AGTRL: AXIS-GUIDED TRUSTED REPRESENTATION LEARNING FOR PASSIVE HIDDEN-SCREEN RECOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

When a screen is completely hidden from direct view, its emitted light can still diffuse onto a nearby wall and encode information about the displayed content. We investigate whether an ordinary camera can recover that content solely by observing these wall light traces. Propagation and diffuse reflection compress screen appearance into low-contrast brightness and color variations while mixing contributions from different screen regions. These screen-related traces are further overlaid with wall texture, shadows, exposure variation, and sensor noise. We propose Axis-Guided Trusted Representation Learning (AGTRL) to recover hidden-screen content by learning both wall-to-screen correspondence and trace reliability. Axis-Guided Representation Learning (AGRL) organizes dispersed wall features along channel and spatial axes to form a screen-oriented representation. Embedded TraceTrust arbitration combines local continuity, artifact suspicion, and a learned gate to prioritize credible screen traces and suppress pseudo traces during reconstruction. We conduct extensive experiments across four screen-content families and under changes in distance, brightness, viewpoint, noise, occlusion, and wall appearance. AGTRL outperforms 20 mainstream image reconstruction methods in all 12 dataset–metric combinations and robustly recovers the principal color distributions, region positions, page layouts, and structural contours of hidden screens. These results demonstrate the initial feasibility of observing and recovering hidden-screen content using an ordinary camera aimed only at wall light traces.

1 INTRODUCTION

A screen can be hidden from a camera while its light remains visible on a nearby wall. Can an ordinary camera recover the displayed content solely from these wall light traces? Kuhn (2002) demonstrated hidden CRT-content recovery from wall-reflected light using a photomultiplier and oscilloscope sampling at 250 MHz. With paired wall–screen training, we investigate recovery of the screen's principal color distributions, region positions, page layouts, and structural contours from RGB wall observations captured by an ordinary camera.

Diffuse propagation spreads light from each screen region over the wall, while contributions from different regions overlap at the observed locations (Saunders et al., 2019; Geng et al., 2022). The resulting wall pattern carries screen appearance as weak brightness and color variations. Wall texture, shadows, exposure changes, and sensor noise can obscure these variations. Recovery therefore requires two connected capabilities: learning how mixed wall features correspond to screen structure, and identifying the traces that should guide reconstruction. Figure 1 illustrates this wall-to-screen observation setting.

To address these challenges, we propose Axis-Guided Trusted Representation Learning (AGTRL), a unified encoder–decoder framework for passive hidden-screen recovery. Axis-Guided Representation Learning (AGRL) first organizes the channel responses and spatial relationships of dispersed wall features to learn their correspondence with screen structure. TraceTrust arbitration then combines local continuity, artifact suspicion, and a learned gate to distinguish credible screen traces from pseudo traces and weight their contribution to reconstruction. Together, the two designs turn weak, mixed wall evidence into a representation that supports the recovery of screen colors, region positions, page layouts, and structural contours. Across four screen-content families, AGTRL outperforms 20 reconstruction methods in all 12 dataset–metric combinations and recovers the principal color distributions, region positions, page layouts, and structural contours of hidden screens. These results establish the initial feasibility of recovering screen content with an ordinary camera observing only wall light.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

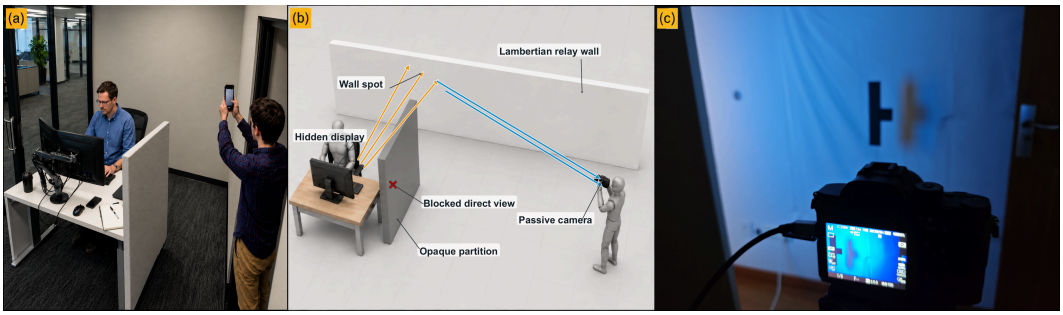

Figure 1: Hidden-screen observation through wall light. (a) Indoor observation scenario. (b) Schematic of light propagation from the screen to the wall and camera. (c) Screen2000 acquisition setup. The screen is hidden from the acquisition camera, while its emitted light forms a visible pattern on the wall.

The main contributions are summarized as follows.

(1) We formulate passive hidden-screen recovery as reconstructing display content solely from an ordinary-camera observation of wall light traces, and propose AGTRL as an end-to-end solution to this task.

(2) We introduce AGRL to learn the correspondence between dispersed wall evidence and screen structure, together with embedded TraceTrust arbitration to distinguish credible screen traces from pseudo traces and form a trusted latent update.

(3) We validate AGTRL on four screen-content families against 20 reconstruction methods, with component ablations, a local-masking analysis, and robustness tests across distance, brightness, viewpoint, noise, occlusion, and wall-appearance conditions.

2 Related Work

Indirect light can preserve visual information about a hidden source. Kuhn (2002) recovered CRT display content from temporally resolved optical emanations after diffuse reflection, establishing a direct connection between emitted screen light and recoverable information. Active imaging subsequently recovered hidden geometry from ultrafast time-of-flight measurements (Velten et al., 2012). Liu et al. (2023) extended time-resolved reconstruction to irregular illumination and detection patterns through joint signal–object regularization, recovering hidden-object albedo and surface normals. Ordinary-camera approaches exploited spatial modulation: Bouman et al. (2017) reconstructed angular videos from corner penumbrae, and Saunders et al. (2019) recovered two-dimensional hidden scenes through computational periscopy. Learning-based methods further connected indirect observations to complex images through manifold embedding and optimal transport (Geng et al., 2022), and through light-transport-conditioned reconstruction (Zhang et al., 2025). Recent hidden-screen work couples physically guided irradiance inversion with irreversibility-constrained semantic reprojection (Zheng et al., 2026). AGTRL develops a representation-learning route that jointly learns wall-to-screen correspondence and trace weighting. Channel and spatial processing organize dispersed wall features to represent screen structure, and arbitration based on local continuity and artifact suspicion weights their contribution to the recovered screen.

Learned image restoration provides complementary mechanisms for organizing degraded visual evidence. DnCNN learns residual corrections for image recovery (Zhang et al., 2017), while Uformer and Restormer use hierarchical spatial attention and channel-oriented attention to model image relationships (Wang et al., 2022; Zamir et al., 2022). Adaptive sparse attention further refines the interactions used for restoration (Zhou et al., 2024). Feature weighting offers another route: SE recalibrates channel responses, and CBAM combines channel and spatial attention (Hu et al., 2018; Woo et al., 2018). Non-local blocks aggregate distant features, while FFC combines local processing with spectral interactions (Wang et al., 2018; Chi et al., 2020). PromptIR and NDR-Restore use prompts and learned degradation representations to guide restoration (Potlapalli et al., 2023; Yao et al., 2024); UniProcessor and OneRestore further condition processing on degradation descriptions and scene attributes (Duan et al., 2025; Guo et al., 2025). Learned score priors provide a complementary route to inverse reconstruction through generative modeling and posterior sampling (Song et al., 2021; Chung et al., 2023). For wall-to-screen recovery, AGTRL couples feature organization with evidence selection: AGRL learns how mixed wall features correspond to screen structure, and TraceTrust

combines local continuity, artifact suspicion, and a learned gate to weight the reconstruction update before bottleneck integration.

3 METHOD

3.1 OVERALL ARCHITECTURE

AGTRL reconstructs the hidden RGB screen from a wall observation. The encoder first extracts a multi-scale wall representation. AGRL forms a candidate update from the bottleneck features, and TraceTrust weights this update using local continuity, artifact suspicion, and a learned gate. The resulting representation passes through the middle transformation and skip-connected decoder to produce an RGB correction d. The final reconstruction is $\widehat{y} = x + d$. Figure 2 illustrates the architecture.

