

# SAiW: Source-Attributable Invisible Watermarking for Proactive Deepfake Defense

Bibek Das, *Student Member, IEEE*, Chandranath Adak, *Senior Member, IEEE*, Soumi Chattopadhyay, *Senior Member, IEEE*, Zahid Akhtar, *Senior Member, IEEE*, Soumya Dutta, *Member, IEEE*

**Abstract**—Deepfakes generated by modern generative models pose a serious threat to information integrity, digital identity, and public trust. Existing detection methods are largely reactive, attempting to identify manipulations after they occur and often failing to generalize across evolving generation techniques. This motivates the need for proactive mechanisms that secure media authenticity at the time of creation. In this work, we introduce SAIW, a *Source-Attributable invisible Watermarking* framework for proactive deepfake defense and media provenance verification. Unlike conventional watermarking methods that treat watermark payloads as generic signals, SAIW formulates watermark embedding as a source-conditioned representation learning problem, where watermark identity encodes the originating source and modulates the embedding process to produce discriminative and traceable signatures. The framework integrates feature-wise linear modulation to inject source identity into the embedding network, enabling scalable multi-source watermark generation. A perceptual guidance module derived from human visual system priors ensures that watermark perturbations remain visually imperceptible while maintaining robustness. In addition, a dual-purpose forensic decoder simultaneously reconstructs the embedded watermark and performs source attribution, providing both automated verification and interpretable forensic evidence. Extensive experiments across multiple deepfake datasets demonstrate that SAIW achieves high perceptual quality while maintaining strong robustness against compression, filtering, noise, geometric transformations, and adversarial perturbations. By binding digital media to its origin through invisible yet verifiable markers, SAIW enables reliable authentication and source attribution, providing a scalable foundation for proactive deepfake defense and trustworthy media provenance.

**Index Terms**—Proactive Deepfake Defense, Invisible Watermarking, Source Attribution, Content Authentication

## I. INTRODUCTION

The rapid advancement of generative models has enabled the creation of highly realistic synthetic media, fundamentally challenging the reliability and authenticity of digital visual content. While these technologies enable creative media generation, they also amplify the threat posed by deepfakes [1]. Deepfakes are AI-generated videos, images, or audio that falsely depict individuals performing actions or making statements they never made. Such synthetic content can

facilitate misinformation, impersonation, and fraud, while eroding trust in digital evidence, particularly in domains where media authenticity is critical, including journalism, public communication, and legal proceedings.

Among different forms of synthetic manipulation, facial deepfakes have attracted significant attention in the biometrics and visual forensics communities [2]–[4]. The human face, one of the most socially and semantically important biometric traits, can now be manipulated using techniques such as identity swapping, attribute editing, and face reenactment [5]–[8]. Although these methods were originally developed for entertainment applications, the widespread availability of pre-trained generative models and consumer applications such as FaceApp [7] enables even non-expert users to generate convincing forged media within minutes. Consequently, deepfakes have become powerful tools for misinformation, political manipulation, harassment, financial fraud, and reputational damage, further undermining public trust in digital media. To address these challenges, most existing research focuses on deepfake detection methods that attempt to identify manipulation artifacts after synthetic media has been generated. These approaches typically rely on visual inconsistencies, blending artifacts, or physiological irregularities [9]–[12]. However, such reactive detection strategies face fundamental limitations. Detection models often struggle to generalize across unseen manipulation techniques, suffer from dataset bias, and remain vulnerable to adversarial attacks. Moreover, the rapid evolution of generative models frequently outpaces the adaptability of artifact-based detection systems [11], [13]. These limitations motivate a shift from reactive detection toward proactive mechanisms that secure media authenticity at the time of creation.

Proactive deepfake defense aims to embed verifiable authenticity signals directly into digital content, enabling reliable verification of origin and integrity throughout the media lifecycle. Recent approaches include cryptographic provenance tracking, identity-preserving perturbations, and robust digital watermarking [14], [15]. In parallel, initiatives such as the C2PA [16] propose embedding signed metadata within digital assets to maintain traceable provenance records. While such metadata-based mechanisms provide useful provenance information, they remain vulnerable to removal or corruption during platform-specific processing or adversarial manipulation [17]. As a result, embedding persistent authenticity markers directly within visual content has emerged as a promising complementary solution. Despite these advances, existing watermarking techniques still face challenges in deepfake defense. Many methods focus mainly on payload recovery without modeling source attribution

B. Das and C. Adak are with Dept. of CSE, Indian Institute of Technology Patna, India – 801106.

S. Chattopadhyay is with Dept. of CSE, Indian Institute of Technology Indore, India – 453552.

Z. Akhtar is with State University of New York Polytechnic Institute, NY, USA – 13502.

S. Dutta is with Dept. of CSE, Indian Institute of Technology Kanpur, India – 208016.

Corresponding author: Chandranath Adak (chandranath@iitp.ac.in)

or supporting scalable multi-source watermarking. Moreover, jointly achieving imperceptibility, robustness to real-world distortions, and reliable source identification remains difficult.

The key insight of our work is that watermark identity can act not only as a hidden payload but also as a conditioning signal that guides the embedding process, enabling a unified framework for scalable media provenance attribution. In this work, we propose SAIW for proactive deepfake defense and media provenance verification. Unlike conventional approaches that treat watermark payloads as generic signals, our method formulates watermark embedding as a source-conditioned representation learning problem in which watermark identity encodes the originating source and modulates the embedding process to produce discriminative and traceable signatures. Trusted entities embed source-specific invisible watermarks into media at creation, establishing persistent authenticity markers that remain verifiable under subsequent transformations. The framework integrates source-conditioned feature modulation with perceptual guidance based on human visual system characteristics to ensure imperceptibility while maintaining robustness to compression, filtering, noise, and other real-world distortions. In addition, a dual-purpose forensic decoder simultaneously reconstructs the embedded watermark and performs source attribution, enabling both automated verification and interpretable forensic evidence. By binding digital media to its origin through invisible yet verifiable markers, SAIW provides a scalable mechanism for proactive media authentication and complements existing deepfake detection methods by enabling reliable source authentication and media provenance. The major **contributions** of this paper are summarized as follows:

(i) **Source-conditioned watermarking formulation:** We reformulate digital watermarking as a source-conditioned representation learning problem, where the watermark identity encodes the originating media source and acts as a conditioning signal for the embedding process. Unlike conventional payload-based designs, this formulation enables a unified model to learn source-aware watermark representations and supports scalable provenance encoding across multiple generative sources.

(ii) **Content-adaptive perceptual embedding with identity modulation:** We propose a content-adaptive embedding mechanism that integrates perceptual guidance based on human visual system characteristics with watermark identity-driven feature modulation. By combining luminance adaptation and contrast masking with channel-wise conditioning, the framework dynamically allocates embedding strength to visually tolerant regions while enabling source-specific watermark patterns, thereby improving both imperceptibility and robustness.

(iii) **Unified forensic decoding for reconstruction and attribution:** We design a dual-purpose decoding architecture that jointly performs watermark reconstruction and source attribution. The shared representation enables reliable watermark recovery under distortions while learning discriminative source embeddings, providing both automated verification and interpretable forensic evidence for media provenance analysis.

The rest of the paper is organized as follows. Section II reviews related work. Section III presents the proposed method. Section IV describes the experiments and results. Finally, Section V concludes the paper.

## II. RELATED WORKS

The rapid proliferation of generative models has exposed fundamental limitations of artifact-based deepfake detection, which remains reactive and struggles to generalize across novel architectures. To address these limitations, recent research has increasingly shifted toward proactive, source-centric watermarking strategies that embed verifiable provenance directly into generated content, enabling integrity verification and source attribution. Existing approaches span multiple paradigms, including model-specific fingerprints for attribution [18], semi-fragile marks for tamper detection [19], [20], robust watermarking methods [14], [21], [22] designed to withstand real-world distortions and GAN-based post-processing [23], hybrid fragile robust frameworks that support joint detection and tracing [24], and identity–semantic watermarking schemes for authenticity verification [25]. Collectively, these approaches demonstrate the potential of watermarking as a proactive mechanism for media authentication while avoiding disruptive perturbations and enabling provenance tracing.

Within this paradigm, several architectures address different challenges through specialized designs. [26] and FaceSigns [19] focus on deepfake detection through controlled fragility, embedding invisible bit strings that become non-recoverable when malicious facial manipulations occur [19]. To mitigate watermark removal vulnerabilities, [26] introduce an adversarial network that improves resilience against both white-box and black-box attacks. SepMark [24] addresses the joint detection–attribution problem using a dual-channel architecture, enabling simultaneous robust source tracing and semi-robust tamper localization with a single embedded watermark. LampMark [27] instead focuses on improving generalization against hyper-realistic synthetic media by proposing a training-free, landmark-based perceptual watermark that exploits structural consistency changes introduced during deepfake generation, comparing recovered watermarks against dynamically generated structural identifiers for reliable detection. Despite these advances, existing approaches rarely provide transparent, human-interpretable forensic evidence and often lack principled multi-objective frameworks that jointly optimize imperceptibility, robustness, and source discriminability.

**Positioning of our work:** Despite significant progress, several challenges remain. Many frameworks require retraining or manual reconfiguration to handle emerging deepfake generators, limiting scalability, while multi-source watermarking remains difficult when embedding generator-specific identifiers without degrading visual fidelity. In addition, existing systems rarely provide dual-verification mechanisms that combine deepfake detection with reliable source tracing and clear, human-interpretable forensic evidence. Unlike conventional watermarking approaches that embed arbitrary payload bits, we introduce a source-conditioned watermark identity formulation in which the watermark itself encodes the generating media source and directly conditions the embedding process, enabling scalable multi-source provenance encoding within a unified framework. To address these limitations, we propose a scalable watermarking framework that embeds visual logos and machine-readable identifiers for different deepfake generators alongside






Fig. 1: Workflow of the proposed framework, SAIW

markers for authentic content. The architecture employs a dual-mode forensic decoder that produces both human-verifiable watermark recovery and automated source classification, enabling robust attribution under real-world distortions and adversarial attacks while jointly optimizing imperceptibility, robustness, and source discriminability.

