All in One: Unifying Deepfake Detection, Tampering Localization, and Source Tracing with a Robust Landmark-Identity Watermark

Supplementary Material

A. Scalability Analysis

Since the 136-D landmark vector W_L is already sufficient for detection and localization, we assess LIDMark's scalability by expanding the identifier payload W_{ID} from 16-D to 32-D. While a 16-D payload supports 2^{16} unique identities and is sufficient for most scenarios, a 32-D payload accommodates $2^{32} \approx 4.3 \times 10^9$ identities to cater to large-scale applications. This expansion increases the total payload to 168-D, comprising the 136-D landmark and the new 32-D identifier. While W_{ID} can represent arbitrary provenance or copyright data, our experiments, for convenience, use the hashed filename to simulate this tracing scenario. Results are detailed in Tabs. 1–3 and Fig. 1. The framework accommodates a 10.5% increase in total payload capacity with only a marginal PSNR cost, while yielding a more stable and robust recovery of the geometric landmark signal against both common distortions and advanced deepfake manipulations. Notably, the expanded identifier creates no interference with the geometric signal, confirming that our factorized decoding strategy effectively disentangles these two heterogeneous tasks. This validates the potential to expand W_{ID} for more complex or granular source tracing.

Table 1. Objective imperceptibility comparison on CelebA-HQ [7]. Results are at 256×256 resolution. We compare the standard payload (W.L. 152) against the expanded (W.L. 168).

W.L.	Image Size	PSNR ↑	SSIM ↑
152	256×256	44.31	0.99
168	256×256	42.36	0.99

Table 2. Quantitative comparison on CelebA-HQ [7] of BER and AED for LIDMark with different capacities (W.L. 152 vs. 168) under common distortions $\mathcal{M}_c$, demonstrating highly robust scalability.

Distortion	**W.L. 152**		**W.L. 168**	
	BER ↓	AED ↓	BER ↓	AED ↓
Identity	0.00%	2.96	0.00%	2.14
Resize	0.00%	3.00	0.00%	2.17
GausBlur	0.00%	3.03	0.00%	2.16
MedBlur	0.00%	3.01	0.02%	2.16
JpegTest	0.57%	3.39	0.00%	2.36
JpegMask	0.01%	2.98	0.01%	2.17
Average	**0.10%**	**3.06**	**0.01%**	**2.19**

Table 3. Quantitative comparison on CelebA-HQ [7] of BER and AED for LIDMark with different capacities (W.L. 152 vs. 168) under deepfake manipulations $\mathcal{M}_d$, demonstrating highly robust scalability.

Manipulation	**W.L. 152**		**W.L. 168**	
	BER ↓	AED ↓	BER ↓	AED ↓
SimSwap [1]	0.97%	3.55	0.59%	2.58
UniFace [9]	2.44%	4.01	0.18%	2.47
CSCS [6]	0.01%	3.05	0.00%	2.21
StarGAN-v2 [2]	8.47%	5.51	14.33%	5.32
InfoSwap [3]	0.84%	3.91	1.19%	2.48
Average	**2.55%**	**4.01**	**3.26%**	**3.01**

Table 4. Comparison of model complexity against dual-function baselines. We report the number of Parameters (M) and FLOPs (G). For fair comparison, the watermark length (W.L.) is standardized to 152-D for all methods.

Model	Image Size	Params (M)	FLOPs (G)
SepMark [8]	128×128	28.310	4.031
WaveGuard [5]	128×128	34.810	27.014
KAD-NET [4]	128×128	8.542	7.835
Ours	128×128	5.815	12.221
SepMark [8]	256×256	28.310	16.059
WaveGuard [5]	256×256	34.810	27.011
KAD-NET [4]	256×256	48.542	31.357
Ours	256×256	21.023	60.622

B. Complexity Analysis

Tab. 4 presents a comprehensive benchmarking of model complexity, contrasting LIDMark with existing dual-function methods. In terms of parameter efficiency, LIDMark demonstrates superior compactness. This is largely attributed to our unified design, where a single shared backbone efficiently extracts rich and versatile features for multiple downstream tasks, avoiding the structural redundancy of dual-branch encoders. At a resolution of 128×128, our model requires only 5.82M parameters, representing approximately 20% and 17% of the capacity of SepMark [8] and WaveGuard [5], respectively. This lightweight footprint minimizes storage requirements, making LIDMark exceptionally suitable for deployment on storage-constrained edge devices and mobile platforms.

Regarding computational cost, we observe that LIDMark incurs higher FLOPs. This increased computational den-