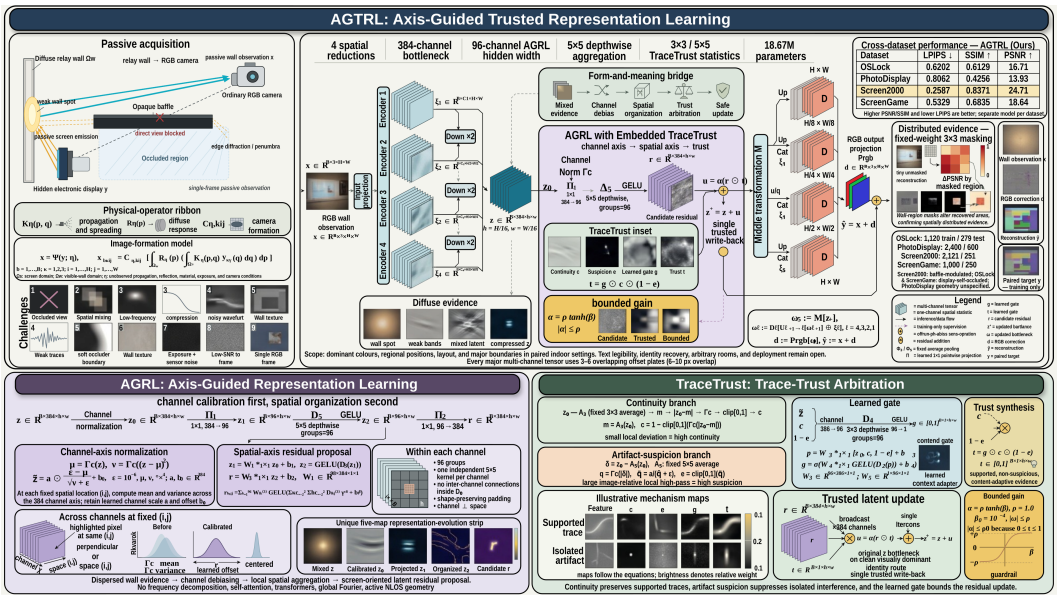

Figure 2: AGTRL encodes a wall observation, forms a screen-oriented bottleneck residual through AGRL and embedded TraceTrust arbitration, and decodes an RGB residual added to the input.

Let $y \in \mathbb{R}^{B\times 3\times H\times W}$ denote the paired hidden-screen target and $x \in \mathbb{R}^{B\times 3\times H\times W}$ the wall observation. Their unknown image-formation relation is written as

$$x = \Psi(y;\eta).$$

The operator Ψ combines propagation from the screen to the visible wall, diffuse reflection, and camera image formation. The condition variable η collects the propagation, surface, and camera parameters. Aittala et al. (2019) recovered hidden videos by jointly fitting a light-transport matrix and a hidden-video representation to the observed illumination sequence through deep matrix factorization. AGTRL learns wall-to-screen correspondence and trace weighting from paired wall observations and screen targets during training, then uses its learned parameters to reconstruct the screen with the RGB wall observation x as its sole inference input.

The four-level encoder maps $v^{(0)} = x$ to a skip feature and a reduced feature at each scale. With learned coefficients $E^{(\ell)}$ and $Q^{(\ell)}$, the indexed transformation is

$$\xi_{\ell,\mathfrak{b}\chi ij} = \varphi_\ell\left(\sum_{\substack{\chi'=1,\dots,C_{\ell-1}\\(\delta_i,\delta_j)\in\mathcal{N}_\ell^E}} E^{(\ell)}_{\chi\chi'\delta_i\delta_j}\, v^{(\ell-1)}_{\mathfrak{b}\chi',\, i+\delta_i,\, j+\delta_j} + \zeta^{(\ell)}_{\chi}\right),$$

$$v^{(\ell)}_{\mathfrak{b}\chi ij} = \sum_{\substack{\chi'=1,\dots,C_\ell\\(\delta_i,\delta_j)\in\mathcal{N}_\ell^Q}} Q^{(\ell)}_{\chi\chi'\delta_i\delta_j}\, \xi_{\ell,\mathfrak{b}\chi',\, 2i+\delta_i,\, 2j+\delta_j} + \upsilon^{(\ell)}_{\chi},$$

$$z := v^{(4)} \in \mathbb{R}^{B\times C\times h\times w}, \qquad \ell = 1,\dots,4.$$

Here, χ, χ' are channel indices, (δ_i, δ_j) are offsets in $\mathcal{N}_\ell^E$ and $\mathcal{N}_\ell^Q$, and φ_ℓ is the level-wise nonlinear transformation. The tensors $\xi_1, \ldots, \xi_4$ form the skip path, while z is the bottleneck supplied to AGRL.

3.2 AXIS-GUIDED REPRESENTATION LEARNING

In passive hidden-screen scenarios, propagation-induced mixing obscures correspondence between wall features and screen structure, while conventional local extraction tends to follow dominant wall appearance. AGRL therefore forms a spatially organized residual proposal through channel-axis normalization followed by spatial-axis neighborhood aggregation. The proposal enters the bottleneck through a trust-weighted residual path.

AGRL first centers and rescales the channel responses at each location, then applies a learned channel scale and offset:

$$\mu_{\mathfrak{b}ij} := \Gamma_c(z)_{\mathfrak{b}ij} = \frac{1}{C}\sum_{\chi=1}^{C} z_{\mathfrak{b}\chi ij},$$

$$\nu_{\mathfrak{b}ij} := \Gamma_c\big((z-\mu)^2\big)_{\mathfrak{b}ij} = \frac{1}{C}\sum_{\chi=1}^{C}\left(z_{\mathfrak{b}\chi ij} - \frac{1}{C}\sum_{\chi'=1}^{C} z_{\mathfrak{b}\chi' ij}\right)^2,$$

$$\widetilde{z}_{\mathfrak{b}\chi ij} := a_\chi \frac{z_{\mathfrak{b}\chi ij} - \mu_{\mathfrak{b}ij}}{\sqrt{\nu_{\mathfrak{b}ij} + \epsilon}} + b_{0,\chi}.$$

Here, Γ_c is the channel mean, $\mu, \nu \in \mathbb{R}^{B\times 1\times h\times w}$, $a, b_0 \in \mathbb{R}^C$ are learned affine parameters, and $\epsilon > 0$ stabilizes normalization.

The spatial-axis stage converts $\widetilde{z}$ into a candidate residual through the following projection–aggregation–projection chain:

$$r_{\mathfrak{b}\chi ij} := [\Pi_2(\mathrm{GELU}(\Delta_5(\Pi_1(\widetilde{z}))))]_{\mathfrak{b}\chi ij}$$
$$= \sum_{k=1}^{K} W_{\chi k}^{(2)}\ \mathrm{GELU}\left(\sum_{\delta_i,\delta_j=-2}^{2} D_{k\delta_i\delta_j}^{(5)}\left[\sum_{\chi'=1}^{C} W_{k\chi'}^{(1)}\widetilde{z}_{\mathfrak{b}\chi',\, i+\delta_i,\, j+\delta_j} + b_k^{(1)}\right] + \gamma_k^{(5)}\right) + b_\chi^{(2)}.$$

Here, Π_1 reduces C channels to K, Δ_5 is 5×5 depthwise spatial aggregation, and Π_2 returns the result to C channels. Thus, r is a latent proposal with the same shape as z, which embedded TraceTrust arbitrates before bottleneck integration. We use $C = 384$, $K = 96$, and a 5×5 spatial neighborhood.

During training, gradients from the final reconstruction objective train AGRL to relate wall features to screen structure. At inference, its input is the bottleneck derived from the current wall image. Channel normalization calibrates feature responses at each location. Spatial aggregation then combines neighboring evidence into the candidate update r. TraceTrust determines how strongly this update contributes to the screen representation.

3.3 TRACETRUST ARBITRATION

TraceTrust distinguishes credible screen traces from pseudo traces by evaluating the local support for the candidate bottleneck update. Continuity measures agreement with neighboring features, while artifact suspicion compares local deviations with the current representation's image-wide average deviation. A learned gate combines these cues with the normalized wall features to form a single-channel trust map. This map weights the candidate residual before bottleneck integration, allowing spatially supported screen evidence to guide reconstruction when wall texture, shadows, and sensor noise affect the observation.

Local continuity measures agreement with the neighboring representation:

$$c = 1 - \mathrm{clip}_{[0,1]}(\Gamma_c(|\widetilde{z} - \Phi_3(\widetilde{z})|))\,.$$

Artifact suspicion compares a wider local deviation with the current image's average deviation:

$$q = \Gamma_c(|\widetilde{z} - \Phi_5(\widetilde{z})|)\,, \qquad e = \mathrm{clip}_{[0,1]}\left(\frac{q}{\Gamma_s(q) + \epsilon}\right).$$

Here, Φ_3 and Φ_5 are local averages, and Γ_c and Γ_s average over channels and spatial positions, respectively. Together, the two statistics provide local support and an image-specific reference for trust estimation. Appendix Equation (A3) gives their indexed formulation.