### III. PROPOSED METHODOLOGY

We propose SAIW, a proactive deepfake defense framework that enables source-aware authentication by embedding verifiable watermarks at the time of content generation (Fig. 1). SAIW integrates feature-wise linear modulation with perceptual guidance to achieve both robustness and visual imperceptibility. It consists of two core components: a watermark embedding network ( $f_E$ ), which embeds source-specific watermarks into images, and a watermark extraction network ( $f_D$ ), which reliably recovers the embedded watermark under various distortions.

#### A. Overview and Problem Formulation

Our framework operates on a collection of cover images  $\mathcal{I} \in \mathbb{R}^{H \times W \times C}$ , each originating from a distinct source and paired with unique source-specific watermark  $\mathcal{I}_l \in \mathbb{R}^{h \times w \times c}$ . The embedding network  $f_E$  integrates these watermarks into the

corresponding images to generate visually indistinguishable outputs  $\mathcal{I}_w$ , while the extraction network  $f_D$  recovers the embedded watermark such that  $\hat{\mathcal{I}}_l \approx \mathcal{I}_l$ . This formulation enables scalable multi-source provenance encoding with high visual fidelity and reliable watermark recovery. Here,  $(H, W, C) = (256, 256, 3)$  and  $(h, w, c) = (64, 64, 1)$  denote the dimensions of the cover image and watermark, respectively.

Watermark embedding is formulated as an additive residual learning task, where the embedding network  $f_E$  generates a residual map  $\mathcal{I}_r$  conditioned on both the cover image  $\mathcal{I}$  and watermark  $\mathcal{I}_l$ . The watermarked image is obtained as  $\mathcal{I}_w = f_{clip}(\mathcal{I} + \mathcal{I}_r; 0, 1)$ , where  $f_{clip}$  constrains pixel intensities to the normalized range  $[0, 1]$ .

To ensure visual imperceptibility, a perceptual guidance map  $\mathcal{P}_n$  is derived from  $\mathcal{I}$  based on human visual priors. This map adaptively allocates stronger perturbations to textured or high-frequency regions while suppressing changes in perceptually sensitive areas, enabling effective concealment of the watermark.

The extraction network  $f_D$  operates on distorted inputs  $\mathcal{I}_d$  to recover the embedded watermark:  $\hat{\mathcal{I}}_l = f_D(\mathcal{I}_d)$ . The recovered watermark serves both for reconstruction and for source identification, enabling reliable attribution of the image origin.

### B. Watermark Embedding Network ( $f_E$ )

The embedding network  $f_E$  adopts a U-Net backbone integrated with a Hybrid Transformer bottleneck, enabling it to capture both fine-grained local features and global contextual dependencies. In addition, Feature-wise linear modulation  $f_M$  enables joint conditioning on the cover image and the source-specific watermark identity while being guided by the perceptual guidance block  $f_P$ .

1) *Source-Conditioned Watermark Encoding*: To achieve source-specific watermark embedding, the watermark  $\mathcal{I}_l$  is first processed by a lightweight logo encoder  $f_{logo}$  composed of  $L_l$  number of stacked Conv2D layers with batch normalization and GELU activations. The resulting feature maps are aggregated using adaptive average pooling followed by flattening to produce a compact global feature representation  $g_l \in \mathbb{R}^{128}$ . Here, we fix  $L_l = 5$ , empirically. Since watermark identities are represented through learned latent embeddings, the framework can scale to large numbers of media sources without redesigning the embedding architecture.

Now  $g_l$  is passed for feature-wise linear modulation module  $f_M$  to predict parameters  $\{(\gamma_i, \beta_i)\}_{i=1}^K$  used to condition multiple layers of the watermark encoder  $f_E$ . Here,  $K = 10$ . Given a feature map  $\mathbf{x}_i$  at the  $i^{\text{th}}$ -encoder stage,  $f_M$ -conditioning is defined as:  $f_M(\mathbf{x}_i; \gamma_i, \beta_i) = \gamma_i \odot \mathbf{x}_i + \beta_i$ , where  $\odot$  denotes element-wise multiplication, and  $\gamma_i$  and  $\beta_i$  represent the learned scaling and shifting parameters [28].  $f_M$  conditioning provides lightweight channel-wise modulation of encoder features across multiple resolutions, enabling the network to learn watermark-specific transformations and adaptively modulate embedding patterns according to the source identity with minimal parameter overhead. Appendix A and Fig. B.1 (*supp. file*) analyze  $f_M$ -based conditioning, highlighting its role in learning source-specific style embeddings.

2) *Content-Adaptive Embedding via Perceptual Guidance*: To balance imperceptibility and robustness, we engage a perceptual guidance block ( $f_P$ ) that estimates spatial perceptual tolerance from luminance structure. The perceptual guidance mechanism follows principles from human visual system models used in perceptual watermarking literature [29]. Given a cover image  $\mathcal{I} \in \mathbb{R}^{H \times W \times C}$ , it is first converted to YUV color space, and the luminance channel  $\mathcal{I}_y \in \mathbb{R}^{H \times W \times C_y}$  is extracted. Here,  $C_y = 1$ . Luminance adaptation is computed by estimating the local background luminance  $\mathbf{L}_{bg}$  from the luminance channel using a  $5 \times 5$  smoothing kernel  $\mathcal{K}_l$ , i.e.,  $\mathbf{L}_{bg} = \mathcal{K}_l * \mathcal{I}_y$ , where  $*$  denotes convolution operation. The perceptual luminance adaptation component  $\mathcal{P}_{LA}$  is then defined as:

$$\mathcal{P}_{LA} = \begin{cases} p_1 \left(1 - (\mathbf{L}_{bg}/p_2)^{1/2}\right) + p_3; & 0 \leq \mathbf{L}_{bg} \leq p_2 \\ p_3(\mathbf{L}_{bg} - p_2)/(p_2 + 1) + p_3; & p_2 < \mathbf{L}_{bg} \leq 255 \end{cases} \quad (1)$$

where  $p_1, p_2$  are scaling factors, and  $p_3$  acts a threshold. In parallel, contrast masking is obtained from spatial gradients computed using Sobel filters  $\mathcal{O}_x$  and  $\mathcal{O}_y$  to estimate local edge strength. Specifically, the gradient magnitude is computed as  $\nabla_s = ((\mathcal{O}_x * \mathcal{I}_y)^2 + (\mathcal{O}_y * \mathcal{I}_y)^2)^{1/2}$ , which captures the spatial variation of luminance. Based on this gradient response, the perceptual contrast masking component  $\mathcal{P}_{CM}$  is computed

as:  $\mathcal{P}_{CM} = (p_4(\nabla_s)^{p_5})/((\nabla_s)^2 + (p_6)^2)$ , where  $p_4$  is a scaling factor and  $p_6$  is a stabilization constant controlling the saturation of masking strength. Empirically, we set  $p_1 = 17$ ,  $p_2 = 127$ ,  $p_3 = 3$ ,  $p_4 = 1.872$ ,  $p_5 = 8$ , and  $p_6 = 26$ . Now,  $\mathcal{P}_{LA}$  and  $\mathcal{P}_{CM}$  are combined to form a perceptual map:  $\mathcal{P} = \mathcal{P}_{LA} + \mathcal{P}_{CM} - \lambda_1 \min(\mathcal{P}_{LA}, \mathcal{P}_{CM})$ , where  $\lambda_1$  is a tunable hyperparameter. The fused perceptual map  $\mathcal{P}$  is rectified using a ReLU activation and normalized to produce the perceptual guidance map  $\mathcal{P}_n = f_{clip}(\lambda_2 \text{ReLU}(\mathcal{P}); 0, 3)$ , where  $\lambda_2$  rescales the perceptual thresholds. Now,  $\mathcal{P}_n \in \mathbb{R}^{H \times W \times 1}$  is injected into the hybrid transformer bottleneck to modulate feature representations based on local perceptual sensitivity.

3) *Hybrid Transformer Bottleneck*: As illustrated in Fig. 1, the input image  $\mathcal{I} \in \mathbb{R}^{H \times W \times C}$  is first processed by an initial feature extractor  $f_d$ , implemented as Conv2D followed by a GELU activation, to extract low-level features. The resulting feature representation is then progressively encoded through a series of down-sampling blocks to obtain feature map  $g_h \in \mathbb{R}^{16 \times 16 \times 1024}$ , which serves as the input to the hybrid transformer bottleneck to capture both local structural cues and long-range spatial dependencies. This feature map  $g_h$  is then passed through Conv2D, followed by GELU, which first projects the feature into the transformer embedding space, producing  $h_1$ . The representation is then modulated through a  $f_M$  layer using dynamically generated parameters  $(\gamma, \beta)$ :  $h_2 = \gamma \odot h_1 + \beta$ , where  $\odot$  denotes element-wise scaling;  $h_1, h_2 \in \mathbb{R}^{16 \times 16 \times 512}$ .

The  $f_M$  conditioned features are flattened and enriched with learnable positional embeddings  $\mathcal{E}_{pos}$  before being processed by a stack of  $L_s$  number of Swin transformer layers [30]:  $h_3 = \mathcal{S}^{(L_s)}(h_2 + \mathcal{E}_{pos})$ , where  $\mathcal{S}^{(L_s)}(\cdot)$  denotes the Swin transformation with  $L_s$  layers;  $h_3 \in \mathbb{R}^{256 \times 512}$ . Here, we empirically set  $L_s = 4$ . The resulting representation  $h_3$  is reshaped and projected back to the original channel dimension using a Conv2D followed by batch norm and GELU, yielding  $h_4$ .

Now the perceptual guidance map  $\mathcal{P}_n$  is converted to  $\mathcal{P}_s$  using bilinear interpolation that scales  $h_4$  as:  $h_5 = (1 - \lambda_3 \cdot \mathcal{P}_s) \cdot h_4$ , where  $\lambda_3$  regulates the trade-off between suppressing perturbations in perceptually fragile regions and amplifying embedding strength in visually tolerant areas;  $h_4, h_5 \in \mathbb{R}^{16 \times 16 \times 1024}$ ,  $\mathcal{P}_s \in \mathbb{R}^{16 \times 16 \times 1}$ .

To incorporate perceptual priors,  $h_5$  is connected with the bottleneck input  $g_h$  through a skip connection:  $g'_h = g_h + h_5$ ;  $g'_h \in \mathbb{R}^{16 \times 16 \times 1024}$ .

