interface allows users to paint over particular artifacts or zones that appear fake. Our system tracks both the position and frame at which an annotation occurs. Users highlighted both large areas and more specific, semantically meaningful areas (e.g. eyebrows). Middle panel: example of single subject’s annotations. Right panel: Examples of artifact maps created from averaging human annotations. Humans consistently highlighted areas with abnormal artifacts.

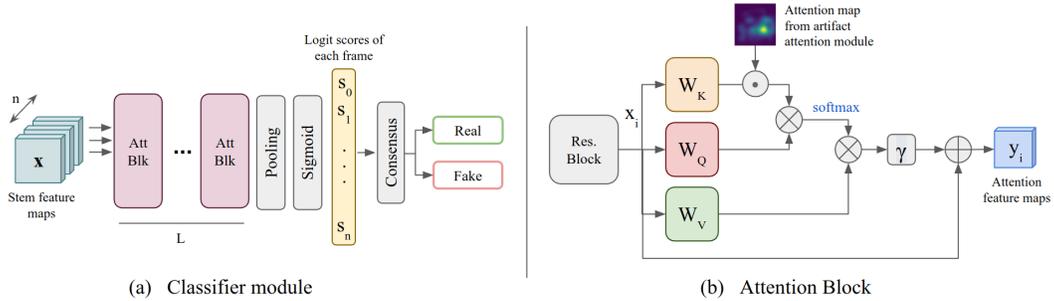


Figure 3: **The classifier module and its attention blocks.** (a) our classifier takes the feature maps outputted by our convolutional stem, passes them through attention blocks modulated by human heatmaps, and computes logit scores for each frame before classifying the video. The consensus operation is instantiated as an average of the logits followed by thresholding to determine the output label. (b) Our attention block: a traditional residual block is followed by key, query and value matrices following the self-attention framework. The key-query product is modulated by our human heatmaps. \times represents matrix multiplication and \cdot represents element-wise multiplication.

4 Model Experiments

4.1 Setup

Datasets. We evaluate models on four benchmarks: First, the *FaceForensics++* (*FF++*) [41]. With a large range of manipulations, *FaceForensics++* is a strong benchmark used by several previous approaches. The dataset contains 1000 real videos and 4000 fake videos generated from four different facial manipulation methods: (Deepfake [1], Face2Face [47], FaceSwap [2] and NeuralTextures [46]). To ensure fair comparisons with [18] and others, we use the first 270 frames of each training video and the first 110 frames for each validation video. Next, we test on the *The Preview Deepfake Detection Challenge Dataset (DFDCp)* [14]. We use the "preview" version of this dataset, containing 5214 videos from 68 different subjects. This dataset features 3754 fakes and 1460 real videos generated using two different facial modification methods. Additionally, we use *Celeb-DF v2* [25]. This dataset contains high quality fake videos of celebrities generated from real Youtube videos. There are 590 original videos and 5639 corresponding deepfakes. Finally, we also use the *DeeperForensics (DFo)* [25]. A large scale dataset, featuring 10000 fake videos obtained by applying new manipulations and augmentations to the 1000 real videos from *FaceForensics++*, mainly swapping in faces of 100 new paid actors.

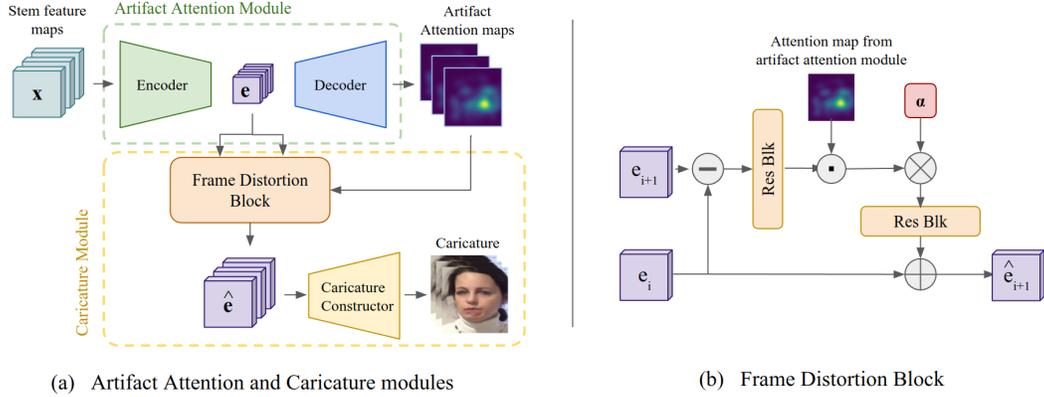


Figure 4: **Our artifact attention and caricature modules.** (a) The artifact attention and caricature modules are intertwined: artifact attention operates with an encoder-decoder architecture to generate artifact heatmaps. Those heatmaps are supervised with the human heatmaps collected through our annotation interface. The caricature module receives both the heatmaps and the internal codes e , distorts those codes according to the artifact attention heatmaps, and generates caricatures by reconstructing the video from the distorted codes. (b) Our frame distortion block computes the difference between codes e_i and e_{i+1} , re-weights it according to the artifact attention maps, and then amplifies it by a factor of α before summing it back to e_i to generate distorted code \hat{e}_{i+1} .