TraceTrust next fuses the normalized feature, continuity, and complementary suspicion through channel-axis concatenation $\oplus_a$, pointwise projections Π_3, Π_4, and 3×3 depthwise aggregation Δ_3:

$$\varpi_{\mathfrak{b}kij} := \sum_{\lambda=1}^{C+2} W^{(3)}_{k\lambda} \left[\widetilde{z} \oplus_a c \oplus_a (1-e)\right]_{\mathfrak{b}\lambda ij} + b^{(3)}_k,$$

$$g_{\mathfrak{b}ij} := \sigma\left(\sum_{k=1}^{K} W^{(4)}_k \ \mathrm{GELU}\left(\sum_{\delta_i=-1}^{1}\sum_{\delta_j=-1}^{1} D^{(3)}_{k\delta_i\delta_j}\, \varpi_{\mathfrak{b}k,\, i+\delta_i,\, j+\delta_j} + \gamma^{(3)}_k\right) + b^{(4)}\right),$$

$$t_{\mathfrak{b}ij} := g_{\mathfrak{b}ij}\, c_{\mathfrak{b}ij}\big(1 - e_{\mathfrak{b}ij}\big) \in [0,1].$$

The learned gate g adjusts the joint weighting of the deterministic factors c and $1-e$ for the current representation, and t is the final TraceTrust weight.

The trust map is finally broadcast over the C proposal channels. With learned scalar β and fixed bound ρ, the trusted update is integrated as

$$\begin{aligned}
\alpha &:= \rho \tanh(\beta), \qquad |\alpha| \le |\rho|,\\
u_{\mathfrak{b}\chi ij} &:= \alpha\, r_{\mathfrak{b}\chi ij}\, t_{\mathfrak{b}ij} = \rho \tanh(\beta)\, r_{\mathfrak{b}\chi ij}\, g_{\mathfrak{b}ij}\, c_{\mathfrak{b}ij}\big(1 - e_{\mathfrak{b}ij}\big),\\
z^{+}_{\mathfrak{b}\chi ij} &:= z_{\mathfrak{b}\chi ij} + u_{\mathfrak{b}\chi ij} = z_{\mathfrak{b}\chi ij} + \rho \tanh(\beta)\, r_{\mathfrak{b}\chi ij}\, g_{\mathfrak{b}ij}\, c_{\mathfrak{b}ij}\big(1 - e_{\mathfrak{b}ij}\big).
\end{aligned}$$

Here, $u, z^{+} \in \mathbb{R}^{B\times C\times h\times w}$. The identity path carries z forward, while αt controls the contribution of the candidate update r. With $\rho = 1$, the trusted contribution is bounded in magnitude by the proposal. This allows the representation to incorporate spatially supported changes through the bounded, trust-weighted update u.

The updated bottleneck z^{+} passes through the middle transformation and the skip-connected decoder. At each scale, the decoder combines its upsampled state with the corresponding encoder feature ξ_ℓ. The RGB projection produces d, and the final recovered screen is $\widehat{y} = x + d$. Appendix Equation (A6) gives the indexed formulation.

Training uses paired mean absolute reconstruction error. For a batch of B RGB image pairs with spatial size $H \times W$, the objective is

$$\mathcal{L}(\theta) = \frac{1}{3BHW} \sum_{\mathfrak{b}=1}^{B}\sum_{\kappa=1}^{3}\sum_{i=1}^{H}\sum_{j=1}^{W} \left|\widehat{y}_{\mathfrak{b}\kappa ij} - y_{\mathfrak{b}\kappa ij}\right|, \qquad \theta \in \Theta.$$

Here, $\mathfrak{b}$, κ, i, and j index the batch, RGB channel, row, and column, respectively. The normalization factor $3BHW$ averages the absolute reconstruction error over all samples, channels, and pixels. The parameter set θ contains all learned coefficients above. Gradients from $\mathcal{L}$ train the complete encoder–AGRL–TraceTrust–decoder path. Because the target is the displayed screen, the reconstruction gradients guide both AGRL's feature proposal and TraceTrust's learned gate toward screen-content recovery. At inference, the network computes $\widehat{y}$ from x with fixed parameters.

4 Experiments and Results

4.1 Experimental setup

We use fixed, non-overlapping paired splits for OSLock (1,120 training and 279 test pairs), PhotoDisplay (2,400 and 600), Screen2000 (2,121 and 251), and ScreenGame (1,000 and 250). A separate model is trained for each dataset. Each sample pairs one RGB wall observation with its simultaneous RGB screen target, and the input and output share the same spatial dimensions.

Computational periscopy recovered hidden two-dimensional scenes from an occluder's penumbra (Saunders et al., 2019), and two-edge-resolved imaging used orthogonal doorway edges for three-dimensional recovery (Czajkowski & Murray-Bruce, 2024). For hidden-screen recovery, we evaluate AGTRL's reconstruction of screen structure in both the baffle-modulated configuration of Screen2000 (Figure 1) and the display-self-occluded configuration of OSLock, PhotoDisplay, and ScreenGame, where the display body itself blocks direct screen viewing while the camera observes the surrounding wall light. Together, the four datasets cover lock-screen interfaces, natural images, desktop and window layouts, and game content, linking complementary acquisition geometries to diverse screen structures.

We compare AGTRL with 20 reconstruction methods using the paired test images and the three reconstruction metrics. The ablations vary the designated component, and the robustness experiments evaluate fixed learned models as the wall observation changes. For AGTRL, the main comparison and component ablation share dataset-specific checkpoints; the robustness sweeps use a separate fixed set. On OSLock, PhotoDisplay, and ScreenGame, these sweeps reuse the main comparison's test samples while varying the wall input. Peak signal-to-noise ratio (PSNR) (Huynh-Thu & Ghanbari, 2008), structural similarity index (SSIM) (Wang et al., 2004), and learned perceptual image patch similarity

(LPIPS) (Zhang et al., 2018) measure pixel fidelity, structural agreement, and perceptual difference, respectively. Higher PSNR and SSIM and lower LPIPS indicate better performance.

Table 1: Comparison across four screen-content datasets. PSNR is reported in dB; higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	Source	OSLock			PhotoDisplay			Screen2000			ScreenGame		
		PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor (Duan et al., 2025)	ECCV'24	**16.60**	0.6057	**0.6307**	13.36	0.4068	0.8116	**23.44**	**0.7990**	**0.3261**	**18.45**	0.6725	0.5514
OneRestore (Guo et al., 2025)	ECCV'24	12.24	0.5221	0.7212	11.11	0.3756	0.8725	10.41	0.4642	0.7142	10.31	0.5062	0.7474
DeepSN-Net (Deng et al., 2025)	TPAMI'25	12.84	0.4228	0.7711	12.22	0.3173	0.8817	10.49	0.2954	0.7787	10.87	0.4440	0.7390
PromptIR (Potlapalli et al., 2023)	NeurIPS'23	15.34	0.5936	0.6374	12.78	0.3989	**0.8066**	21.56	0.7552	0.3804	17.78	0.6599	0.5708
Uformer (Wang et al., 2022)	CVPR'22	15.15	0.5652	0.6743	12.10	0.3802	0.8338	19.20	0.7068	0.4639	17.55	0.6470	0.6142
MIRNetv2 (Zamir et al., 2023)	TPAMI'23	15.21	0.6042	0.6575	11.98	0.4086	0.8436	18.24	0.6917	0.4923	17.68	0.6559	0.6120
OKNet (Cui et al., 2024b)	AAAI'24	14.66	0.5553	0.6781	12.58	0.3944	0.8316	20.58	0.7384	0.4445	17.89	0.6625	0.5756
AST (Zhou et al., 2024)	CVPR'24	15.07	0.5366	0.6920	12.33	0.3622	0.8210	19.39	0.7050	0.4668	17.46	0.6453	0.6158
ConvIR (Cui et al., 2024a)	TPAMI'24	14.91	0.5925	0.6867	11.69	0.3904	0.8447	20.27	0.7283	0.4495	17.64	0.6515	0.6013
DnCNN (Zhang et al., 2017)	TIP'17	11.82	0.4787	0.7221	10.29	0.3171	0.8386	8.88	0.3466	0.7245	8.74	0.4275	0.7258
U-Net (Ronneberger et al., 2015)	MICCAI'15	12.18	0.4874	0.7618	10.82	0.3243	0.8533	9.18	0.2401	0.8061	10.29	0.4715	0.7581
DExNet (Huang et al., 2025)	TPAMI'25	15.52	0.6084	0.6520	13.27	**0.4193**	0.8189	20.00	0.7323	0.4382	17.63	0.6586	0.5954
HVI-CIDNet (Yan et al., 2025)	CVPR'25	15.61	0.5758	0.6714	13.25	0.4055	0.8119	19.69	0.7178	0.4724	17.77	0.6555	0.5979
NDR-Restore (Yao et al., 2024)	TIP'24	15.32	0.5954	0.6592	12.00	0.3929	0.8440	22.04	0.7812	0.3637	18.34	**0.6795**	**0.5332**
DehazeFormer (Song et al., 2023)	TIP'23	14.87	0.5375	0.6687	12.48	0.3589	0.8230	19.37	0.6931	0.4794	17.77	0.6494	0.6074
Ada-Deblur (Wu et al., 2023)	TIP'23	16.27	**0.6086**	0.6608	**13.49**	0.4145	0.8223	19.43	0.7067	0.4637	17.75	0.6577	0.5878
DEA-Net (Chen et al., 2024)	TIP'24	15.96	0.6069	0.6742	13.16	0.4076	0.8398	20.65	0.7500	0.4240	17.73	0.6579	0.5879
IRNeXt (Cui et al., 2023)	ICML'23	16.13	0.6082	0.6785	12.07	0.3816	0.8499	19.96	0.7234	0.4542	18.14	0.6681	0.5735
SFHformer (Jiang et al., 2025)	ECCV'24	12.46	0.5095	0.7404	11.61	0.3768	0.8867	8.82	0.3940	0.8066	10.01	0.4620	0.7598
DarkIR (Feijoo et al., 2025)	CVPR'25	8.79	0.0949	0.8040	6.37	0.0522	0.9073	9.22	0.1749	0.8673	4.46	0.0409	0.9361
AGTRL	Ours	**16.71**	**0.6129**	**0.6202**	**13.93**	**0.4256**	**0.8062**	**24.71**	**0.8371**	**0.2587**	**18.64**	**0.6835**	**0.5329**