This perceptually guided bottleneck allocates stronger watermark embedding in textured or edge-rich regions while suppressing distortions in visually sensitive smooth areas. The decoder stage then mirrors the encoder structure with  $f_M$  conditioning at each resolution level to progressively reconstruct the residual watermark map  $\mathcal{I}_r$ . Specifically, the decoded feature map is passed through a refinement block  $f_u$  consisting of a  $3 \times 3$  Conv2D followed by GELU activation and a  $1 \times 1$  Conv2D layer to project the features into the residual space. A tanh activation is then applied to constrain the residual values within a bounded range. The resulting residual map is further scaled by a learnable factor  $\lambda_4$  to control the embedding strength, yielding the final residual watermark representation:  $\mathcal{I}_r = \lambda_4 \cdot f(\mathcal{I}, \mathcal{I}_l)$ . Finally, the scaled residual  $\mathcal{I}_r$  is added to

the input image  $\mathcal{I}$  to produce the watermarked image  $\mathcal{I}_w$ ;  $\mathcal{I}_w = \mathcal{I}_r + \mathcal{I}$ . The scalar hyperparameters are empirically set to  $\lambda_1=0.3$ ,  $\lambda_2=\frac{1}{30}$ , and  $\lambda_3=0.5$ , while  $\lambda_4$  is learned during training to control the overall embedding strength. Appendix B further analyzes the contributions of  $f_M$ ,  $f_P$ , and the Swin transformer  $\mathcal{S}$  in enabling effective watermark conditioning.

### C. Watermark Extraction Network ( $f_D$ )

The extractor  $f_D$  is designed to robustly recover the embedded watermark and reconstruct it from distorted image  $\mathcal{I}_d$ . As shown in Fig. 1, the network employs a ConvNeXt backbone [31], selected for its powerful hierarchical feature representation.

1) *Distortion Transformation*: To simulate realistic degradations arising from post-processing operations, distribution pipelines, and potential adversarial manipulations, we apply a set of stochastic, non-differentiable transformations  $\mathcal{T}(\cdot)$  (e.g., compression, resizing, blurring, and noise injection). Formally, the transformations are applied to both the original image  $\mathcal{I}$  and the watermarked image  $\mathcal{I}_w$ , yielding  $\mathcal{I}'_w = \mathcal{T}(\mathcal{I}_w)$  and  $\mathcal{I}' = \mathcal{T}(\mathcal{I})$ . The resulting Image  $\mathcal{I}_d = (\mathcal{I}' | \mathcal{I}'_w)$  is then used for feature encoding.

2) *Feature Encoding*: The distorted input image  $\mathcal{I}_d$  is first processed by a Conv2D layer with a  $7 \times 7$  kernel and stride 4 to generate an initial low-level embedding. This representation is subsequently passed through four hierarchical ConvNeXt stages. Each stage progressively enlarges the receptive field while enhancing the semantic richness of the representation. Following the ConvNeXt design, each stage employs depthwise convolutions, layer normalization, GELU activations, and skip connections that enable efficient feature extraction and stable optimization [31]. After the final ConvNeXt stage (Stage 4), a compact feature map  $g_{map} \in \mathbb{R}^{8 \times 8 \times 1024}$  is obtained. This feature map encodes high-level semantic information and is fed into two parallel branches: a watermark reconstruction head ( $f_R$ ) and a source identification head ( $f_{id}$ ).

a) *Reconstruction Head*: The reconstruction head  $f_R$  progressively upsamples  $g_{map}$  through a convolutional decoder composed of multiple ConvTranspose2D layers followed by an upsampling block. This hierarchical decoding process gradually restores spatial resolution while preserving watermark-related structures. A final sigmoid activation produces the reconstructed binary watermark  $\hat{\mathcal{I}}_l$ , i.e.,  $\hat{\mathcal{I}}_l = f_R(g_{map})$ .

b) *Source Identification Head*: In parallel, the feature map  $g_{map}$  is transformed into a compact embedding for source identification. Specifically, global average pooling aggregates spatial information to produce a 1024-dimensional feature vector, which is then projected through a linear layer to obtain a latent embedding  $v_i = f_{id}(g_{map})$ ;  $v_i \in \mathbb{R}^{256}$ . To learn discriminative source-aware representations, the embedding is optimized using an angular margin-based identification loss  $\mathcal{L}_{id}$ . The classification is performed using cosine similarity between the embedding vector  $v_i$  and the corresponding class weights. This objective encourages embeddings from the same source to form compact clusters while increasing angular separation between different source identities. In conjunction with watermark reconstruction, this identification objective

provides complementary supervision: reconstruction ensures faithful recovery of the embedded watermark, while the classification head enables reliable provenance attribution even when the watermark is partially degraded.

### D. Training Loss

The proposed SAIW is trained under a composite objective that jointly optimizes for three criteria: (i) imperceptibility of the embedded watermark, (ii) robustness against distortions and attacks, and (iii) discriminative identity preservation. The total loss is formulated as:

$$\mathcal{L}_{total} = w_{imp}\mathcal{L}_{imp} + w_{rob}\mathcal{L}_{rob} + w_{id}\mathcal{L}_{id} \quad (2)$$

where  $w_{imp}$ ,  $w_{rob}$ , and  $w_{id}$  are non-negative weights balancing the respective objectives;  $w_{imp} + w_{rob} + w_{id} = 1$ .

1) *Imperceptibility Loss ( $\mathcal{L}_{imp}$ )*: To preserve perceptual fidelity between the cover image  $\mathcal{I}$  and the watermarked image  $\mathcal{I}_w$ , we use a weighted combination of pixel-wise  $L_1$  loss and the Learned Perceptual Image Patch Similarity (LPIPS) metric:

$$\mathcal{L}_{imp} = w_1 \mathcal{L}_{L_1}(\mathcal{I}, \mathcal{I}_w) + w_{lrips} \mathcal{L}_{lrips}(\mathcal{I}, \mathcal{I}_w), \quad (3)$$

where  $\mathcal{L}_{lrips}$  captures perceptual similarity more consistently with human vision, complementing the low-level fidelity enforced by  $L_1$  loss;  $w_1 + w_{lrips} = 1$ .

2) *Robustness Loss ( $\mathcal{L}_{rob}$ )*: The robustness objective ensures accurate recovery of the watermark  $\mathcal{I}_l$  from  $\mathcal{I}_d$ . Let  $\hat{\mathcal{I}}_l = f_D(\mathcal{I}_d)$  denote the reconstructed watermark logo. The robustness loss is defined as:

$$\mathcal{L}_{rob} = \mathbb{E}_{\mathcal{A}} \left[ w_2 \mathcal{L}_{L_1}(\hat{\mathcal{I}}_l, \mathcal{I}_l) + w_{ssim} \mathcal{L}_{ssim}(\hat{\mathcal{I}}_l, \mathcal{I}_l) + w'_{lrips} \mathcal{L}_{lrips}(\hat{\mathcal{I}}_l, \mathcal{I}_l) \right] \quad (4)$$

where  $\mathcal{L}_{ssim}$  promotes structural similarity,  $\mathbb{E}_{\mathcal{A}}[\cdot]$  denotes expectation over sampled attack conditions;  $w_2 + w_{ssim} + w'_{lrips} = 1$ .

3) *Identification Loss ( $\mathcal{L}_{id}$ )*: We employ the additive angular margin loss [32] to ensure intra-class compactness and inter-class separability of identity vectors. On the hypersphere,  $\mathcal{L}_{id}$  imposes an angular margin  $m$  between classes. For a batch of  $L_2$ -normalized vectors  $\{v_i\}_{i=1}^N$  with class labels  $\{y_i\}_{i=1}^C$ , the loss is given by:

$$\mathcal{L}_{id} = \frac{-1}{N} \sum_{i=1}^N \log \frac{\exp(\lambda_5 \cdot \cos(\theta_{y_i} + m))}{\exp(\lambda_5 \cdot \cos(\theta_{y_i} + m)) + \sum_{\substack{j=1 \\ j \neq y_i}}^C \exp(\lambda_5 \cdot \cos(\theta_j))} \quad (5)$$

where  $\theta_j$  is the angle between  $v_i$  and the class weight,  $\lambda_5$  is a scale parameter,  $m$  is the additive angular margin, and  $C$  is the total number of watermark classes plus 1 (for “no watermark”). This formulation encourages intra-class compactness and inter-class separation on the hypersphere. Empirically, we fix  $m = 0.4$ , mini-batch size  $N = 16$ ,  $\lambda_5 = 30$ . To train the overall SAIW framework, we engaged AdamW optimizer.

## IV. EXPERIMENTS AND DISCUSSIONS

This section presents the experimental evaluation of SAIW. All experiments were implemented in PyTorch 2.2.2 with CUDA 11.8 and conducted on a system equipped with an Intel Xeon W7-2495X (24 cores), 256 GB RAM, and an NVIDIA RTX A6000 GPU with 48 GB VRAM.



## A. Datasets

To evaluate the proposed framework under diverse manipulation scenarios, we conducted experiments on multiple publicly available datasets: IndicSideFace [33], FaceForensics (FF) [34], FaceForensics++ (FF++) [35], and CelebA [36]. The IndicSideFace dataset [33] served as our primary dataset and contains 984 genuine images from 164 Indian subjects across frontal and side-profile views, along with 21648 deepfake images generated using five identity-swapping methods (Ghost [37], SimSwap [6], SimSwap++ [38], FaceDancer [39], and InsightFace [5]) and one attribute-modification tool (FaceApp [7]). To assess cross-dataset generalization, we further evaluated SAIW on FF and FF++, sampling 60 frames per video with an 80:20 train-test split. Additionally, the CelebA [36] dataset was used to generate additional watermarked facial samples by sampling 202599 images across more than 200 identities with a 70:30 train-test split. For logo-conditioned watermarking, LOGO-30K [40] was used, where each logo acts as a unique identifier representing a specific manipulation source or platform, enabling structured watermark embedding and systematic source attribution. This multi-dataset evaluation protocol allows the framework to be evaluated under both controlled and cross-domain conditions, enabling comprehensive analysis of watermark robustness and source attribution capability.