In order to standardize faces across datasets, video were preprocessed so that each video was 360×360 pixels, and showed a single face, with a minimum size of 50 pixels and a minimum margin of 100 pixels from the edge of the frame (see supplement for details).

Baseline comparisons We compare CariNet to several baselines from the literature. **(1)** First we consider the *XceptionNet* model, which was trained on the FaceForensics++ dataset [41]. **(2)** We also tested a *Frame-based ResNet18 and ResNet34* [20], which we use as a more general yet simple baseline for frame-based detection. **(3)** Additionally, we use a *Clip-based ResNet* which uses 3D kernels inflated from ResNet18 pretrained on ImageNet [13] to learn over multiple frames at once [10, 19, 49], as well as a version pretrained on the Kinetics dataset [9]. We report results over our best ResNet: a 3DResnet-34 pretrained on Kinetics. **(4)** We also use the *Face X-ray model*, a model that attends to the seams in face-swapping to make its predictions. We use existing performance values from [18] for the FF++ trained results, and evaluate an HRNet-W48 trained on synthetic blended images in conjunction with FF++ videos for our DFDCp results. **(5)** Additionally, we use the *CNN-GRU* from [42], where we train a DenseNet-161 followed by a GRU on FF++ to evaluate on DFDCp, and report existing performance values for FF++. **(6)** We also use *DSP-FWA*, a model that tries to detect face warping artifacts. We evaluate the publicly available pretrained model for FF++, and train a ResNet50 with the pyramidal setup introduced in [24] for DFDCp. **(7)** We also test the *Multi-task* model from [35]. We report existing numbers from [18] and retrain using the publicly available implementation. **(8)** We also consider the *Two-branch* model from [30]. We report existing numbers for CelebDFv2 evaluation, and retrain on the rest. **(9)** Finally, we test the *Multi-attention* model from [55]. We report the values from the existing paper, and retrain following the setup in the paper on FF++ for DFDCp results.

Learning and Optimization. Our models were implemented in PyTorch. Our CariNets were optimized with Rectified Adam [27] with LookAhead [54], an extension of the Adam optimizer informally referred to as the Ranger optimizer. We use a batch size of 32 and an initial learning rate of 0.001. Cosine annealing learning rate scheduling was applied with half period of 100 epochs. We chose an early stopping strategy to terminate training, stopping if validation accuracy stagnates in a 10 epoch window. We apply flipping (probability 0.5) and random cropping (224×224) augmentations.

4.2 Deepfake Detection Performance

Generalization to unseen datasets. At present, the performance of deepfake classifiers are assessed primarily based on their ability to generalize well to datasets built with different techniques from

Model	CelebDFv2	DFDCp	FShifter	DFo	Overall
3DResNet-34 (Kinetics)	72.1	64.9	72.2	74.5	70.9
Xception [41]	73.7	65.7	72.0	84.5	75.3
Face X-ray [23]	79.5	62.1	92.8	86.8	81.2
CNN-GRU [42]	69.8	63.7	80.8	74.1	73.4
Multi-task [35]	75.7	63.9	66.0	77.7	71.9
DSP-FWA [24]	69.5	64.5	65.5	50.2	63.1
Two-branch [30]	73.4	64.0	-	-	-
Multi-attention [55]	67.4	67.1	-	-	-
LipForensics [55]	82.4	70.0	97.1	97.6	86.8
CariNet18 (ours)	80.1	72.9	90.4	94.3	84.4
CariNet34 (ours)	83.2	74.4	93.9	97.9	87.4

Table 1: **Detection performance results on unseen datasets.** We report Video-level AUC (%) on four tested benchmarks. All models are pretrained on FF++ (all manipulations), except CariNet18 and CariNet34, which are also trained with human artifact maps from DFDCp videos. We only report the best performing ResNet baseline (3DResnet-34).

Method	Train on remaining			
	DF	F2F	NT	Avg
Xception [41]	93.9	86.8	79.7	77.9
CNN-GRU [42]	97.6	85.8	86.6	79.4
Face X-ray [23]	99.5	94.5	92.5	94.9
LipForensics	99.7	99.7	99.1	99.5
CariNet34 (ours)	99.7	99.8	97.5	99.0

Table 2: **Cross-forgery performance.** Video-level AUC (%) on three of the forgery types of FaceForensics++ (Deepfakes (DF), Face2Face (F2F) and Neural Textures (NT)). As CariNet uses DFDC data with a manipulation most similar to FaceSwap during training, we exclude the FS manipulation from the set.

what the classifier was trained on. Thus, following the work of [18], we report our results as our model’s ability to train on one dataset and perform well on another. We train models on FF++ and evaluate their performance on CelebDFv2, DFDCp, FaceShifter and DeeperForensics. Our CariNets show strong performance in this task, surpassing all previous alternatives in most benchmarks except for FaceShifter. We hypothesize that the attentional framework proposed allows CariNet to build more robust representation of artifact location and properties, focusing as needed on lips, eyes or other telling features. (Table 1).