AGTRL uses four spatial reductions, a 384-channel bottleneck, and an AGRL hidden projection width of 96 channels. It uses channelwise normalization, pointwise projections, 5×5 depthwise aggregation, 3×3 and 5×5 TraceTrust averages, and an RGB residual readout. Training uses mean absolute error and AdamW (Loshchilov & Hutter, 2019); the learning rate and weight decay are both 10^{-4}, and the maximum gradient norm is 1.0. The learnable residual parameter β is initialized to 10^{-4}, while the fixed residual-scale bound is set to $\rho = 1.0$.

4.2 COMPARISON WITH STATE OF THE ART

Table 1 compares AGTRL with 20 reconstruction methods across four screen-content families.

AGTRL achieves the best results in all 12 dataset–metric combinations across four screen-content families (Table 1). PromptIR guides restoration with degradation-specific prompts, while UniProcessor uses text-conditioned processing (Potlapalli et al., 2023; Duan et al., 2025); in our Screen2000 comparison, they achieve 21.56 and 23.44 dB PSNR, respectively (Table 1). AGTRL couples wall-to-screen correspondence with trace reliability and achieves 24.71 dB PSNR, exceeding UniProcessor by 1.27 dB while increasing SSIM from 0.7990 to 0.8371 and reducing LPIPS from 0.3261 to 0.2587. These gains accompany the recovery of window placement, large color regions, and major boundaries in Figure 3.

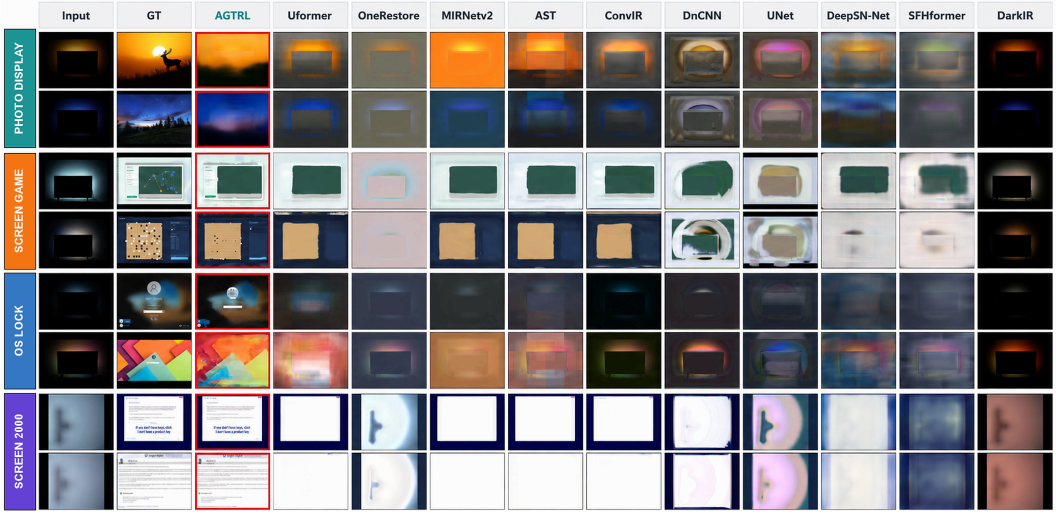

Figure 3: Qualitative comparison on two fixed examples from each of PhotoDisplay, ScreenGame, OSLock, and Screen2000. The columns show the wall input, ground truth, AGTRL, and the labeled comparison methods at a common scale.

The recovered information takes different forms across the four content families. Natural-image examples reveal broad color composition; interface and desktop examples recover major region positions and window layouts; game examples retain the dominant arrangement of the displayed scene. Across these contents, the weak wall pattern supports recovery of an organized screen image. The following ablation examines how correspondence learning and trace arbitration bring this structure into the reconstruction.

4.3 ABLATION STUDY

We test the contributions of correspondence learning and trace arbitration through the quantitative main-component ablation in Table 2 and the corresponding qualitative comparison in Figure 5.

Table 2: Main-component ablation across four fixed paired test sets. PSNR is reported in dB; higher PSNR and SSIM and lower LPIPS indicate better reconstruction. AGTRL is shown in bold red, and the strongest ablation values, including ties, are shown in bold blue. A pale red background highlights AGTRL, and up to three ablation rows with the largest numbers of strongest-alternative entries are shaded pale blue.

Method	OSLock			PhotoDisplay			Screen2000			ScreenGame		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
AGRL Variants												
w/o AGRL	16.636	0.612	0.647	13.836	0.422	0.833	19.418	**0.714**	0.462	17.596	0.655	0.603
AGRL → plain residual	16.616	0.610	0.642	13.754	**0.423**	0.827	19.272	0.703	0.466	**17.794**	0.654	0.604
AGRL → frequency branch	16.609	0.611	0.642	13.720	0.422	0.830	19.208	0.702	0.468	17.752	**0.656**	0.604
AGRL → attention branch	16.611	0.609	0.656	13.676	0.420	0.831	19.437	0.708	0.462	17.674	0.653	0.603
TraceTrust Variants												
w/o TraceTrust	16.225	0.609	0.649	13.869	0.422	0.834	19.325	0.703	0.464	17.672	0.646	**0.600**
TraceTrust → sigmoid gate	16.334	0.600	0.644	13.757	**0.423**	**0.824**	19.474	**0.714**	0.458	17.653	0.652	0.601
TraceTrust → channel gate	**16.651**	**0.613**	**0.638**	13.637	0.419	0.826	19.374	0.697	0.466	17.541	0.649	0.605
TraceTrust → convolutional spatial gate	16.598	0.606	0.655	**13.897**	**0.423**	0.830	**19.530**	0.705	**0.456**	17.554	0.651	0.601
AGTRL	**16.710**	**0.613**	**0.620**	**13.925**	**0.426**	**0.806**	**24.706**	**0.837**	**0.259**	**18.639**	**0.684**	**0.533**

Combining correspondence learning with trace arbitration yields the clearest gain in recovering screen organization. On Screen2000, the complete model exceeds the strongest ablation by 5.176 dB PSNR, accompanied by more coherent color regions and clearer major boundaries in Figure 5. Across Table 2, AGTRL leads 11 metric columns and shares the best displayed OSLock SSIM of 0.613. At this shared displayed SSIM, it also reduces LPIPS from 0.638 for the channel-gate variant to 0.620. These complementary gains in layout and appearance support the joint design: AGRL organizes mixed wall evidence into a screen-oriented update, and TraceTrust selects how strongly that update contributes to reconstruction.

4.4 ROBUSTNESS TO GAUSSIAN INPUT NOISE

We evaluate fixed models on OSLock as Gaussian input-noise SNR varies from 25 to 40 dB. The quantitative Gaussian-noise results for OSLock are reported in Table 3. Because the perturbation is applied directly to the wall observation while model parameters remain fixed, the sweep isolates sensitivity to distributed measurement noise rather than adaptation through retraining. Figure 4 summarizes the metric trends across the tested SNR range, while Table 3 provides the corresponding numerical values for all methods.