## B. Evaluation Criteria

We evaluate SAIW according to three key criteria: *imperceptibility*, which measures the visual fidelity of the watermarked image; *robustness*, which evaluates the ability to recover the embedded watermark after the host image undergoes distortions; and *source identification*, which assesses the model's capability to correctly determine the origin of an image.

1) *Imperceptibility*: This quantifies the visual distortion introduced during watermark embedding [41]. We compare the original cover image  $\mathcal{I}$  and the corresponding watermarked image  $\mathcal{I}_w$  using three widely adopted perceptual metrics.

a) *PSNR (Peak Signal-to-Noise Ratio)*: PSNR measures the ratio between the maximum possible pixel value ( $\max_{\mathcal{I}}$ ) and the reconstruction distortion. For images normalized to  $[0, 1]$ ,  $\max_{\mathcal{I}} = 1$ .

$$\text{PSNR}(\mathcal{I}, \mathcal{I}_w) = 10 \log_{10} (\max_{\mathcal{I}}^2 / \text{MSE}) \quad (6)$$

where,  $\text{MSE} = \frac{1}{H \cdot W \cdot C} \sum_{i=1}^H \sum_{j=1}^W \sum_{k=1}^C (\mathcal{I}^{(i,j,k)} - \mathcal{I}_w^{(i,j,k)})^2$ ;  $\mathcal{I}, \mathcal{I}_w \in \mathbb{R}^{H \times W \times C}$ ;  $i, j, k$  denote height, width, channel indices. Higher PSNR values indicate better visual fidelity.

b) *SSIM (Structural Similarity Index Measure)*: SSIM evaluates perceptual similarity by comparing luminance, contrast, and structural information across image patches between  $\mathcal{I}$  and  $\mathcal{I}_w$ .

$$\text{SSIM}(\mathcal{I}, \mathcal{I}_w) = \frac{(2\mu_{\mathcal{I}}\mu_{\mathcal{I}_w} + c_1)(2\sigma_{\mathcal{I}, \mathcal{I}_w} + c_2)}{(\mu_{\mathcal{I}}^2 + \mu_{\mathcal{I}_w}^2 + c_1)(\sigma_{\mathcal{I}}^2 + \sigma_{\mathcal{I}_w}^2 + c_2)} \quad (7)$$

where  $\mu_{\mathcal{I}}, \mu_{\mathcal{I}_w}$  denote pixel sample means of  $\mathcal{I}, \mathcal{I}_w$ , respectively;  $\sigma_{\mathcal{I}}^2, \sigma_{\mathcal{I}_w}^2$  denote sample variances of  $\mathcal{I}, \mathcal{I}_w$ , respectively;  $\sigma_{\mathcal{I}, \mathcal{I}_w}$  represents their sample co-variance; and  $c_1, c_2$  are stabilization constants. The final score is obtained by averaging SSIM across all image patches.

## c) LPIPS (Learned Perceptual Image Patch Similarity):

LPIPS measures perceptual similarity using deep features extracted from a pretrained network  $\mathcal{N}$  [42]. Feature maps from multiple layers are channel-wise normalized, and the weighted squared  $L_2$  distance between corresponding activations is computed across spatial locations and channels. The final perceptual distance is obtained by averaging over spatial dimensions and aggregating across layers:

$$\text{LPIPS}(\mathcal{I}, \mathcal{I}_w) = \sum_l \frac{1}{H_l W_l} \sum_{i=1}^{H_l} \sum_{j=1}^{W_l} \sum_{k=1}^{C_l} w_l(k) \|\hat{\mathcal{N}}_l(\mathcal{I}^{(i,j,k)}) - \hat{\mathcal{N}}_l(\mathcal{I}_w^{(i,j,k)})\|_2^2 \quad (8)$$

where  $\hat{\mathcal{N}}_l(\cdot)$  denotes normalized features at layer  $l$ ;  $H_l, W_l, C_l$  are the feature dimensions; and  $w_l(k)$  are learned channel-wise weights. Lower LPIPS indicates higher perceptual similarity.

2) *Robustness*: This measures the ability of the extractor to faithfully recover the embedded watermark logo  $\mathcal{I}_l \in \mathbb{R}^{h \times w \times c}$  from a distorted watermarked image  $\mathcal{I}_w'$ . The quality of the recovered logo  $\hat{\mathcal{I}}_l$  is evaluated using:

a) *SSIM ( $\mathcal{A}_{ssim}$ )*: SSIM (Eq. 7) is applied to evaluate the structural integrity of the recovered logo  $\hat{\mathcal{I}}_l$  relative to the original logo  $\mathcal{I}_l$ ;  $\mathcal{A}_{ssim} = \text{SSIM}(\mathcal{I}_l, \hat{\mathcal{I}}_l)$ .

b) *Bit Recovery Accuracy ( $\mathcal{A}_{br}$ )*: To quantitatively assess the fidelity of the recovered watermark signal, we employ  $\mathcal{A}_{br}$  that measures the proportion of correctly reconstructed bits in the recovered watermark logo after binarization.

$$\mathcal{A}_{br}(\mathcal{I}_l, \hat{\mathcal{I}}_l) = \frac{1}{h \cdot w} \sum_{i=1}^h \sum_{j=1}^w \left[ 1 - \left| B(\mathcal{I}_l^{(i,j)}, \tau) - B(\hat{\mathcal{I}}_l^{(i,j)}, \tau) \right| \right] \quad (9)$$

where  $B(\mathbf{x}^{(i,j)}, \tau) = \mathbb{1}[\mathbf{x}^{(i,j)} > \tau]$  denotes a binarization operator that thresholds an image  $\mathbf{x}$  at a scalar threshold  $\tau$ ;  $\mathbb{1}[\cdot]$  is the indicator function.  $\mathcal{A}_{br}$  measures the normalized agreement between binarized ground-truth and reconstructed logos, estimating bit-level recovery fidelity.  $\mathcal{A}_{br} = 1$  indicates perfect reconstruction and is appropriate for binary payloads requiring exact bit recovery.

3) *Source Identification*: This component evaluates the model's ability to identify the generative or manipulation source of an input image using identification accuracy ( $\mathcal{A}_{id}$ ). In the multi-class setting, predictions are summarized using a confusion matrix, where each entry  $c_{ij}$  denotes the number of samples with ground-truth label  $i$  predicted as class  $j$ .  $\mathcal{A}_{id}$  is defined as:  $\mathcal{A}_{id} = (\sum_i c_{ii}) / (\sum_{i,j} c_{ij})$ , which represents the proportion of correctly classified samples. This metric reflects the model's effectiveness in attributing images to their true generative source under different attack scenarios.

## C. Quantitative Analysis

As shown in Table I, SAIW achieves a superior trade-off between imperceptibility and robustness compared to existing watermarking methods, including DwtDct, DwtDctSvd [43], MBRS [22], StegaStamp [14], RivaGAN [44], SepMark [24], and WAM [45]. In terms of imperceptibility, SAIW consistently produces high-fidelity watermarked images across all generators, achieving PSNR values above 50 dB, SSIM close to 1, and extremely low LPIPS scores. In contrast, existing methods operate at substantially lower PSNR ranges

TABLE I: Quantitative comparison of watermarking methods across multiple deepfake generators on IndicSideFace [33]