Generalization to unseen manipulations. To measure our model’s ability to perform well on unseen manipulations, we train on a subset of the data containing three of the FF++ manipulations as well as our DFDC videos with human maps, and evaluate on the remaining unseen manipulations. We consider that DFDC contains face-swap manipulations only, so we report performance exclusively on the Deepfake manipulation (DF), Face2Face, and NeuralTextures. We compare this to alternative models from the literature and show the results in Table 2. Overall, our technique is competitive with the best-performing model.

Robustness to unseen perturbations Another key aspect of a forgery detector is the ability to natively adapt to the presence of new, unseen perturbations. Thus we analyse the robustness of our model’s detection capabilities in the presence of unseen perturbations to the input deepfakes. To analyze CariNet’s behavior across types of perturbation, we test our model trained with FF++ on test videos from FF++ with 4 perturbations: Saturation, Contrast, Gaussian Noise, Gaussian blur and Pixelation. We follow the framework of [18] and apply perturbations at 5 severity levels for our comparisons. We report average performance over the 5 severity levels in Table 3. We observe that our method outperforms previous approaches at most severity levels. We hypothesize that our human-guided attention framework might be of help in this setting, as humans are naturally capable of adapting to different lighting conditions, blurs, resolutions and other photometric modifications. This adaptability might be captured in the ground truth maps that guide the learning process of our artifact attention module.

4.3 Ablation studies

Role of the artifact attention module. We confirmed the contribution of our artifact attention module by performing ablation studies across the different datasets under study. We observe performance

Method	Clean	Saturation	Contrast	Noise	Blur	Pixel
Xception	99.8	99.3	98.6	53.8	60.2	74.2
Face X-ray	99.8	97.6	88.5	49.8	63.8	88.6
CNN-GRU	99.9	99.0	98.8	47.9	71.5	86.5
LipForensics	99.9	99.9	99.6	73.8	96.1	95.6
CariNet34 (ours)	99.9	99.9	99.8	74.5	91.8	96.3

Table 3: **Generalization performance over unseen perturbations.** Video-level AUC (%) on FF++ over videos perturbed with 5 different modifications. We report averages across severity levels.

Model	DFDCp	FF++	CelebDFv2	DFo	Overall
CariNet18 w/o attention blocks	70.92	93.73	73.9	84.56	80.78
CariNet18 w/o attention maps	72.34	94.82	76.5	91.21	83.72
CariNet18 w/ fixed attention	68.15	90.15	71.1	82.23	77.91
CariNet18 (ours)	72.90	96.81	80.1	94.33	86.04

Table 4: **Ablation study results.** We show how certain components of our approach affect video-level AUC (%). We replace attention blocks with regular residual blocks, we remove attention maps altogether, and we replace our attention maps with a fixed center bias. All modifications yield lower performance than our proposed network.

decreases following ablation in two different scenarios: 1) removing our artifact attention module from the classification pipeline, which yields regular attention blocks with no key modulation, and 2) replacing artifact attention with a fixed center bias, operationalized as a 2D gaussian kernel with mean $\mu = (W/2, H/2)$ and standard deviation $\sigma = (20, 20)$. We show our results in Table 4. Interestingly, adding a fixed attention map degrades performance: this forces attention blocks to attend to parts of the input that don't necessary contain valuable information related to the input's manipulation.

Role of attention blocks. We also observe the performance of our classification module when replacing our attention blocks with simple 3D residual blocks in CariNet18. The residual blocks contain two 3x3x3 convolutions followed by batch normalization. This classification module receives no information from our artifact attention module. As we can see in Table 4, the performance of our model is greatly improved when our attention blocks are used.

5 Human Experiments

Next, we tested whether the caricature technique improves human detection of deepfakes. In a sample of 400 videos from the DFDCp (200 fake, 200 real), we compared the detectability of *standard* (i.e. unmodified) deepfakes with deepfakes which have passed through the Caricature Generation Module, which we will call *caricature* deepfakes. Videos were shown to human participants one at a time, and participants indicated if the video was real or fake. Crucially, we tested whether caricatures could improve deepfake detection even under challenging conditions such as very fast presentation times or low attentiveness. Videos were presented for 6 different durations: 300ms, 500ms, 1000ms, 3000ms, 5000ms. See Supplement for experiment design details.