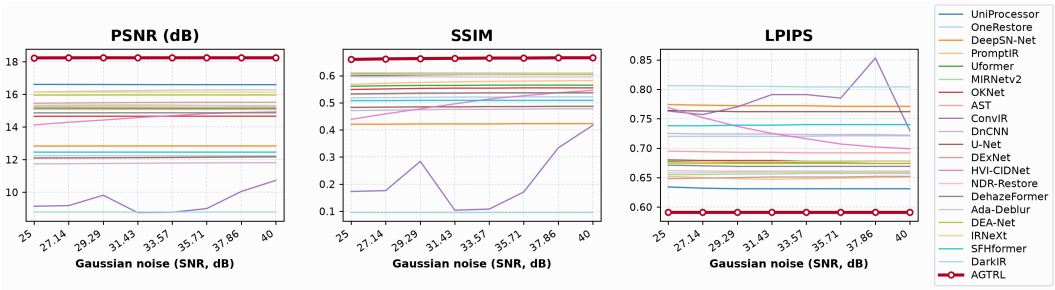

Figure 4: Gaussian-noise robustness on OSLock. The panels show PSNR, SSIM, and LPIPS as input SNR varies from 25 to 40 dB, with model weights fixed.

Table 3: Quantitative results for Gaussian noise on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	25			27.14			29.29			31.43			33.57			35.71			37.86			40		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**16.610**	0.599	**0.634**	**16.608**	0.601	**0.632**	**16.605**	0.603	**0.631**	**16.600**	0.604	**0.631**	**16.597**	0.605	**0.631**	**16.596**	0.605	**0.631**	**16.595**	0.605	**0.631**	**16.595**	0.605	**0.631**
OneRestore	12.243	0.518	0.720	12.242	0.519	0.720	12.241	0.520	0.720	12.240	0.521	0.720	12.239	0.521	0.720	12.238	0.522	0.721	12.238	0.522	0.721	12.237	0.522	0.721
DeepSN-Net	12.833	0.421	0.774	12.834	0.421	0.773	12.835	0.422	0.772	12.836	0.422	0.772	12.836	0.422	0.772	12.836	0.423	0.771	12.836	0.423	0.771	12.836	0.423	0.771
PromptIR	15.325	0.568	0.652	15.343	0.572	0.650	15.353	0.575	0.648	15.355	0.577	0.647	15.345	0.579	0.648	15.323	0.580	0.649	15.290	0.581	0.650	15.257	0.582	0.651
Uformer	15.149	0.561	0.677	15.151	0.563	0.676	15.150	0.564	0.675	15.150	0.564	0.675	15.149	0.565	0.675	15.149	0.565	0.675	15.149	0.565	0.674	15.149	0.565	0.674
MIRNetv2	15.235	0.604	0.655	15.227	0.604	0.655	15.221	0.604	0.656	15.216	0.604	0.656	15.213	0.604	0.657	15.212	0.604	0.657	15.211	0.604	0.657	15.211	0.604	0.657
OKNet	14.655	0.549	0.680	14.658	0.551	0.679	14.660	0.553	0.679	14.661	0.554	0.679	14.662	0.554	0.678	14.663	0.555	0.678	14.663	0.555	0.678	14.664	0.555	0.678
AST	15.095	0.531	0.695	15.090	0.533	0.694	15.085	0.534	0.693	15.082	0.535	0.693	15.079	0.536	0.692	15.077	0.536	0.692	15.075	0.536	0.692	15.074	0.536	0.692
ConvIR	9.143	0.173	0.763	9.182	0.176	0.757	9.817	0.284	0.770	8.753	0.104	0.791	8.778	0.108	0.791	8.999	0.171	0.785	10.049	0.334	0.853	10.728	0.417	0.730
DnCNN	11.739	0.471	0.725	11.750	0.473	0.724	11.764	0.474	0.724	11.778	0.476	0.724	11.790	0.476	0.723	11.800	0.477	0.723	11.808	0.478	0.723	11.813	0.478	0.722
U-Net	12.101	0.483	0.764	12.106	0.484	0.763	12.112	0.485	0.762	12.122	0.485	0.762	12.136	0.486	0.762	12.150	0.486	0.762	12.162	0.487	0.762	12.170	0.487	0.762
DExNet	15.454	**0.609**	0.648	15.472	**0.609**	0.649	15.487	**0.609**	0.650	15.498	**0.609**	0.651	15.505	**0.609**	0.651	15.509	**0.609**	0.651	15.512	**0.609**	0.652	15.515	**0.609**	0.652
HVI-CIDNet	14.129	0.439	0.769	14.292	0.459	0.752	14.425	0.478	0.737	14.577	0.496	0.725	14.699	0.514	0.716	14.819	0.526	0.707	14.889	0.537	0.702	14.928	0.546	0.699
NDR-Restore	15.306	0.596	0.658	15.310	0.596	0.658	15.313	0.596	0.659	15.314	0.595	0.659	15.317	0.595	0.659	15.318	0.595	0.659	15.319	0.595	0.659	15.319	0.595	0.659
DehazeFormer	14.865	0.533	0.671	14.868	0.534	0.670	14.869	0.536	0.669	14.870	0.536	0.669	14.870	0.537	0.669	14.871	0.537	0.669	14.872	0.537	0.669	14.872	0.537	0.669
Ada-Deblur	16.145	0.604	0.662	16.183	0.605	0.661	16.208	0.606	0.661	16.232	0.607	0.661	16.256	0.607	0.661	16.272	0.608	0.661	16.280	0.608	0.661	16.280	0.608	0.661
DEA-Net	15.959	0.607	0.674	15.958	0.607	0.674	15.963	0.607	0.674	15.963	0.607	0.674	15.962	0.607	0.674	15.964	0.607	0.674	15.962	0.607	0.674	15.962	0.607	0.674
IRNeXt	16.138	0.604	0.675	16.136	0.605	0.676	16.136	0.606	0.677	16.135	0.607	0.677	16.136	0.607	0.678	16.134	0.608	0.678	16.134	0.608	0.678	16.134	0.608	0.678
SFHformer	12.460	0.508	0.738	12.460	0.508	0.738	12.461	0.509	0.739	12.461	0.509	0.739	12.461	0.509	0.740	12.461	0.509	0.740	12.461	0.509	0.740	12.461	0.509	0.740
DarkIR	8.789	0.095	0.806	8.788	0.095	0.806	8.788	0.095	0.805	8.788	0.095	0.805	8.788	0.095	0.804	8.788	0.095	0.804	8.788	0.095	0.804	8.787	0.095	0.804
AGTRL	**18.235**	**0.660**	**0.591**	**18.240**	**0.662**	**0.591**	**18.243**	**0.663**	**0.591**	**18.240**	**0.664**	**0.591**	**18.242**	**0.665**	**0.591**	**18.242**	**0.665**	**0.591**	**18.241**	**0.666**	**0.591**	**18.241**	**0.666**	**0.591**

Gaussian noise perturbs the wall signal at many locations simultaneously. AGTRL retains the leading reconstruction scores throughout the sweep, with a PSNR span of 0.008 dB. The narrow range indicates stable recovery of interface content under the tested noise levels.

Screen recovery remains stable as random fluctuations are added to the wall observation. Across input SNR levels from 25 to 40 dB, AGTRL leads all three reconstruction metrics, and its PSNR varies by just 0.008 dB in Table 3 and Figure 4. Together with the component ablation, this stable response supports the joint use of spatially organized features and trust-weighted updates for screen-structure recovery from weak, noisy wall observations.

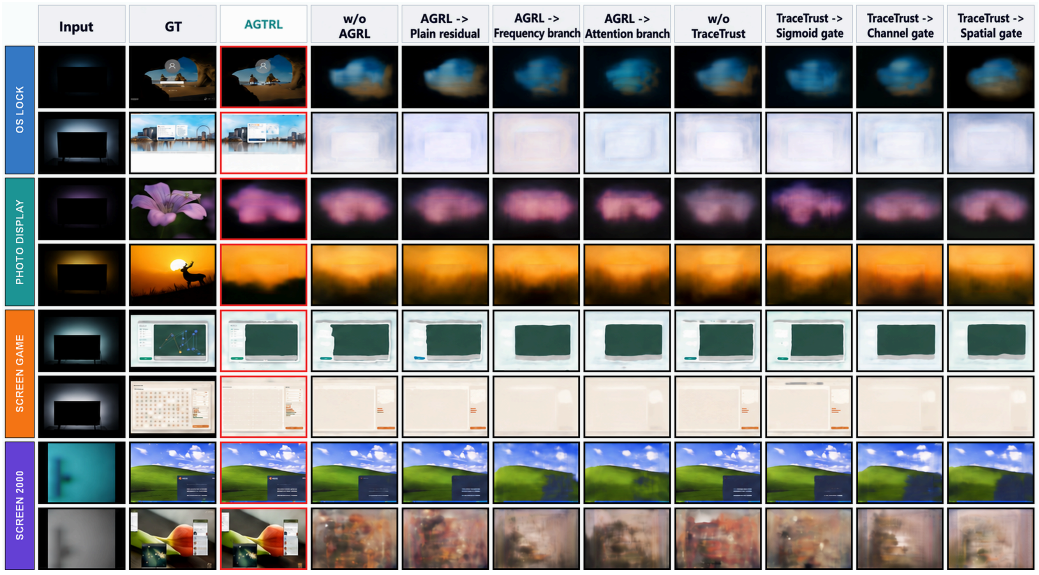

Figure 5: Qualitative main-component ablation of AGTRL on OSLock, PhotoDisplay, ScreenGame, and Screen2000, corresponding to Table 2. Two fixed examples are shown for each dataset. From left to right, each row contains the wall-reflection input, ground-truth screen, AGTRL, four AGRL variants, and four TraceTrust variants. The AGRL variants remove AGRL or replace it with a plain residual, frequency branch, or attention branch. The TraceTrust variants remove TraceTrust or replace it with sigmoid, channel, or spatial gating. All variants use the same wall input and ground truth within each example.