| Generator       | Methods         | Imperceptibility |                 |                    | Pre-Attack<br>( $\mathcal{A}_{br} \uparrow$ ) | Post-Attack ( $\mathcal{A}_{br} \uparrow$ ) |             |             |             |             |
|-----------------|-----------------|------------------|-----------------|--------------------|-----------------------------------------------|---------------------------------------------|-------------|-------------|-------------|-------------|
|                 |                 | PSNR $\uparrow$  | SSIM $\uparrow$ | LPIPS $\downarrow$ |                                               | Brightness                                  | Contrast    | JPEG        | G-Noise     | G-Blur      |
| FaceApp [7]     | DwtDet          | 44.40            | 0.981           | 0.0014             | 0.75                                          | 0.74                                        | 0.74        | 0.74        | 0.74        | 0.75        |
|                 | DwtDetSvd [43]  | 40.92            | 0.981           | 0.0048             | 0.87                                          | 0.74                                        | 0.74        | 0.75        | 0.82        | 0.82        |
|                 | MBRS [22]       | 43.72            | 0.989           | 0.0024             | <b>0.99</b>                                   | 0.83                                        | 0.91        | 0.88        | 0.87        | 0.98        |
|                 | StegaStamp [14] | 34.97            | 0.966           | 0.0254             | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.80        | 0.95        | <b>0.99</b> |
|                 | RivaGAN [44]    | 41.47            | 0.995           | 0.0370             | <b>0.99</b>                                   | 0.88                                        | 0.86        | 0.96        | 0.98        | 0.98        |
|                 | SepMark [24]    | 35.24            | 0.939           | 0.0120             | 0.86                                          | 0.87                                        | 0.87        | 0.94        | 0.93        | 0.87        |
|                 | WAM [45]        | 42.25            | 0.994           | 0.0139             | <b>0.99</b>                                   | 0.93                                        | 0.91        | 0.98        | 0.96        | <b>0.99</b> |
|                 | SAiW            | <b>55.81</b>     | <b>0.999</b>    | <b>0.0009</b>      | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | <b>0.99</b> | <b>0.99</b> | <b>0.99</b> |
| FaceDancer [39] | DwtDet          | 45.50            | 0.986           | 0.0013             | 0.75                                          | 0.74                                        | 0.74        | 0.74        | 0.74        | 0.75        |
|                 | DwtDetSvd [43]  | 41.56            | 0.988           | 0.0046             | 0.85                                          | 0.74                                        | 0.74        | 0.75        | 0.81        | 0.80        |
|                 | MBRS [22]       | 44.01            | 0.988           | 0.0024             | <b>0.99</b>                                   | 0.81                                        | 0.89        | 0.87        | 0.87        | 0.98        |
|                 | StegaStamp [14] | 35.34            | 0.969           | 0.0253             | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.81        | 0.95        | <b>0.99</b> |
|                 | RivaGAN [44]    | 41.26            | 0.996           | 0.0380             | 0.98                                          | 0.87                                        | 0.85        | 0.94        | 0.98        | 0.98        |
|                 | SepMark [24]    | 35.21            | 0.932           | 0.0150             | 0.86                                          | 0.85                                        | 0.86        | 0.95        | 0.91        | 0.84        |
|                 | WAM [45]        | 41.69            | 0.994           | 0.0168             | <b>0.99</b>                                   | 0.95                                        | 0.92        | <b>0.99</b> | 0.96        | <b>0.99</b> |
|                 | SAiW            | <b>56.16</b>     | <b>0.999</b>    | <b>0.0005</b>      | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.97        | <b>0.99</b> | <b>0.99</b> |
| Ghost [37]      | DwtDet          | 45.42            | 0.985           | 0.0015             | 0.75                                          | 0.74                                        | 0.74        | 0.74        | 0.73        | 0.75        |
|                 | DwtDetSvd [43]  | 41.61            | 0.988           | 0.0050             | 0.85                                          | 0.74                                        | 0.75        | 0.75        | 0.81        | 0.80        |
|                 | MBRS [22]       | 44.08            | 0.988           | 0.0025             | <b>0.99</b>                                   | 0.81                                        | 0.89        | 0.87        | 0.86        | 0.98        |
|                 | StegaStamp [14] | 35.30            | 0.969           | 0.0247             | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.81        | 0.95        | <b>0.99</b> |
|                 | RivaGAN [44]    | 41.28            | 0.996           | 0.0380             | 0.98                                          | 0.86                                        | 0.85        | 0.95        | <b>0.98</b> | 0.98        |
|                 | SepMark [24]    | 35.13            | 0.929           | 0.0160             | 0.89                                          | 0.86                                        | 0.88        | 0.96        | 0.90        | 0.82        |
|                 | WAM [45]        | 41.75            | 0.994           | 0.0149             | <b>0.99</b>                                   | 0.95                                        | 0.93        | <b>0.98</b> | 0.95        | <b>0.99</b> |
|                 | SAiW            | <b>57.49</b>     | <b>0.999</b>    | <b>0.0003</b>      | <b>0.99</b>                                   | 0.97                                        | <b>0.99</b> | <b>0.98</b> | <b>0.98</b> | <b>0.99</b> |
| SimSwap [6]     | DwtDet          | 45.46            | 0.986           | 0.0012             | 0.75                                          | 0.74                                        | 0.74        | 0.74        | 0.75        | 0.75        |
|                 | DwtDetSvd [43]  | 41.62            | 0.988           | 0.0042             | 0.85                                          | 0.75                                        | 0.75        | 0.75        | 0.82        | 0.80        |
|                 | MBRS [22]       | 43.81            | 0.989           | 0.0023             | <b>0.99</b>                                   | 0.82                                        | 0.90        | 0.88        | 0.87        | 0.98        |
|                 | StegaStamp [14] | 35.10            | 0.968           | 0.0256             | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.81        | 0.95        | <b>0.99</b> |
|                 | RivaGAN [44]    | 41.21            | 0.996           | 0.0390             | 0.96                                          | 0.86                                        | 0.85        | 0.95        | 0.98        | 0.98        |
|                 | SepMark [24]    | 34.80            | 0.930           | 0.0150             | 0.91                                          | 0.91                                        | 0.88        | 0.94        | 0.90        | 0.87        |
|                 | WAM [45]        | 42.04            | 0.995           | 0.0154             | <b>0.99</b>                                   | 0.96                                        | 0.94        | <b>0.98</b> | 0.96        | <b>0.99</b> |
|                 | SAiW            | <b>51.40</b>     | <b>0.998</b>    | <b>0.0008</b>      | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | <b>0.98</b> | <b>0.99</b> | <b>0.99</b> |
| SimSwap++ [38]  | DwtDet          | 45.43            | 0.985           | 0.0012             | 0.74                                          | 0.74                                        | 0.74        | 0.74        | 0.74        | 0.75        |
|                 | DwtDetSvd [43]  | 41.62            | 0.987           | 0.0040             | 0.85                                          | 0.74                                        | 0.74        | 0.75        | 0.82        | 0.80        |
|                 | MBRS [22]       | 43.83            | 0.988           | 0.0025             | <b>0.99</b>                                   | 0.81                                        | 0.89        | 0.87        | 0.87        | 0.97        |
|                 | StegaStamp [14] | 35.14            | 0.967           | 0.0254             | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.81        | 0.95        | <b>0.99</b> |
|                 | RivaGAN [44]    | 41.27            | 0.996           | 0.0380             | <b>0.99</b>                                   | 0.89                                        | 0.87        | 0.96        | 0.98        | 0.98        |
|                 | SepMark [24]    | 34.91            | 0.932           | 0.0140             | 0.88                                          | 0.86                                        | 0.86        | 0.96        | 0.91        | 0.82        |
|                 | WAM [45]        | 42.03            | 0.995           | 0.0143             | <b>0.99</b>                                   | 0.95                                        | 0.95        | <b>0.98</b> | 0.96        | <b>0.99</b> |
|                 | SAiW            | <b>53.21</b>     | <b>0.999</b>    | <b>0.0006</b>      | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.97        | <b>0.99</b> | <b>0.99</b> |
| InsightFace [5] | DwtDet          | 45.55            | 0.986           | 0.0015             | 0.75                                          | 0.74                                        | 0.74        | 0.74        | 0.74        | 0.75        |
|                 | DwtDetSvd [43]  | 41.67            | 0.988           | 0.0046             | 0.85                                          | 0.75                                        | 0.75        | 0.75        | 0.82        | 0.80        |
|                 | MBRS [22]       | 43.98            | 0.98            | 0.0024             | <b>0.99</b>                                   | 0.83                                        | 0.90        | 0.87        | 0.87        | 0.98        |
|                 | StegaStamp [14] | 35.26            | 0.967           | 0.0256             | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | 0.81        | 0.95        | <b>0.99</b> |
|                 | RivaGAN [44]    | 41.27            | 0.996           | 0.0376             | 0.98                                          | 0.87                                        | 0.86        | 0.95        | 0.98        | 0.98        |
|                 | SepMark [24]    | 35.15            | 0.927           | 0.0160             | 0.91                                          | 0.87                                        | 0.88        | 0.96        | 0.91        | 0.85        |
|                 | WAM [45]        | 41.82            | 0.992           | 0.0138             | <b>0.99</b>                                   | 0.96                                        | 0.95        | <b>0.99</b> | 0.97        | <b>0.99</b> |
|                 | SAiW            | <b>51.30</b>     | <b>0.998</b>    | <b>0.0005</b>      | <b>0.99</b>                                   | <b>0.99</b>                                 | <b>0.99</b> | <b>0.99</b> | <b>0.99</b> | <b>0.99</b> |

TABLE II: Cross-dataset robustness evaluation on public datasets

| Dataset          |                    | CelebA [36]                   |                               | FF++ [35] | FF [34] |
|------------------|--------------------|-------------------------------|-------------------------------|-----------|---------|
| Imperceptibility | PSNR $\uparrow$    | 55.89                         | 55.92                         | 55.83     |         |
|                  | SSIM $\uparrow$    | 0.9994                        | 0.9996                        | 0.9993    |         |
|                  | LPIPS $\downarrow$ | 0.0016                        | 0.0007                        | 0.0019    |         |
| Pre-Attack       |                    | $\mathcal{A}_{ssim} \uparrow$ | 0.99                          | 0.99      | 0.99    |
|                  |                    | $\mathcal{A}_{br} \uparrow$   | 0.99                          | 0.99      | 0.99    |
| Post-Attack      | Instagram Filters  | A                             | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |
|                  |                    | B                             | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |
|                  |                    | C                             | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |
|                  | JPEG               | A+B+C                         | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |
|                  |                    | QF = 50%                      | $\mathcal{A}_{ssim} \uparrow$ | 0.93      | 0.97    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.96      | 0.98    |
|                  | G-Blur             | QF = 75%                      | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |
|                  |                    | $\mathcal{K} = 3 \times 3$    | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |
|                  |                    | $\mathcal{K} = 5 \times 5$    | $\mathcal{A}_{ssim} \uparrow$ | 0.99      | 0.99    |
|                  |                    |                               | $\mathcal{A}_{br} \uparrow$   | 0.99      | 0.99    |

Instagram filters: A (Aden), B (Brooklyn), C (Clarendon)

and exhibit higher perceptual distortion, indicating that the proposed perceptual embedding strategy effectively minimizes visible artifacts while preserving structural integrity. For pre-attack evaluation, SAIW achieves near-perfect bit recovery accuracy ( $\mathcal{A}_{br}$ ) across all generators, matching the strongest baselines such as MBRS, StegaStamp, and WAM. Importantly, this performance is achieved while maintaining significantly higher visual fidelity, highlighting a clear advantage in the

fidelity-robustness trade-off. Under post-attack conditions, SAIW maintains consistently high recovery performance across diverse distortions, including brightness, contrast, JPEG compression, Gaussian noise, and blur, with  $\mathcal{A}_{br}$  remaining above 0.97 in nearly all cases. In comparison, several baselines exhibit noticeable degradation under challenging perturbations; for instance, StegaStamp shows reduced robustness under JPEG compression, while MBRS and SepMark experience performance drops under stochastic noise. Although WAM demonstrates strong robustness, SAIW achieves comparable recovery performance while significantly improving perceptual quality. Notably, the consistent robustness across heterogeneous distortions indicates that the model learns distortion-invariant watermark representations rather than overfitting to specific perturbation patterns. This property is critical for reliable deployment in real-world scenarios, where distortions are often unpredictable and compositional. These results demonstrate that SAIW effectively integrates perceptual-aware embedding with robust decoding, achieving high visual fidelity (PSNR  $>51$  dB, SSIM  $>0.998$ , LPIPS  $<0.001$ ) alongside reliable watermark recovery ( $\mathcal{A}_{br} \geq 0.97$ ) across diverse generators and attack scenarios. This balance is essential for scalable and dependable source attribution in deepfake media.