In figure 5A, we show the detectability of fake videos in each condition (green for caricature, orange for standard deepfake). For each video, we calculated the proportion of time it was correctly detected in a sample of 10 participants. Detectability was substantially higher for caricatures than standard deepfakes, averaged over all timepoints (main effect of condition, $F(2,1630)=17.12$, $p<0.001$). Figure 5B plots the difference between the detectability of caricature videos and standard deepfakes for each timepoint. The advantage for caricature over standard deepfakes was present at every timepoint tested (statistically significant at a Bonferroni-corrected alpha of $p=0.0083$ for all timepoints). With as little as 300ms exposure, detection was better for caricatures by 14 percentage points, and this advantage increases to 43 percentage points for a 5 second exposure (significant interaction, $F(10,1630)=8.88$, $p<0.001$). Crucially, while standard deepfakes were detected no better than chance, caricatures boosted the likelihood of detection above 50% across all presentation times. Thus, caricatures increase the likelihood that a deepfake will be detected as fake, even when participants had less than a second of exposure. See Supplement for a full statistical reporting.

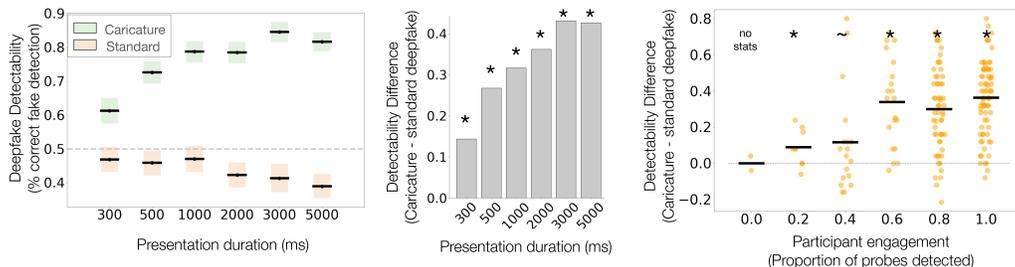


Figure 5: Behavioral results showing the effect of caricatures on human deepfake detection. A) deepfake detectability by humans for standard (orange) and caricature (green) conditions. Colored boxes represent 95% confidence intervals. B) Replotting of the difference between standard and caricatures from pane A. Caricatures are more detectable, even at extremely short presentation times. C) Caricatures are more detectable than standard deepfakes even for distracted participants.

In Figure 5C, we show how participant engagement impacts the effectiveness of the caricature technique. For artifact magnification to be a viable method of signaling deepfakes to humans, it should be effective even when users are distracted or inattentive. Unbeknownst to participants, we embedded *engagement probes* in the experiment (5 trials per 100-trial HIT), consisting of standard deepfakes with very large artifacts which were extremely easy to detect. We reasoned that highly-engaged participants would succeed on all of these trials, but medium to low-engagement participants would miss some proportion of them. We binned participants based on the number of engagement probes they correctly identified as fake, and examined the difference between caricatures and standard fakes for each bin. We find that caricatures yield higher detection at all but the lowest level of engagement (no stats were conducted for the lowest engagement bin, because it contained too few participants; all other comparisons were significant at $p < 0.05$, or marginal at $p = 0.051$ in one case).

Overall, these behavioral results demonstrate that the caricature method is extremely effective at signalling to a human user that a video is fake. Not only does it allow for the correct detection of deepfakes from the first glimpse (as quickly as 300 ms), but it also protects people even when they are not being attentive and vigilant.

6 Conclusion

The goal of this work is to increase trust in video-based media, by developing an automated procedure that detects manipulated videos and exposes them to human observers from the first glance. Our CariNet shows excellent performance on four different datasets, and crucially, creates novel Deepfake Caricatures which allow for above-chance detection by human observers as early as 300ms after video onset.

This work establishes the importance of integrating computer vision and human factors solutions for deepfake mitigation, and demonstrates the feasibility of improving human performance with machine assistance. More broadly, these results point to the power of using motion signals to alert human users that a video is fake. As decades of human vision research has shown, motion is quickly and automatically detectable for people [38, 4, 29], and thus may provide a more powerful signal than simply placing a authenticity warning near a questionable video.

One weakness of the current work is that deepfakes are constantly and rapidly improving. As such, our detector architecture will need to be re-trained on newer methods as they emerge to remain useful. We note, however, that the broader effectiveness of magnifying artifacts to increase detectability by human users is constant, and independent of how this effect is achieved.

As with any misinformation detection system, there is risk that our network could be reverse engineered to produce higher quality deepfakes. However, a system which allows humans to directly detect if a video is doctored will empower more individuals to assess for themselves whether to trust and engage with a given video. Aggregated over the many millions of people who watch videos online everyday, we believe that the benefits of such a system outweigh the risks.