4.5 LOCAL MASKING ANALYSIS

We conduct a local masking diagnostic on Screen2000 to examine how different wall regions contribute to hidden-screen recovery in a T-shadowed wall observation. The observation is divided into a 3×3 grid. Using fixed learned weights, AGTRL reconstructs the screen from the unmasked input and nine single-region masks.

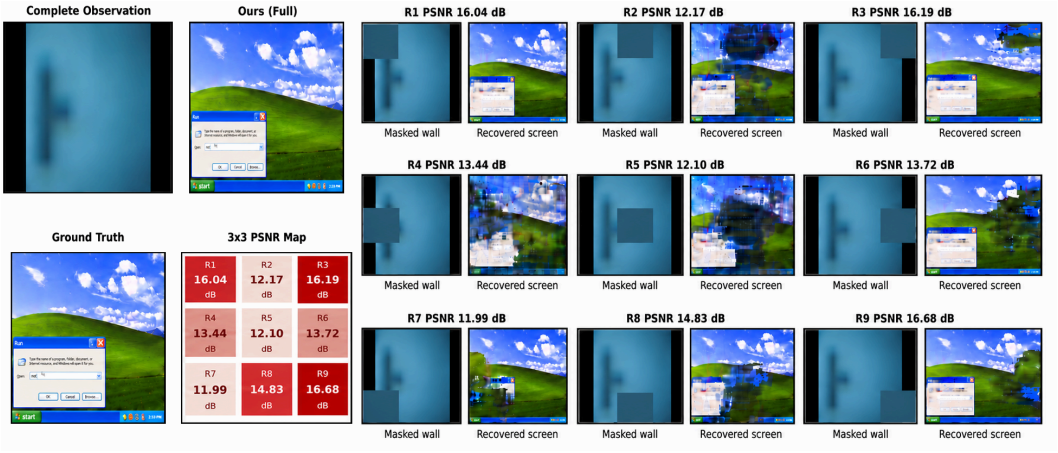

Figure 6: Local 3×3 masking diagnostic on Screen2000 with fixed AGTRL weights. The complete wall observation, ground-truth screen, unmasked recovery, and nine masked wall inputs with their recovered screens are shown. Each heatmap cell gives the reconstruction PSNR after masking the corresponding wall region.

Indirect image formation distributes scene contributions across the observed surface (Saunders et al., 2019). Figure 6 links a local mask on the wall to changes in several separated regions of the recovered screen. Masking R2 or R5 alters both the scene appearance and the window region, yielding 12.17 and 12.10 dB PSNR; masking R9 preserves a more coherent central layout and yields 16.68 dB. The changes across the sky, hillside, window, and taskbar reveal the spatially distributed evidence used by AGTRL. AGRL aggregates neighboring wall-derived features to assemble this evidence, and TraceTrust weights its local support before reconstruction.

4.6 RECOVERY ACROSS OBSERVATION CONDITIONS

The supplementary experiments extend the recovery evidence along three connected directions: weaker wall signals, altered spatial arrangements, and reduced visible evidence. On ScreenGame, AGTRL retains 0.664 SSIM at the -10% wall-image brightness setting; on OSLock, it achieves 17.347 dB PSNR at the labeled $+15^\circ$ pitch setting. The shared U-Net representation-module comparisons further examine recovery with masking ratios up to 75% and across eight wall-appearance proxies. These evaluations keep the learned reconstruction model fixed within each test and alter only the observation condition, so the resulting differences reflect changes in the available wall evidence. The combined photometric, geometric, masking, and wall-appearance tests also examine whether the learned correspondence remains useful beyond a single acquisition configuration. Together with the noise, contrast, distance, and other geometric tests, these results show how screen structure remains recoverable as the strength, arrangement, and availability of wall evidence change.

5 CONCLUSION

Our results show that an ordinary camera can recover hidden-screen content from the light it casts on a wall. AGTRL makes this possible through two connected operations: AGRL learns the correspondence between mixed wall features and screen structure, and TraceTrust selects the traces that guide reconstruction. Across four screen-content families, AGTRL outperforms 20 reconstruction methods in all 12 dataset–metric combinations. Additional experiments extend the recovery analysis across changes in distance, brightness, viewpoint, noise, occlusion, and wall appearance. Optical emanation studies recovered display information from temporally resolved reflected light, and computational periscopy established ordinary-camera recovery of hidden scenes (Kuhn, 2002; Saunders et al., 2019). AGTRL develops this connection through joint correspondence learning and trace arbitration, recovering principal color distributions, region positions, page layouts, and structural contours from RGB wall observations. Together, these findings establish the initial feasibility of recovering hidden-screen content with an ordinary camera observing only wall light.

REFERENCES

Miika Aittala, Prafull Sharma, Lukas Murmann, Adam B. Yedidia, Gregory W. Wornell, William T. Freeman, and Frédo Durand. Computational mirrors: Blind inverse light transport by deep matrix factorization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/5a2afca61e35f45a7dd44ca46e0225f4-Paper.pdf.

Katherine L. Bouman, Vickie Ye, Adam B. Yedidia, Fredo Durand, Gregory W. Wornell, Antonio Torralba, and William T. Freeman. Turning corners into cameras: Principles and methods. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2289–2297. IEEE, 2017. doi: 10.1109/ICCV.2017.249. URL https://doi.org/10.1109/ICCV.2017.249.

Zixuan Chen, Zewei He, and Zhe-Ming Lu. DEA-Net: Single image dehazing based on detail-enhanced convolution and content-guided attention. *IEEE Transactions on Image Processing*, 33: 1002–1015, 2024. doi: 10.1109/TIP.2024.3354108.

Lu Chi, Borui Jiang, and Yadong Mu. Fast Fourier convolution. In *Advances in Neural Information Processing Systems*, volume 33, pp. 4479–4488, 2020.

Hyungjin Chung, Jeongsol Kim, Michael T. McCann, Marc L. Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/forum?id=OnD9zGAGT0k.

Yuning Cui, Wenqi Ren, Sining Yang, Xiaochun Cao, and Alois Knoll. IRNeXt: Rethinking convolutional network design for image restoration. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, volume 202 of *Proceedings of Machine Learning Research*, pp. 6545–6564, 2023. URL https://proceedings.mlr.press/v202/cui23d.html.

Yuning Cui, Wenqi Ren, Xiaochun Cao, and Alois Knoll. Revitalizing convolutional network for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 9423–9438, 2024a. doi: 10.1109/TPAMI.2024.3419007.

Yuning Cui, Wenqi Ren, and Alois Knoll. Omni-kernel network for image restoration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1426–1434, 2024b. doi: 10.1609/aaai.v38i2.27907.

Robinson Czajkowski and John Murray-Bruce. Two-edge-resolved three-dimensional non-line-of-sight imaging with an ordinary camera. *Nature Communications*, 15(1):1162, 2024. doi: 10.1038/s41467-024-45397-7.

Xin Deng, Chenxiao Zhang, Lai Jiang, Jingyuan Xia, and Mai Xu. DeepSN-Net: Deep semi-smooth Newton driven network for blind image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(4):2632–2646, 2025. doi: 10.1109/TPAMI.2024.3525089.

Huiyu Duan, Xiongkuo Min, Sijing Wu, Wei Shen, and Guangtao Zhai. UniProcessor: A text-induced unified low-level image processor. In *Computer Vision – ECCV 2024*, volume 15125 of *Lecture Notes in Computer Science*, pp. 180–199. Springer, Cham, 2025. doi: 10.1007/978-3-031-72855-6_11. URL https://doi.org/10.1007/978-3-031-72855-6_11.

Daniel Feijoo, Juan C. Benito, Alvaro Garcia, and Marcos V. Conde. DarkIR: Robust low-light image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10879–10889, 2025. doi: 10.1109/CVPR52734.2025.01016.