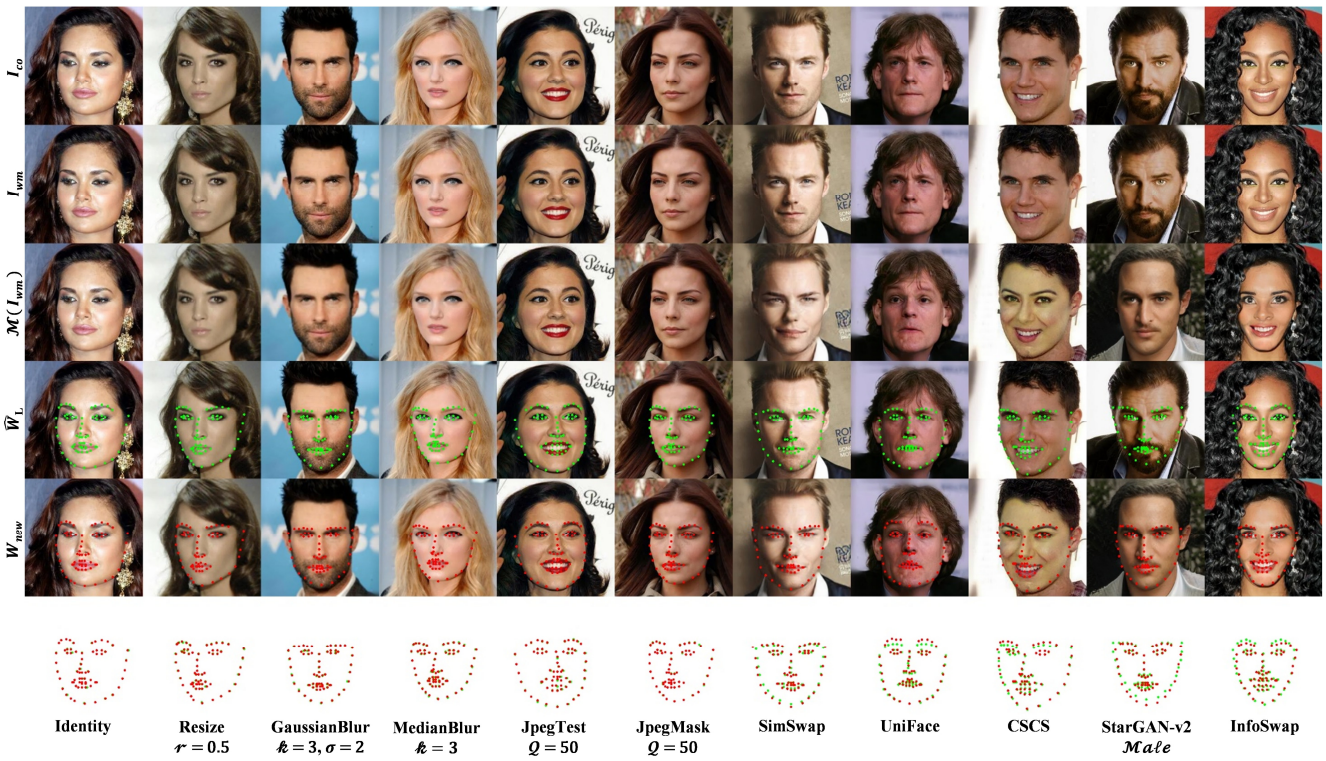

Figure 1. Visual assessment of the 168-D LIDMark robustness and imperceptibility. Comparing rows 1 and 2 shows the watermarked image I_{wm} is indistinguishable from the cover image I_{co}. Row 3 displays the manipulation results $\mathcal{M}(I_{wm})$. The "intrinsic-extrinsic" consistency check compares the green dots in row 4, representing FHD-recovered intrinsic landmarks $\hat{W}_L$, against the red dots in row 5, representing re-detected extrinsic landmarks W_{new}. Row 6 merges these two landmark sets into a combined image to visualize their spatial differences.

sity is a deliberate and necessary investment to achieve trifunctional capability. Unlike baselines that are limited to coarse-grained detection or tracing, our FHD performs dense, pixel-level feature processing to ensure precise tampering localization and high-fidelity watermark recovery. While methods like SepMark achieve lower FLOPs, they sacrifice the granularity required for localization and the robustness needed for challenging deepfake attacks. Consequently, we argue that the marginal increase in inference latency is a worthy exchange for the significant gains in forensic security and the unification of three distinct tasks within a single framework.

References

[1] Renwang Chen, Xuanhong Chen, Bingbing Ni, and Yanhao Ge. Simswap: An efficient framework for high fidelity face swapping. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2003–2011, 2020. 1

[2] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. 1

[3] Gege Gao, Huaibo Huang, Chaoyou Fu, Zhaoyang Li, and Ran He. Information bottleneck disentanglement for identity swapping. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3404–3413, 2021. 1

[4] Sijia He, Yunfeng Diao, Yongming Li, Chen Sun, Liejun Wang, and Zhiqing Guo. Kad-net: Kolmogorov-arnold and differential-aware networks for robust and sensitive proactive deepfake forensics. *Knowledge-Based Systems*, page 114692, 2025. 1

[5] Ziyuan He, Zhiqing Guo, Liejun Wang, Gaobo Yang, Yunfeng Diao, and Dan Ma. Waveguard: Robust deepfake detection and source tracing via dual-tree complex wavelet and graph neural networks. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2025. 1

[6] Ziyao Huang, Fan Tang, Yong Zhang, Juan Cao, Chengyu Li, Sheng Tang, Jintao Li, and Tong-Yee Lee. Identity-preserving face swapping via dual surrogate generative models. *ACM Transactions on Graphics*, 43(5):1–19, 2024. 1

[7] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 1

[8] Xiaoshuai Wu, Xin Liao, and Bo Ou. Sepmark: Deep separable watermarking for unified source tracing and deepfake detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1190–1201, 2023. 1

[9] Chao Xu, Jiangning Zhang, Yue Han, Guanzhong Tian, Xianfang Zeng, Ying Tai, Yabiao Wang, Chengjie Wang, and Yong Liu. Designing one unified framework for high-fidelity face reenactment and swapping. In *European conference on computer vision*, pages 54–71. Springer, 2022. 1