TABLE III: Comparison of watermark imperceptibility and robustness under Instagram-style photometric distortions on CelebA [36]

| Model                  | Imperceptibility |                 | Pre-Attack<br>( $\mathcal{A}_{br} \uparrow$ ) | Post-Attack ( $\mathcal{A}_{br} \uparrow$ ) |               |               |               |               |               |               |               | JPEG          | G-Blur |
|------------------------|------------------|-----------------|-----------------------------------------------|---------------------------------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|--------|
|                        | PSNR $\uparrow$  | SSIM $\uparrow$ |                                               | A                                           | B             | C             | A+B           | B+C           | A+C           | A+B+C         |               |               |        |
| SemiFragile-DCT [46]   | 20.29            | 0.846           | 0.9943                                        | 0.9347                                      | 0.9579        | 0.9512        | 0.9239        | 0.9342        | 0.9474        | 0.9156        | 0.5386        | 0.7624        |        |
| HiDDeN [21]            | 24.96            | 0.928           | 0.9765                                        | 0.9461                                      | 0.9382        | 0.9413        | 0.9137        | 0.9129        | 0.9223        | 0.8979        | 0.6719        | 0.8497        |        |
| StegaStamp [14]        | 28.64            | 0.922           | 0.9962                                        | 0.9948                                      | 0.9926        | 0.9909        | 0.9718        | 0.9628        | 0.9579        | 0.9413        | 0.9524        | 0.9914        |        |
| FaceSigns [19]         | 36.08            | 0.975           | 0.9949                                        | 0.9945                                      | 0.9922        | 0.9915        | 0.9853        | 0.9786        | 0.9751        | 0.9636        | 0.9375        | 0.9824        |        |
| Nadimpalli et al. [26] | 35.57            | 0.970           | 0.9945                                        | 0.9939                                      | 0.9925        | 0.9919        | 0.9887        | 0.9854        | 0.9869        | 0.9718        | 0.9452        | 0.9918        |        |
| SAiW                   | <b>55.89</b>     | <b>0.999</b>    | <b>0.9974</b>                                 | <b>0.9970</b>                               | <b>0.9967</b> | <b>0.9963</b> | <b>0.9932</b> | <b>0.9898</b> | <b>0.9942</b> | <b>0.9967</b> | <b>0.9955</b> | <b>0.9968</b> |        |

Instagram filters: A (Aden), B (Brooklyn), C (Clarendon)

Instagram filters: A (Aden), B (Brooklyn), C (Clarendon)

1) *Performance of Proactive Detector on Cross Datasets:* SAIW was evaluated on three widely used public benchmarks, namely FF [34], FF++ [35], and CelebA [36], to assess both imperceptibility and robustness under cross-dataset settings, while trained on IndicSideFace [33]. As summarized in Table II, SAIW achieved consistently high perceptual fidelity, with PSNR values above 55 dB, SSIM scores exceeding 0.99, and LPIPS ranging from 0.0007 to 0.0019, indicating negligible visual distortion across diverse datasets. In addition to high imperceptibility, SAIW demonstrated strong robustness under various real-world distortions. Under Instagram-style photometric perturbations (A: Aden, B: Brooklyn, and C: Clarendon), both structural similarity ( $\mathcal{A}_{ssim}$ ) and bit recovery accuracy ( $\mathcal{A}_{br}$ ) remained consistently at 0.99 across all filters and their combinations. Under JPEG compression, SAIW maintained  $\mathcal{A}_{br}$  values of 0.96-0.98 at 50% QF (quality factor), and near-perfect recovery at 75% QF, demonstrating resilience even under strong compression. Similarly, under Gaussian blur with varying kernel sizes, both  $\mathcal{A}_{ssim}$  and  $\mathcal{A}_{br}$  remained stable at 0.99, indicating robustness to convolutional degradations.

These results comprehend that SAIW generalizes effectively across heterogeneous datasets while preserving high perceptual quality and robust watermark recovery. The consistent performance under diverse distortion scenarios highlights its suitability for practical and scalable deepfake forensics and proactive media integrity verification.

2) *Robustness under Photometric Filter Distortions:* We evaluated the robustness of SAIW under real-world photometric distortions by applying Instagram-style filters (A: Aden, B: Brooklyn, and C: Clarendon, and their combinations) to images from the CelebA dataset [36]. As shown in Table III, SAIW obtained substantially higher imperceptibility compared to existing baselines, with the highest PSNR (55.89 dB) and SSIM (0.999), indicating minimal visual distortion after watermark embedding. More importantly, SAIW demonstrated consistently strong robustness under both individual and compositional filter perturbations. While prior methods exhibit noticeable degradation in  $\mathcal{A}_{br}$  from pre-attack to A+B+C post-attack (e.g., SemiFragile-DCT [46] drops to 0.9156, HiDDeN [21] to 0.8979, and FaceSigns [19] and [26] to 0.9636 and 0.9718, respectively), SAIW maintains near-perfect recovery, with  $\mathcal{A}_{br}$  consistently above 0.99 across all filter settings. This includes challenging scenarios involving multiple filter compositions, where nonlinear photometric transformations typically degrade watermark integrity. Furthermore, SAIW exhibited strong resilience under additional distortions such as JPEG compression and Gaussian blur, achieving recovery accuracies of 0.9955 and 0.9968, respectively, on CelebA [36].

These results highlight the robustness of the learned em-

bedding and decoding strategy in preserving watermark information under diverse real-world transformations. Notably, the strong performance under filter compositions suggests that SAIW learns invariant watermark representations rather than overfitting to specific distortion types, thereby effectively balancing high imperceptibility with robust watermark recovery and enabling strong generalization under complex photometric distortions commonly encountered in social media pipelines.

#### D. Qualitative Analysis

SAiW embeds a source-conditioned watermark  $\mathcal{I}_l$  into images, enabling fine-grained provenance encoding and reliable generator-level attribution. Each image, whether genuine or manipulated, is associated with a distinct binary watermark that serves as a persistent and verifiable source signature. Fig. 2 presents qualitative results demonstrating the end-to-end embedding and recovery pipeline on [33]. The input image  $\mathcal{I}$  and corresponding watermark  $\mathcal{I}_l$  are jointly processed, while the perceptual map  $\mathcal{P}_n$  identifies visually tolerant regions that guide imperceptible embedding. The bottleneck representation  $g'_h$  captures the fused content-watermark features, and the residual map  $\mathcal{I}_r$  reveals that embedding perturbations are spatially localized and low in magnitude. The resulting watermarked image  $\mathcal{I}_w$  remains visually indistinguishable from the original, even after undergoing distortions to produce  $\mathcal{I}_d$ . Crucially, the recovered watermark  $\hat{\mathcal{I}}_l$  retains strong structural fidelity to the original logo.

Across diverse poses, textures, and imaging conditions, the model consistently concentrates watermark energy in perceptually insensitive regions, ensuring minimal visual degradation while preserving robustness. The stability of reconstruction under distortions demonstrates that the learned representation captures invariant watermark features rather than distortion-specific artifacts. This unified embedding-decoding behavior enables reliable source attribution and interpretable forensic validation, highlighting the practical viability of SAIW for real-world media authentication.

1) *Robustness under Diverse Real-World Distortions:* We comprehensively evaluate the robustness of SAIW under a wide spectrum of real-world distortions, with qualitative results illustrated in Fig. 3. The evaluation encompasses geometric transformations (rotation, cropping, flipping, resizing), photometric and color variations (noise, color jitter, grayscale), filtering operations (blurring and sharpening), and severe lossy compression. As observed, the embedded watermark is consistently recovered with high fidelity across all scenarios on both IndicSideFace [33] and CelebA [36], achieving  $\mathcal{A}_{ssim}$  of approximately 0.99. Notably, the method remains stable even under challenging conditions such as aggressive JPEG





Fig. 2: Qualitative analysis of the proposed framework SAIW on IndicSideFace [33]. Softcopy exhibits better display.




Fig. 3: Robustness evaluation of SAIW under diverse real-world distortions. *Top*: The embedded “Aol.” logo is perfectly recovered from an IndicSideFace [33] sample after multiple distortions. *Bottom*: The “©” watermark is recovered from CelebA [36] sample after distortions.




Fig. 4: (a) Confusion matrix for source identification, (b) UMAP of extractor features between deepfake generators and originals.

compression, substantial cropping, and complex perturbations, including text overlays and Instagram-style filters. These results demonstrate that the learned embedding-decoding mechanism captures distortion-invariant watermark representations, ensuring reliable recovery under diverse transformations. Overall, the strong resilience across heterogeneous distortions highlights the practical suitability of the proposed framework for real-world

media authentication and content protection.

2) *Watermark-Aware Source Attribution and Feature Space Analysis*: To assess the discriminative capability of the learned representations, we extend the evaluation to a multi-class source attribution task, where each image is assigned to its exact generative origin. The label space includes multiple manipulation methods (e.g., FaceApp, FaceDancer, Ghost, InsightFace, SimSwap, and SimSwap++), along with two control categories: watermarked genuine, and no-watermark class. The performance of SAIW is presented using a confusion matrix and latent space visualization in Fig. 4:(a)-(b). The confusion matrix exhibits strong diagonal dominance with minimal cross-class confusion, indicating reliable separation among generators, with an overall identification accuracy  $\mathcal{A}_{id}$  of 84.06%. Notably, the no-watermark class is distinctly recognized, reducing false positives and ensuring high specificity, which is critical for real-world forensic deployment. The UMAP projection in Fig. 4:(b) reveals a well-structured feature space, where samples form compact and well-separated clusters across sources. This

indicates that SAIW captures generator-specific signatures beyond watermark presence, enabling discriminative and robust source attribution under challenging conditions.

Additional analyses on semantic-style disentanglement, margin-based losses, and texture effects on  $\mathcal{I}_r$  are presented in Appendices C-E (*supp. file*).

## V. CONCLUSION

This work introduces SAIW, a source-conditioned deep watermarking framework for proactive media authentication and provenance verification. Unlike conventional watermarking methods that embed generic payloads, the proposed formulation treats watermark identity as a conditioning signal that directly guides the embedding process, enabling a unified model to learn source-aware watermark representations. By integrating a hybrid convolution–transformer embedding network with perceptual guidance inspired by human visual system principles, the framework distributes watermark perturbations adaptively while maintaining visual fidelity. Furthermore, the proposed dual-purpose forensic decoder jointly performs watermark reconstruction and source attribution, providing both automated verification and interpretable provenance evidence. Through this design, SAIW establishes a scalable and reliable mechanism for binding digital media to its origin, offering a practical foundation for trustworthy media provenance and strengthening proactive defenses against synthetic media manipulation. Future work will extend the framework to video and multimodal media and explore robustness against generative regeneration pipelines to support large-scale provenance tracking in evolving generative ecosystems.