Ruixu Geng, Yang Hu, Zhi Lu, Cong Yu, Houqiang Li, Hengyu Zhang, and Yan Chen. Passive non-line-of-sight imaging using optimal transport. *IEEE Transactions on Image Processing*, 31: 110–124, 2022. doi: 10.1109/TIP.2021.3128312.

Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. OneRestore: A universal restoration framework for composite degradation. In *Computer Vision – ECCV 2024*, volume 15077 of *Lecture Notes in Computer Science*, pp. 255–272. Springer, Cham, 2025. doi: 10.1007/978-3-031-72655-2_15. URL https://doi.org/10.1007/978-3-031-72655-2_15.

Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7132–7141, 2018.

Jun-Jie Huang, Tianrui Liu, Zihan Chen, Xinwang Liu, Meng Wang, and Pier Luigi Dragotti. A lightweight deep exclusion unfolding network for single image reflection removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(6):4957–4973, 2025. doi: 10.1109/TPAMI.2025.3548148.

Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of PSNR in image/video quality assessment. *Electronics Letters*, 44(13):800–801, 2008. doi: 10.1049/el:20080522.

Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 13899–13909. IEEE, 2021. doi: 10.1109/ICCV48922.2021.01366. URL https://doi.org/10.1109/ICCV48922.2021.01366.

Xingyu Jiang, Xiuhui Zhang, Ning Gao, and Yue Deng. When fast Fourier transform meets transformer for image restoration. In *Computer Vision – ECCV 2024*, volume 15103 of *Lecture Notes in Computer Science*, pp. 381–402. Springer, Cham, 2025. doi: 10.1007/978-3-031-72995-9_22. URL https://doi.org/10.1007/978-3-031-72995-9_22.

Markus G. Kuhn. Optical time-domain eavesdropping risks of CRT displays. In *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 3–18, 2002. doi: 10.1109/SECPRI.2002.1004358.

Xintong Liu, Jianyu Wang, Leping Xiao, Zuoqiang Shi, Xing Fu, and Lingyun Qiu. Non-line-of-sight imaging with arbitrary illumination and detection pattern. *Nature Communications*, 14:3230, 2023. doi: 10.1038/s41467-023-38898-4.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. PromptIR: Prompting for all-in-one image restoration. In *Advances in Neural Information Processing Systems*, volume 36, pp. 71275–71293. Curran Associates, Inc., 2023. doi: 10.52202/075280-3121. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/e187897ed7780a579a0d76fd4a35d107-Abstract-Conference.html.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 234–241, 2015. doi: 10.1007/978-3-319-24574-4_28.

Charles Saunders, John Murray-Bruce, and Vivek K. Goyal. Computational periscopy with an ordinary digital camera. *Nature*, 565(7740):472–475, 2019. doi: 10.1038/s41586-018-0868-6.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Yuda Song, Zhuqing He, Hui Qian, and Xin Du. Vision transformers for single image dehazing. *IEEE Transactions on Image Processing*, 32:1927–1941, 2023. doi: 10.1109/TIP.2023.3256763.

Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Moungi G. Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3(1):745, 2012. doi: 10.1038/ncomms1747.

Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7794–7803, 2018.

Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general U-shaped transformer for image restoration. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17662–17672. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01716. URL https://doi.org/10.1109/CVPR52688.2022.01716.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. doi: 10.1109/TIP.2003.819861.

Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. CBAM: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision*, pp. 3–19, 2018.

Qiucheng Wu, Yifan Jiang, Junru Wu, Victor Kulikov, Vidit Goel, Nikita Orlov, Humphrey Shi, Zhangyang Wang, and Shiyu Chang. Broad spectrum image deblurring via an adaptive supernetwork. *IEEE Transactions on Image Processing*, 32:5270–5282, 2023. doi: 10.1109/TIP.2023.3312912.

Qingsen Yan, Yixu Feng, Cheng Zhang, Guansong Pang, Kangbiao Shi, Peng Wu, Wei Dong, Jinqiu Sun, and Yanning Zhang. HVI: A new color space for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5678–5687, 2025. doi: 10.1109/CVPR52734.2025.00533.

Mingde Yao, Ruikang Xu, Yuanshen Guan, Jie Huang, and Zhiwei Xiong. Neural degradation representation learning for all-in-one image restoration. *IEEE Transactions on Image Processing*, 33:5408–5423, 2024. doi: 10.1109/TIP.2024.3456583.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5718–5729. IEEE, 2022. doi: 10.1109/CVPR52688.2022.00564. URL https://doi.org/10.1109/CVPR52688.2022.00564.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for fast image restoration and enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1934–1948, 2023. doi: 10.1109/TPAMI.2022.3167175.

Jiarui Zhang, Ruixu Geng, Xiaolong Du, Yan Chen, Houqiang Li, and Yang Hu. Passive non-line-of-sight imaging with light transport modulation. *IEEE Transactions on Image Processing*, 34: 410–424, 2025. doi: 10.1109/TIP.2024.3518097.

Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26 (7):3142–3155, 2017. doi: 10.1109/TIP.2017.2662206.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018. doi: 10.1109/CVPR.2018.00068.

Zhiwen Zheng, Yuheng Qiao, Xiaoshuai Zhang, Zhao Huang, Tao Zhang, Huiyu Zhou, Shaowei Jiang, Jin Liu, Wenwen Tang, and Xingru Huang. Physically-guided optical inversion enable non-contact side-channel attack on isolated screens, 2026. URL https://arxiv.org/abs/2604.13419. arXiv preprint arXiv:2604.13419.

Shihao Zhou, Duosheng Chen, Jinshan Pan, Jinglei Shi, and Jufeng Yang. Adapt or perish: Adaptive sparse transformer with attentive feature refinement for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2952–2963, 2024. doi: 10.1109/CVPR52733.2024.00285.

SUPPLEMENTARY MATERIAL FOR AGTRL: AXIS-GUIDED TRUSTED REPRESENTATION LEARNING FOR PASSIVE HIDDEN-SCREEN RECOVERY

This supplement expands the evidence chain of the main paper. It follows the observed wall-reflected screen signal from passive acquisition through axis-guided correspondence learning and TraceTrust arbitration to the recovered screen. The results show how AGTRL recovers principal color distributions, regional layouts, page structures, and structural contours from weak wall observations across the reported datasets and observation conditions.

A ADDITIONAL SCENARIO INFORMATION

A.1 CAPTURE SETTINGS AND INDIRECT-LIGHT OBSERVATION

We use paired RGB wall observations and screen targets to study passive hidden-screen recovery under two acquisition configurations. In both configurations, the camera records light reflected from a white diffuse relay wall, while the screen remains outside its direct view. Acquisition is passive and uses the display's own emission. The paired screen image provides training supervision, and the RGB wall observation provides the inference input. The two configurations differ in the object that blocks direct observation: an external baffle in Scene 1 and the display body itself in Scene 2.

A.1.1 SCENE 1: BAFFLE OCCLUSION

Scene 1 represents an office environment in which an external baffle obstructs the camera's view of a self-emissive LED display. The display and its user are on one side of the baffle, and the acquisition camera is on the other. The user can view the screen directly, whereas the camera observes only the nearby white diffuse relay wall. The observation geometry is illustrated in Figure 1, and photographs of the physical setup are provided in Supplementary Figure 1.

Light emitted by the screen first illuminates the relay wall, where diffuse reflection forms a spatially smooth, low-frequency pattern. The baffle further occludes and spatially modulates the reflected light reaching the camera, introducing shadow structure into the captured wall observation. The recorded image is also affected by diffraction, diffuse scattering, sensor noise, and photometric drift. Each sample pairs this wall observation with the corresponding ground-truth screen image. We recover the principal color distributions, region positions, page layouts, and structural contours of hidden-screen content, including webpages, desktop layouts, charts, and login interfaces, from the corresponding wall observations. Screen2000 uses this baffle-modulated configuration.

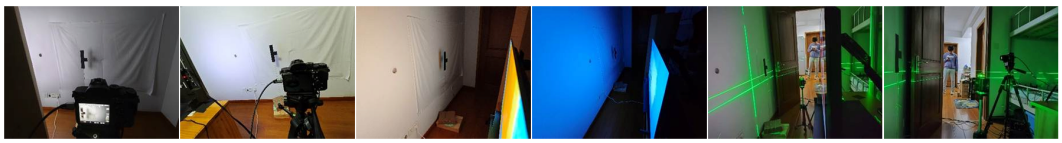

Supplementary Figure 1: Real photographs of Scene 1 with baffle occlusion. From left to right, the first two panels show the camera facing the relay wall, the next two show the display beside the illuminated wall, and the final two show wider views of the corridor arrangement. These photographs document the setup from different viewpoints; the acquisition camera observes wall-reflected screen light while the baffle blocks its direct view of the display. The green reference lines indicate setup alignment.