References

- [1] Deepfake. <https://github.com/deepfakes/faceswap>, 2020.
- [2] Faceswap. <https://github.com/MarekKowalski/FaceSwap/>, 2020.
- [3] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen. Mesonet: a compact facial video forgery detection network. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE, 2018.
- [4] T. D. Albright and G. R. Stoner. Visual motion perception. *Proceedings of the National Academy of Sciences*, 92(7):2433–2440, 1995.
- [5] B. Bayar and M. C. Stamm. A deep learning approach to universal image manipulation detection using a new convolutional layer. In *Proceedings of the 4th ACM Workshop on Information Hiding and Multimedia Security*, pages 5–10, 2016.
- [6] I. M. Begg, A. Anas, and S. Farinacci. Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121(4):446, 1992.
- [7] I. Bello, B. Zoph, A. Vaswani, J. Shlens, and Q. V. Le. Attention augmented convolutional networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3286–3295, 2019.
- [8] P. J. Benson and D. I. Perrett. Perception and recognition of photographic quality facial caricatures: Implications for the recognition of natural images. *European Journal of Cognitive Psychology*, 3(1):105–135, 1991.
- [9] J. Carreira, E. Noland, C. Hillier, and A. Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019.
- [10] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [11] F. Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.
- [12] U. A. Ciftci and I. Demir. Fakecatcher: Detection of synthetic portrait videos using biological signals. *arXiv preprint arXiv:1901.02212*, 2019.
- [13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [14] B. Dolhansky, R. Howes, B. Pflaum, N. Baram, and C. C. Ferrer. The deepfake detection challenge (dfdc) preview dataset. *arXiv preprint arXiv:1910.08854*, 2019.
- [15] R. Durall, M. Keuper, F.-J. Pfreundt, and J. Keuper. Unmasking deepfakes with simple features. *arXiv preprint arXiv:1911.00686*, 2019.
- [16] M. Groh, Z. Epstein, C. Firestone, and R. Picard. Deepfake detection by human crowds, machines, and machine-informed crowds. *Proceedings of the National Academy of Sciences*, 119(1), 2022.
- [17] D. Güera and E. J. Delp. Deepfake video detection using recurrent neural networks. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018.
- [18] A. Haliassos, K. Vougioukas, S. Petridis, and M. Pantic. Lips don’t lie: A generalisable and robust approach to face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5039–5049, 2021.
- [19] K. Hara, H. Kataoka, and Y. Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [21] H. Li, B. Li, S. Tan, and J. Huang. Identification of deep network generated images using disparities in color components. *Signal Processing*, 174:107616, 2020.
- [22] J. Li, T. Shen, W. Zhang, H. Ren, D. Zeng, and T. Mei. Zooming into face forensics: A pixel-level analysis. *arXiv preprint arXiv:1912.05790*, 2019.
- [23] L. Li, J. Bao, T. Zhang, H. Yang, D. Chen, F. Wen, and B. Guo. Face x-ray for more general face forgery detection. *arXiv preprint arXiv:1912.13458*, 2019.
- [24] Y. Li and S. Lyu. Exposing deepfake videos by detecting face warping artifacts, 2018.
- [25] Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu. Celeb-df: A new dataset for deepfake forensics. *arXiv preprint arXiv:1909.12962*, 2019.
- [26] C. Liu, A. Torralba, W. T. Freeman, F. Durand, and E. H. Adelson. Motion magnification. *ACM transactions on graphics (TOG)*, 24(3):519–526, 2005.
- [27] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [28] M. E. Lloyd, D. L. Westerman, and J. K. Miller. The fluency heuristic in recognition memory: The effect of repetition. *Journal of Memory and Language*, 48(3):603–614, 2003.
- [29] Z.-L. Lu and G. Sperling. Three-systems theory of human visual motion perception: review and update. *JOSA A*, 18(9):2331–2370, 2001.
- [30] I. Masi, A. Killekar, R. M. Mascarenhas, S. P. Gurudatt, and W. AbdAlmageed. Two-branch recurrent network for isolating deepfakes in videos. In *European Conference on Computer Vision*, pages 667–684. Springer, 2020.