## REFERENCES

- [1] Z. Akhtar *et al.*, “Video and audio deepfake datasets and open issues in deepfake technology: being ahead of the curve,” *Forensic Sciences*, vol. 4, no. 3, pp. 289–377, 2024.
- [2] T. T. Nguyen *et al.*, “Deep learning for deepfakes creation and detection: A survey,” *Computer Vision and Image Understanding*, vol. 223, p. 103525, 2022.
- [3] M. Westerlund, “The Emergence of Deepfake Technology: A Review,” *Technology Innovation Management Review*, vol. 9, pp. 40–53, 2019.
- [4] B. Das *et al.*, “Securing AI-Generated Media: Rethinking Deepfake Vulnerabilities in Side-Face Perspectives,” in *BuildSEC*, 2025, pp. 1–7.
- [5] J. Guo *et al.*, “InsightFace: 2D and 3D Face Analysis Project,” <https://github.com/deepinsight/insightface>.
- [6] R. Chen *et al.*, “SimSwap: An efficient framework for high fidelity face swapping,” in *ACM MM*, 2020, pp. 2003–2011.
- [7] FaceApp Technology Ltd., “FaceApp,” <https://www.faceapp.com>.
- [8] Y. Nirkin *et al.*, “FSGAN: Subject Agnostic Face Swapping and Reenactment,” in *ICCV*, 2019, pp. 7183–7192.
- [9] Y. Li *et al.*, “In Ictu Oculi: Exposing AI Created Fake Videos by Detecting Eye Blinking,” in *IEEE WIFS*, 2018, pp. 1–7.
- [10] B. Dolhansky *et al.*, “The DeepFake Detection Challenge (DFDC) Dataset,” *arXiv:2006.07397*, 2020.
- [11] F. Chollet, “Xception: Deep Learning with Depthwise Separable Convolutions,” in *CVPR*, 2017, pp. 1800–1807.
- [12] A. Haliassos *et al.*, “Lips Don’t Lie: A Generalisable and Robust Approach to Face Forgery Detection,” in *CVPR*, 2021, pp. 5037–5047.
- [13] R. Wang *et al.*, “Anti-Forgery: Towards a Stealthy and Robust DeepFake Disruption Attack via Adversarial Perceptual-aware Perturbations,” in *IJCAI*, 2022, pp. 761–767.
- [14] M. Tancik *et al.*, “StegaStamp: Invisible Hyperlinks in Physical Photographs,” in *CVPR*, 2020.
- [15] A. Nadimpalli *et al.*, “ProActive DeepFake Detection using GAN-based Visible Watermarking,” *ACM TOMM*, vol. 20, no. 11, 2024.
- [16] J. MO *et al.*, “Towards Trustworthy Digital Media In The Aigc Era: An Introduction To The Upcoming IsoJpegTrust Standard,” *IEEE Communications Standards Magazine*, vol. 7, no. 4, pp. 2–5, 2023.
- [17] A. Black *et al.*, “VPN: Video Provenance Network for Robust Content Attribution,” in *ACM CVMP*, 2021, pp. 1–10.
- [18] N. Yu *et al.*, “Artificial Fingerprinting for Generative Models: Rooting Deepfake Attribution in Training Data,” in *ICCV*, 2021, pp. 14448–14457.
- [19] P. Neehara *et al.*, “FaceSigns: Semi-fragile Watermarks for Media Authentication,” *ACM TOMM*, vol. 20, no. 11, 2024.
- [20] Y. Yang *et al.*, “Faceguard: Proactive deepfake detection,” *arXiv*, 2021.
- [21] J. Zhu *et al.*, “HiDDeN: Hiding Data With Deep Networks,” in *ECCV*, 2018, p. 682–697.
- [22] Z. Jia *et al.*, “MBRS: Enhancing Robustness of DNN-based Watermarking by Mini-Batch of Real and Simulated JPEG Compression,” in *ACM MM*, 2021, pp. 41–49.
- [23] J. Huang *et al.*, “ARWGAN: Attention-Guided Robust Image Watermarking Model Based on GAN,” *IEEE TIM*, vol. 72, pp. 1–17, 2023.
- [24] X. Wu *et al.*, “Sepmark: Deep separable watermarking for unified source tracing and deepfake detection,” in *ACM MM*, 2023, pp. 1190–1201.
- [25] T. Wang *et al.*, “Robust identity perceptual watermark against deepfake face swapping,” *arXiv*, 2023.
- [26] A.V.Nadimpalli *et al.*, “Social Media Authentication and Combating Deepfakes Using Semi-Fragile Invisible Image Watermarking,” *Digital Threats*, vol. 5, no. 4, Dec. 2024.
- [27] T. Wang *et al.*, “LampMark: Proactive Deepfake Detection via Training-Free Landmark Perceptual Watermarks,” in *ACM MM*, 2024, p. 10515–10524.
- [28] E. Perez *et al.*, “FiLM: Visual Reasoning with a General Conditioning Layer,” in *AAAI*, vol. 32, no. 1, 2018.
- [29] J. Wu *et al.*, “Enhanced Just Noticeable Difference Model for Images With Pattern Complexity,” *IEEE TIP*, vol. 26, no. 6, pp. 2682–2693, 2017.
- [30] Z. Liu *et al.*, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *ICCV*, 2021, pp. 10012–10022.
- [31] Z. Liu, H. Mao *et al.*, “A ConvNet for the 2020s,” in *CVPR*, 2022, pp. 11976–11986.
- [32] J. Deng *et al.*, “ArcFace: Additive Angular Margin Loss for Deep Face Recognition,” in *CVPR*, 2019, pp. 4690–4699.
- [33] A. Deo *et al.*, “IndicSideFace: A dataset for advancing deepfake detection on side-face perspectives of indian subjects,” in *FG*, IEEE, 2025.
- [34] A. Rössler *et al.*, “FaceForensics: A large-scale video dataset for forgery detection in human faces,” *arXiv*, 2018.
- [35] A. Rössler *et al.*, “FaceForensics++: Learning to Detect Manipulated Facial Images,” in *ICCV*, 2019.
- [36] Z. Liu *et al.*, “Deep Learning Face Attributes in the Wild,” in *ICCV*, 2015.
- [37] A. Groshev *et al.*, “GHOST: a new face swap approach for image and video domains,” *IEEE Access*, vol. 10, pp. 83452–83462, 2022.
- [38] X. Chen *et al.*, “SimSwap++: Towards faster and high-quality identity swapping,” *IEEE TPAMI*, vol. 46, no. 1, pp. 576–592, 2023.
- [39] F. Rosberg *et al.*, “Facedancer: Pose-and occlusion-aware high fidelity face swapping,” in *WACV*, 2023, pp. 3454–3463.
- [40] X. Cun and C.-M. Pun, “Split then refine: stacked attention-guided resunets for blind single image visible watermark removal,” in *AAAI*, vol. 35, no. 2, 2021, pp. 1184–1192.
- [41] H. Zhao *et al.*, “Loss functions for image restoration with neural networks,” *IEEE TCI*, vol. 3, no. 1, pp. 47–57, 2016.
- [42] R. Zhang *et al.*, “The unreasonable effectiveness of deep features as a perceptual metric,” in *CVPR*, 2018, pp. 586–595.
- [43] K. A. Navas *et al.*, “DWT-DCT-SVD based watermarking,” in *COM-SWARE*, 2008, pp. 271–274.
- [44] K. A. Zhang *et al.*, “Robust Invisible Video Watermarking with Attention,” *arXiv*, 2019.
- [45] T. Sander *et al.*, “Watermark Anything with Localized Messages,” in *ICLR*, 2025, pp. 1–17.
- [46] C. K. Ho and C.-T. Li, “Semi-fragile watermarking scheme for authentication of JPEG images,” in *ITCC*, 2004, pp. 7–11.
- [47] H. Wang *et al.*, “CosFace: Large Margin Cosine Loss for Deep Face Recognition,” in *CVPR*, 2018, pp. 5265–5274.
- [48] W. Liu *et al.*, “SphereFace: Deep Hypersphere Embedding for Face Recognition,” in *CVPR*, 2017, pp. 212–220.

## VI. SUPPLEMENTARY APPENDIX

Appendices A, B, C, D, and E can be found in <https://github.com/bibek-cse/SAIW>.

## APPENDIX A

### ANALYSIS OF FEATURE-WISE LINEAR MODULATION FOR WATERMARK CONDITIONING

To investigate the role of feature-wise linear modulation ( $f_M$ ) in watermark conditioning, we analyze its behavior across the embedding network  $f_E$  (Section III-B). The results in Fig. A.1 provide quantitative and structural evidence that  $f_M$  enables discriminative, watermark-aware feature transformations.

Fig. A.1:(a) presents the gradient magnitude distribution across different layers of  $f_E$ , revealing a non-uniform sensitivity profile. The modulation impact is significantly stronger in deeper encoder and bottleneck layers (e.g.,  $f_E^{e3}$ ,  $f_E^{e4}$ ,  $f_F^b$ , and  $f_E^{d0}$ ), indicating that watermark conditioning is primarily governed by high-level semantic features. This behavior suggests that the embedding process leverages semantically meaningful regions, which contributes to improved robustness under complex distortions. Fig. A.1:(b) shows the mutual information matrix of the modulation parameters ( $\gamma, \beta$ ). The matrix exhibits strong diagonal dominance, with high self-information (approximately 2.3 nats) and comparatively lower cross-information (typically  $< 0.7$  nats). This indicates that each watermark identity induces a distinct parameter distribution, while retaining limited overlap that supports generalization across related sources. Fig. A.1:(c) illustrates a PCA projection of the learned modulation parameters. The resulting clusters are compact and well separated, with the first two principal components capturing 53.5% of the total variance. This demonstrates that  $f_M$  learns a structured latent space in which watermark identities are encoded as distinct yet semantically meaningful representations.