- [31] F. Matern, C. Riess, and M. Stamminger. Exploiting visual artifacts to expose deepfakes and face manipulations. In *2019 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, pages 83–92. IEEE, 2019.
- [32] R. Mauro and M. Kubovy. Caricature and face recognition. *Memory & Cognition*, 20(4):433–440, 1992.
- [33] D. Misra. Mish: A self regularized non-monotonic neural activation function. *arXiv preprint arXiv:1908.08681*, 2019.
- [34] D. M. Montserrat, H. Hao, S. Yarlagadda, S. Baireddy, R. Shao, J. Horváth, E. Bartusiak, J. Yang, D. Güera, F. Zhu, et al. Deepfakes detection with automatic face weighting. *arXiv preprint arXiv:2004.12027*, 2020.
- [35] H. H. Nguyen, F. Fang, J. Yamagishi, and I. Echizen. Multi-task learning for detecting and segmenting manipulated facial images and videos. *arXiv preprint arXiv:1906.06876*, 2019.
- [36] H. H. Nguyen, J. Yamagishi, and I. Echizen. Capsule-forensics: Using capsule networks to detect forged images and videos. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2307–2311. IEEE, 2019.
- [37] S. J. Nightingale and H. Farid. Ai-synthesized faces are indistinguishable from real faces and more trustworthy. *Proceedings of the National Academy of Sciences*, 119(8):e2120481119, 2022.
- [38] S. Nishida, T. Kawabe, M. Sawayama, and T. Fukiage. Motion perception: From detection to interpretation. *Annual review of vision science*, 4:501–523, 2018.
- [39] T.-H. Oh, R. Jaroensri, C. Kim, M. Elgharib, F. Durand, W. T. Freeman, and W. Matusik. Learning-based video motion magnification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 633–648, 2018.
- [40] G. Pennycook, T. D. Cannon, and D. G. Rand. Prior exposure increases perceived accuracy of fake news. *Journal of experimental psychology: general*, 147(12):1865, 2018.
- [41] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1–11, 2019.
- [42] E. Sabir, J. Cheng, A. Jaiswal, W. AbdAlmageed, I. Masi, and P. Natarajan. Recurrent convolutional strategies for face manipulation detection in videos. *Interfaces (GUI)*, 3:1.
- [43] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [44] P. Sinha, B. Balas, Y. Ostrovsky, and R. Russell. Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11):1948–1962, 2006.
- [45] J. Stehouwer, H. Dang, F. Liu, X. Liu, and A. Jain. On the detection of digital face manipulation. *arXiv preprint arXiv:1910.01717*, 2019.
- [46] J. Thies, M. Zollhöfer, and M. Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [47] J. Thies, M. Zollhofer, M. Stamminger, C. Theobalt, and M. Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016.
- [48] R. Tolosana, S. Romero-Tapiador, J. Fierrez, and R. Vera-Rodriguez. Deepfakes evolution: Analysis of facial regions and fake detection performance, 2020.
- [49] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [50] B. Tversky and D. Baratz. Memory for faces: Are caricatures better than photographs? *Memory & cognition*, 13(1):45–49, 1985.
- [51] N. Wadhwa, M. Rubinstein, F. Durand, and W. T. Freeman. Phase-based video motion processing. *ACM Trans. Graph. (Proceedings SIGGRAPH 2013)*, 32(4), 2013.
- [52] X. Yang, Y. Li, and S. Lyu. Exposing deep fakes using inconsistent head poses. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8261–8265. IEEE, 2019.
- [53] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*, 2018.
- [54] M. R. Zhang, J. Lucas, G. Hinton, and J. Ba. Lookahead optimizer: k steps forward, 1 step back. *arXiv preprint arXiv:1907.08610*, 2019.
- [55] H. Zhao, W. Zhou, D. Chen, T. Wei, W. Zhang, and N. Yu. Multi-attentional deepfake detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2185–2194, 2021.

Supplemental Materials

1 Caricatures in Video Form

The effect of our caricatures is best experienced in video form. We provide a compilation of examples in the supplementary mp4 files *caricature_applied_on_real_vids.mp4* and *caricature_applied_on_fake_vids.mp4*. The first one corresponds to the effect of applying our caricature model to real videos. Little to no impact on the underlying video can be seen. On the second movie, we show the caricature impact on deepfakes. The artifacts are amplified and the video is generally distorted. We show CariNet34 caricatures in both cases.

We additionally provide an anonymized web gallery, *gallery.html*, that allows for easy visualization of several additional examples. Once the *gallery.html* file is extracted from the provided zip, along with its accompanying videos, double click on the html file to view it on your browser.

2 Video Selection for Human Artifact Annotations

Selecting hard fakes according to humans. Selecting target videos for our human artifact annotations poses a challenge: we would like videos to be hard to detect as fakes by humans, to ensure that our artifacts are a) non-trivial, and b) capture the subtleties of well-functioning face alteration techniques. To facilitate this selection, we formulated a crowd sourced experiment to find fake videos that are relatively difficult to detect as fake, without scrutiny. The difficult videos would serve as good candidates to test the efficacy of our caricature method.

We randomly selected 2000 videos (1000 real, 1000 fake) from the Deepfake Detection Challenge Dataset [14] and used FaceNet [43] to crop the face regions, resulting in videos of 360×360 pixels. 1000 video pairs were divided into five sessions, and were presented in a randomized order. One task consisted of 100 video pairs divided in 20 different levels.

In our task, participants were instructed to maintain their fixation at a center cross while a pair of videos were presented, one on the left and one on the right (Figure 6). One of those videos was fake, while the other one was real. Participants were asked to select which video, left or right, was more likely to be fake by pressing a left or right keyboard key.

At the start of the experiment, people had the possibility to adjust the distance between the videos (left and right videos were always at the same distance from the center). To help with the adjustment, the following instructions were given "The videos should not be too close to disturb you from fixating at the center, and not too far away to notably reduce your ability to distinguish them. Find the distance that feels more comfortable to you." After every 10 videos shown to users, users took a quick break and then had 3 seconds to refocus their attention on the fixation cross before examining the next set of videos. Whenever a gaze shifted towards either side of the videos, participants were asked to report and the trial was discarded.

We obtained 10 responses per video from 522 individual users. We computed accuracy by averaging the responses over videos. Based on those responses, we selected 500 fake videos that were closest to having a 50% recognition rate, which indicated that human detection performance was close to chance. Those videos were considered "hard to detect".

Selecting hard fakes according to models. We would additionally like to represent the distribution of hard fakes for models, and collect human artifact maps on these fakes to ensure. For this, we selected 500 fake videos from our DFDCp training split that minimized the "fake" classification confidence (output of the softmax) of an Xception network trained on FF++, while not being a duplicate of the 500 videos selected through our human experiment.

3 Video preprocessing

All datasets were downloaded at full resolution, except FaceForensics++, which was downloaded in the provided *c23* compressed format to avoid large storage costs. To streamline our training pipeline, we performed multi-person face extraction on all videos using a pretrained PyTorch implementation of FaceNet[43], with a minimum face size of 50 pixels and a face margin size of 100 pixels. For

DeepFake Detection

INSTRUCTIONS

Current Level: 2 / 21

Level Progress: 20%

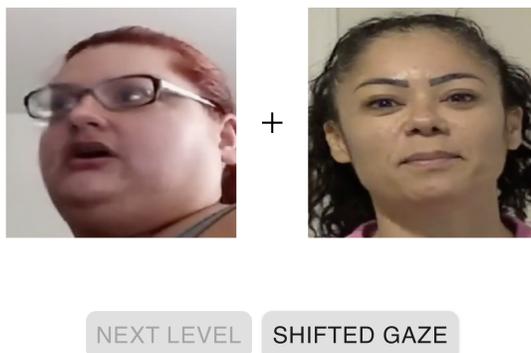


Figure 6: Screenshot of the interface used to collect human detection difficulty. Humans were tasked with indicating which of the two videos was fake without shifting their gaze from the cross.

frames without a detected face, bounding box coordinates were produced via linear interpolation of the next available neighboring bounding boxes. To avoid abrupt motion artifacts (e.g. jittering) from face extraction, the 4 corner coordinates of the detected bounding boxes were temporally smoothed using hanning smoothing with a window size equal to 20% the video length. These face frames were resized to 360×360 pixels and saved as an MP4 file.

4 Video annotations details and results:

4.1 Task design

Videos were annotated in an online task that used a paintbrush interface. Participants were instructed to paint over the regions that were most informative about whether a video was fake. Videos were paused during painting, and the interface contained a separate visualization of where the participant has painted, so they could keep track of their annotations.

4.2 Qualitative results

In total we collected over 11K annotations for 1000 videos. On average, each video had annotations from 22.6 participants, resulting in 4.1K sample points per video. For each video clip, we aggregated all annotation data and generated one 3D attention map: An anisotropic Gaussian kernel of size $(cs, cs, 6)$ in x,y and time dimensions was first applied. c is set to a constant of 20 and s corresponds to the brush size. We normalize the attention map of each time frame to sum to one.

Humans tend to annotate areas of the face that are near the eyes or mouth, which are more salient facial features. Our model picks up on these features, and more fake indication areas. For example, there is attention to inflection regions of the face, such as the jawline or near the hairline. Humans also gravitate towards inconsistencies around the seams generated by the face swapping pipeline, as well as flickering artifacts generated by tracking errors during the deepfake process.

5 Training details

Our CariNets were trained on several datasets for our experiments: the DeepFake Detection Challenge (DFDCp) dataset, FaceForensics++, CelebDFv2 and DeeperForensics. For all datasets, we train on both the pristine and deepfake videos from the train split of each dataset. Evaluation happens on the

validation or test set of each dataset. We oversample real frames during training to achieve real/fake balance. We randomly sample 32 frames from videos during training to feed our model.

As mentioned in the main text, our optimization parameters are as follows: batch size of 32, learning rate of 0.001 with cosine annealing, Ranger21 optimizer, and early stopping based on validation accuracy stagnation. These hyperparameters were selected through a simple grid search with 5 alternatives, by observing performance after 50 epochs on a CariNet18 trained on DFDCp.

6 Behavioral Experiment - Methods and full statistics reporting

6.1 Stimuli

Stimuli consisted of 400 videos selected from DFDCp [14]. Half of the videos contained real, unaltered faces, while the other half contained deepfake faces. Videos were randomly selected, then deepfakes with obvious artifacts were discarded. All videos were subjected to the pre-processing and cropping procedure described above, and were presented at a resolution of 300x300 pixels. These 200 fake videos made up the stimuli for the "standard deepfake" condition. These videos were then passed through the Caricature Creation Module to create stimuli for the "Caricature" condition.