Importantly, these observations are consistent with the quantitative results reported in Table I (main paper). The ability of  $f_M$  to encode watermark identities into a structured and well-separated feature space directly translates into high robustness, with  $\mathcal{A}_{br} \geq 0.97$  under diverse perturbations. This indicates that SAIW learns distortion-invariant watermark representations rather than overfitting to specific attack types, thereby ensuring reliable recovery across unseen transformations. Compared to prior watermarking approaches that rely

on fixed embedding strategies or distortion-specific training (e.g., MBRS, StegaStamp, RivaGAN), the proposed modulation mechanism introduces a conditional and adaptive embedding paradigm. By dynamically modulating intermediate features based on watermark identity,  $f_M$  enables the network to jointly optimize imperceptibility and robustness, effectively shifting the fidelity–robustness trade-off. This is reflected in the simultaneous achievement of high visual quality (PSNR  $> 51$  dB, SSIM  $> 0.998$ , LPIPS  $< 0.001$ ) and strong recovery performance across multiple generators and attack scenarios. This analysis establishes that  $f_M$  is not merely an architectural component but a key mechanism that underpins the discriminative, robust, and generalizable behavior of the proposed watermarking framework.

## APPENDIX B

### MODULE ABLATION STUDY

To quantify the contribution of main architectural modules, we perform a systematic ablation study on CelebA [36] by individually removing the feature-wise linear modulation ( $f_M$ ), Swin transformer ( $\mathcal{S}$ ), and perceptual guidance block ( $f_P$ ). The results, summarized in Table B.1, reveal the distinct and complementary roles of these components in achieving robust and imperceptible watermarking. Ablating  $f_M$  leads to a substantial degradation in both imperceptibility and robustness, with PSNR dropping from 55.89 dB to 28.21 dB and  $\mathcal{A}_{sim}$  falling below 0.70 across most distortions. This confirms that  $f_M$  is critical for learning discriminative and structured watermark embeddings. Consistent with this observation, Fig. B.1 shows that without  $f_M$ , the residual maps  $\mathcal{I}_r$  exhibit weak spatial coherence and reduced localization, making source-specific encoding less reliable. The removal of the Swin Transformer ( $\mathcal{S}$ ) leads to a noticeable degradation in robustness (e.g., JPEG: 0.64, Gaussian noise: 0.60) and perceptual consistency (LPIPS), while PSNR and SSIM remain largely unaffected. This suggests that  $\mathcal{S}$  plays a crucial role in modeling long-range spatial dependencies and maintaining semantic coherence between the host image and the embedded watermark, which is essential for reliable watermark recovery




Fig. A.1: Feature-wise linear modulation analysis: (a) Gradient magnitude distribution across  $f_E$  layers, showing non-uniform and hierarchical embedding strength; (b) mutual information matrix of modulation parameters, indicating identity-specific parameterization; (c) PCA projection of modulation parameters, demonstrating structured, well-separated, and watermark-specific representations.






Fig. B.1: Visualization of source-conditioned watermark ( $\mathcal{I}_i$ ) embedding across different platform sources. The last row shows zoomed-in residual patterns of  $\mathcal{I}_r$  (with  $f_M$ ).

TABLE B.1: Module ablation analysis on CelebA [36]

| Metric                                                       | SAiW   | SAiW- $f_M$ | SAiW- $\mathcal{S}$ | SAiW- $f_P$ |
|--------------------------------------------------------------|--------|-------------|---------------------|-------------|
| <b>Imperceptibility</b>                                      |        |             |                     |             |
| PSNR $\uparrow$                                              | 55.89  | 28.21       | 59.65               | 59.34       |
| SSIM $\uparrow$                                              | 0.999  | 0.950       | 0.999               | 0.999       |
| LPIPS $\downarrow$                                           | 0.0016 | 0.0715      | 0.0005              | 0.0007      |
| <b>Robustness (<math>\mathcal{A}_{ssim} \uparrow</math>)</b> |        |             |                     |             |
| Pre-Attack                                                   | 0.99   | 0.67        | 0.82                | 0.96        |
| Brightness                                                   | 0.99   | 0.77        | 0.81                | 0.96        |
| Contrast                                                     | 0.99   | 0.69        | 0.83                | 0.96        |
| JPEG                                                         | 0.93   | 0.62        | 0.64                | 0.67        |
| Gaussian blur                                                | 0.99   | 0.67        | 0.82                | 0.96        |
| Gaussian noise                                               | 0.99   | 0.59        | 0.60                | 0.73        |

SAiW- $f_M$ , SAiW- $\mathcal{S}$ , SAiW- $f_P$  denote ablating feature-wise linear modulation, Swin transformer, perceptual guidance block, respectively.

under complex distortions. Similarly, ablating the perceptual guidance module ( $f_P$ ) leads to a noticeable drop in robustness under compression and noise (e.g., JPEG: 0.67, Gaussian noise: 0.73), while maintaining high perceptual quality. This suggests that  $f_P$  plays a key role in guiding watermark embedding into perceptually insensitive regions, thereby enhancing resilience to real-world distortions. Overall, SAIW demonstrates a favorable balance between imperceptibility (PSNR >55 dB, SSIM  $\geq 0.999$ , LPIPS <0.002) and robustness ( $\mathcal{A}_{ssim} \geq 0.93$  across distortions). These results suggest that the combined use of  $f_M$  for adaptive conditioning,  $\mathcal{S}$  for global spatial modeling, and  $f_P$  for perceptual masking contributes to learning distortion-invariant watermark representations, rather than relying on distortion-specific adaptations, thereby supporting consistent performance across diverse attack scenarios.

#### APPENDIX C

##### DISENTANGLING SEMANTICS AND STYLE UNDER SIMILAR WATERMARK INSTANCES

To evaluate the behavior of SAIW under challenging conditions of high watermark similarity, we consider a subset of watermark variants that are intentionally designed to be visually indistinguishable (Fig. C.1:(a)). These variants share the same semantic identity, resulting in highly correlated ID embeddings with cosine similarity approaching  $\sim 0.9$  (Fig. C.1:(b)). While this confirms that the identity representation effectively captures shared semantic content, it also highlights a key limitation: relying solely on identity features would collapse such variants

into an indistinguishable latent space, thereby hindering reliable watermark decoding. The proposed feature-wise modulation ( $f_M$ ) addresses this limitation by encoding complementary fine-grained stylistic information. As shown in the PCA projection (Fig. C.1:(c)), the  $f_M$  embeddings corresponding to visually similar watermark instances are well separated, despite their high semantic similarity. This demonstrates that  $f_M$  captures discriminative stylistic patterns that are intentionally abstracted away in the identity representation. As a result, even when watermark instances exhibit strong semantic overlap, their corresponding  $f_M$  representations remain distinct, enabling accurate decoding and reconstruction. This disentanglement between semantic identity and stylistic encoding is critical for robust watermark attribution, as it prevents representation collapse and ensures reliable identification of visually similar yet distinct watermark instances, further confirming that the conditioning mechanism supports identity-consistent yet instance-discriminative watermark representations.

#### APPENDIX D

##### IMPACT OF VARIOUS LOSS FUNCTION

We analyze the effect of margin-based loss functions on the geometry of the learned embedding space by comparing (a) ArcFace [32], (b) CosFace [47], and (c) SphereFace [48] (Fig. D.1). These losses impose different angular constraints, which directly influence intra-class compactness and inter-class separability. ArcFace enforces a strict additive angular margin, resulting in highly compact intra-class clusters and well-separated decision boundaries. CosFace introduces a cosine margin that provides moderate discrimination, but exhibits comparatively reduced compactness and weaker separation. In contrast, SphereFace applies multiplicative angular constraints, which lead to less stable optimization and noticeable inter-class overlap. These observations indicate that stronger and more stable angular margin formulations, such as ArcFace, significantly improve feature compactness and class separability. This directly contributes to learning more discriminative and robust representations, which is critical for reliable watermark identification and source attribution under challenging conditions. This highlights the importance of margin design in




Fig. C.1: Analysis of visually similar watermark instances: (a) watermark variants with high visual similarity; (b) ID embedding similarity showing strong semantic correlation; (c) PCA projection of  $f_M$  parameters demonstrating clear separation despite high semantic similarity, indicating effective stylistic disentanglement.




Fig. D.1: UMAP visualization of embedding distributions under different margin-based loss functions: (a) ArcFace, (b) CosFace, and (c) SphereFace.

|          | Input image ( $\mathcal{I}$ ) | Perceptual map ( $\mathcal{P}_n$ ) | Residual map ( $\mathcal{I}_r$ ) | Watermarked image ( $\mathcal{I}_w$ ) |
|----------|-------------------------------|------------------------------------|----------------------------------|---------------------------------------|
| Sample-1 |                               |                                    |                                  |                                       |
| Sample-2 |                               |                                    |                                  |                                       |
| Sample-3 |                               |                                    |                                  |                                       |

Fig. E.1: Texture-adaptive embedding: perceptual maps  $\mathcal{P}_n$  and residuals  $\mathcal{I}_r$  across varying texture levels in input images.

shaping discriminative embedding spaces for robust watermark decoding.

## APPENDIX E

### IMPACT OF INPUT IMAGE TEXTURE ON $\mathcal{I}_r$

To analyze the effect of input image texture on watermark embedding, we examine the relationship between the perceptual

map  $\mathcal{P}_n$  and the corresponding residual patterns  $\mathcal{I}_r$  generated by the encoder  $f_E$ . The same watermark logo  $\mathcal{I}_l$  is consistently used across all three samples from IndicSideFace [33] presented in Fig. E.1, enabling a controlled analysis of the impact of input texture on embedding behavior. The perceptual map guides the allocation of embedding strength based on local visual sensitivity, allowing adaptive watermark placement. Sample-1, with relatively low texture variation (standard deviation: 33.35), produces a watermarked image  $\mathcal{I}_w$  with PSNR 51.65 dB and SSIM 0.9970. As texture complexity increases in Sample-2 (standard deviation: 62.17), the model allocates stronger embedding capacity to textured regions, resulting in improved perceptual quality (PSNR 52.64 dB, SSIM 0.9996). In Sample-3, which exhibits the highest texture variability (standard deviation: 95.88), the encoder distributes residuals more aggressively in high-texture regions, maintaining high structural similarity (SSIM 0.9992) despite a moderate reduction in PSNR (48.76 dB). These observations indicate that  $f_E$ , guided by  $\mathcal{P}_n$ , learns a texture-adaptive embedding strategy in which watermark perturbations are preferentially concentrated in perceptually tolerant regions, enabling effective concealment while preserving visual fidelity. This behavior supports consistent imperceptibility across varying texture

conditions and contributes to robust watermark recovery under diverse real-world scenarios.