6.2 Behavioral Task

Participants performed a two-alternative forced choice detection task. They were shown one video at a time and were asked to indicate if the video was real or fake by using the mouse to click on buttons labeled "Real" and "Fake". Prior to the tasks, participants were shown instructions which included 3 example deepfakes, to inform them about what a fake video looks like.

Each task session contained 100 videos, divided into 5 blocks of 20 videos with breaks in between. Videos could be presented at 6 different presentation durations: 300ms, 500ms, 1000ms, 3000ms and 5000ms. Each trial started with a 3 second countdown. In addition, each block contained one compliance check, and one engagement probe. Compliance checks consisted of real or fake videos (not included in the experimental set) with the text "THIS IS AN ATTENTION CHECK, PLEASE CLICK REAL" written across the video in light grey font. These trials were included to provide an objective basis by which to exclude low-performing participants (see Analysis section below). Engagement probes consisted of deepfakes (not included in the experimental set) containing extremely obvious artifacts, such as large patches of noise covering the face. These trials were included to provide an online measure of participants attentiveness. We reasoned that while highly-engaged participants would succeed on all of these trials, medium to low-engagement participants would miss some proportion of them.

Trials were divided into experimental sessions (or "HITs") of 100 trials. Trials were assigned to HITs pseudo randomly, with the exception that no fake or caricature could be shown twice in the same HIT. Data were collected for 19 HITs of randomly selected videos, leading to an average of 72 videos per presentation time for both standard fakes and caricatures. We collected 10 participants per each HIT.

6.3 Participants

Participants consisted of 170 workers on Amazon Mechanical Turk. Participation was restricted to individuals over the age of eighteen, located in the US, with approval ratings over 99%. Participants were recruited and compensated according to procedure approved by the [University's] Committee on the Use of Humans as Experimental Subjects. Participants were paid \$1 per HIT, amounting to an hourly rate of \$7/hour.

6.4 Analysis

We collected 10 judgments per video per condition. This number was selected after a power analysis (performed on a pilot set of 100 HITs), which suggested that 6 responses was sufficient to obtain stable estimates of the mean detectability of a given video.

Our primary analysis was to examine the difference in correct detection rates between the standard deepfake condition and the caricature condition, both pooling over presentation times and for each

presentation time individually. As such, we pre-registered an ANOVA analysis to test for a main effect of condition on accuracy, with the hypothesis that correct detection would be higher in the caricature than standard deepfake condition. We additionally hypothesized that overall detection (pooling across conditions) would be higher for longer presentation times (main effect of presentation time), and that the difference between standard and caricature deepfakes would increase with higher presentation time (significant interaction), although these analyses were orthogonal to the main question of this experiment. We also planned further comparisons between the standard and caricature conditions for each time point individually. We pre-registered individual unpaired t-tests at a conservative alpha level of $p = 0.0083$ (Bonferroni correction for performing 6 individual ttests, one per time point). To ensure quality data for this analysis, we excluded any sessions where participants used only one key, or failed three or more compliance checks. These trimming and analysis procedures were developed in an initial pilot dataset of 100 HITs, then pre-registered for use in the experimental dataset. We elected to use the conservative Bonferroni correction even though these were a-priori analyses because this experiment had high enough power to detect effects even at conservative thresholds, and applying this threshold would reduce the chance of Type I error.

We also performed a second analysis to examine the effects of condition across levels of participant attentiveness. We divided HITs into participant engagement bins based on the number of engagement probes that were successfully detected throughout the HIT. As the dependent measure, we subtracted the correct detection rate in the standard condition from the caricature condition, to get a measure of the relative benefit of the caricature condition, and tested whether this difference was greater than 0 using a one-sided one-sample t-test at an alpha level of $p = 0.05$. We used a traditional rather than conservative p values because some conditions were likely to contain a low number of trials, and thus have lower power. Statistics were only carried out for conditions containing 10 or more datapoints. This analysis used all experimental sessions.

6.5 Complete statistical reporting

Analysis 1: ANOVA testing for an effect of condition

	SS	df	F	p
Condition	2.171582	4.0	16.457731	3.126305e-13
Presentation time	2.433297	5.0	14.752950	3.503118e-14
Interaction	6.151269	20.0	9.323703	3.978575e-20
Residual	54.791847	1661.0	-	-

Analysis 2: planned comparisons of detectability difference between standard fake and caricatures for each timepoint.

	300ms	500ms	1000ms	2000ms	3000ms	5000ms
df	155	160	159	162	129	172
t	4.58	7.97	9.11	9.45	11.23	12.86
p	9.2e-6	2.9e-13	3.1e-17	3.5e-17	7.7e-21	5.9e-27

Analysis 3: Significance of the difference between standard fake and caricatures for each level of participant engagement.

	0	0.2	0.4	0.6	0.8	1
df	-	9	18	20	58	76
t	-	2.63	1.73	6.27	9.82	15.64
p	-	0.015	0.051	2.5e-06	3.6e-14	1.3e-25