A.1.2 SCENE 2: DISPLAY SELF-OCCLUSION

Scene 2 represents an office arrangement in which a self-emissive LED display faces a nearby white diffuse relay wall and the user faces the screen. The acquisition camera is placed behind the display and points toward the same wall. The opaque display body blocks the camera's direct view of the emitting screen surface, so the camera records the wall illumination visible around the display.

Supplementary Figure 2 illustrates this observation geometry, and Supplementary Figure 3 shows photographs of the physical setup.

Screen light undergoes diffuse reflection at the relay wall before reaching the camera, spreading the displayed content into low-frequency brightness and color variations. The observation is further affected by the diffuse light-transport process, sensor noise, and photometric drift. Each sample consists of a wall observation paired with its ground-truth screen image. We recover screen structure from wall observations for content including webpages, desktop layouts, charts, user interfaces, games, and photographs. OSLock, PhotoDisplay, ScreenGame, and ScreenFigure use this display-self-occluded configuration.

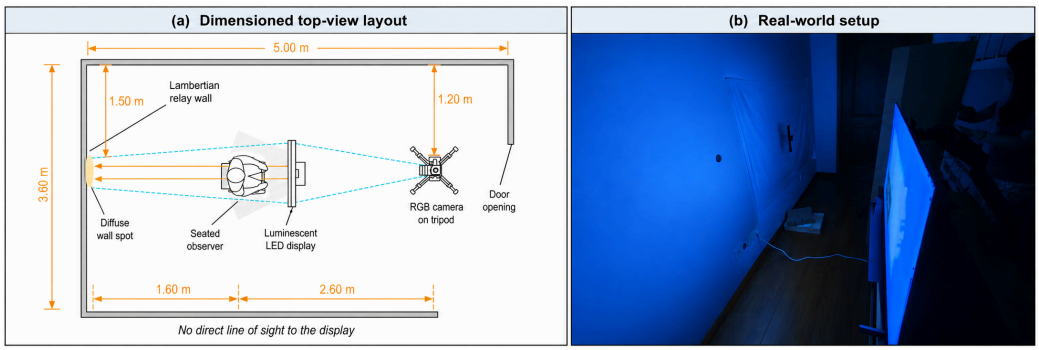

Supplementary Figure 2: Observation geometry of Scene 2 with display self-occlusion. The display faces the diffuse relay wall, and the camera observes that wall from behind the display. The opaque display body blocks direct observation of the screen surface, while reflected screen light remains visible on the surrounding wall and provides the input for hidden-screen recovery.



Supplementary Figure 3: Real photographs of Scene 2 with display self-occlusion. From left to right, the first two panels show the display, camera, and operator during setup; the next two show the camera–display arrangement and the illuminated wall behind the display; and the final two show wider views with different wall-light patterns. The display body obscures the screen from the acquisition camera, which records the indirect illumination on the wall.

The main experiments use OSLock, PhotoDisplay, Screen2000, and ScreenGame, with fixed training and test splits and a separate model trained for each dataset. ScreenFigure is an additional dataset used in the supplementary comparisons described below.

A.1.3 OSLock

OSLock contains lock-screen backgrounds with sparse interface elements, acquired using the display-self-occluded configuration. Its screen targets combine broad background color regions with localized interface structures. This content family supports evaluation of background appearance, region position, and coarse interface layout recovered from the corresponding wall observations. The dataset contains 1,120 training pairs and 279 test pairs.

Example pairs are shown in Supplementary Figure 4.

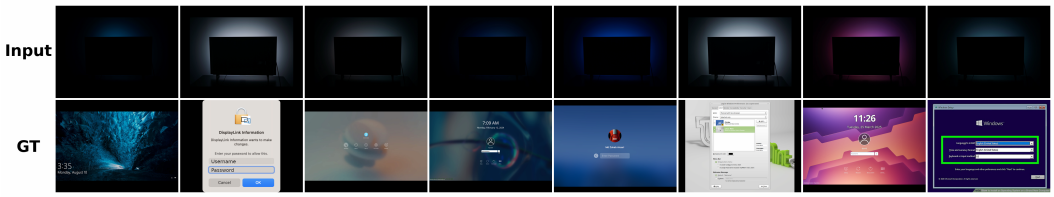

Supplementary Figure 4: Eight test pairs from OSLock, acquired using display self-occlusion. Top row (Input): RGB wall observations. Bottom row (GT): corresponding screen targets. Each column contains one matched pair.

A.1.4 PHOTODISPLAY

PhotoDisplay contains natural photographs presented on the display and acquired using the display-self-occluded configuration. The screen targets provide varied image compositions, color distributions, and major object contours. The paired wall observations support evaluation of these visual structures after diffuse reflection. The dataset contains 2,400 training pairs and 600 test pairs.
Example pairs are shown in Supplementary Figure 5.



Supplementary Figure 5: Eight test pairs from PhotoDisplay, acquired using display self-occlusion. Top row (Input): RGB wall observations. Bottom row (GT): corresponding screen targets. Each column contains one matched pair.

A.1.5 SCREEN2000

Screen2000 contains desktop and window layouts acquired using the baffle-modulated configuration. The wall observations exhibit uneven illumination and baffle-induced shadow structure, while the screen targets contain spatially arranged windows and interface regions. This dataset supports evaluation of layout recovery from spatially modulated wall light and is also used for the local-masking analysis. The dataset contains 2,121 training pairs and 251 test pairs.
Example pairs are shown in Supplementary Figure 6.

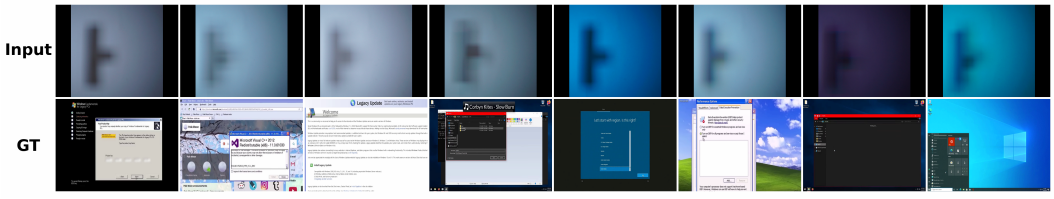

Supplementary Figure 6: Eight test pairs from Screen2000, acquired using baffle modulation. Top row (Input): RGB wall observations. Bottom row (GT): corresponding screen targets. Each column contains one matched pair.

A.1.6 SCREENGAME

ScreenGame contains game content acquired using the display-self-occluded configuration, including scenes with dark image regions and weak contrast in the corresponding wall observations. This content family supports evaluation of broad color regions and major scene contours when the available wall traces are weak. The dataset contains 1,000 training pairs and 250 test pairs.
Example pairs are shown in Supplementary Figure 7.

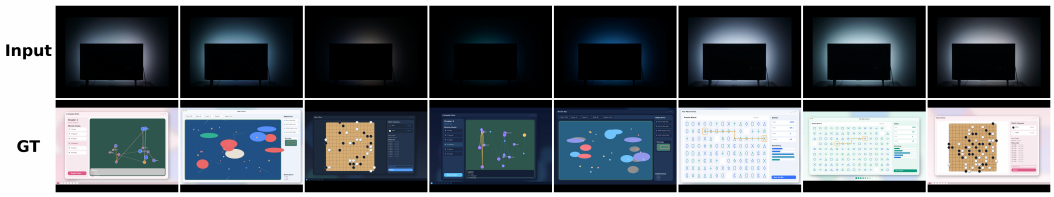

Supplementary Figure 7: Eight test pairs from ScreenGame, acquired using display self-occlusion. Top row (Input): RGB wall observations. Bottom row (GT): corresponding screen targets. Each column contains one matched pair.

A.1.7 SCREENFIGURE

ScreenFigure contains displayed charts, plots, and diagram-like content acquired using the display-self-occluded configuration. These screen targets emphasize graphical elements and their spatial arrangement. ScreenFigure provides an additional content family for the representation-module and training-objective comparisons in Supplementary Tables 24 and 25.
Example pairs are shown in Supplementary Figure 8.



Supplementary Figure 8: Eight test pairs from ScreenFigure, acquired using display self-occlusion. Top row (Input): RGB wall observations. Bottom row (GT): corresponding screen targets. Each column contains one matched pair.

A.2 AGRL AND TRACETRUST FORMULATION

Let the encoder bottleneck be $z \in \mathbb{R}^{B\times C\times h\times w}$, with $C = 384$, and let $K = 96$ be the AGRL hidden width. The indices $\mathfrak{b}$, χ, k, and (i, j) denote batch, bottleneck channel, hidden channel, and spatial position, respectively. The operator $\oplus_a$ denotes concatenation along the channel axis. The following pointwise projection learns how to combine the concatenated features. The sets $\mathcal{N}_5 = \{-2,-1,0,1,2\}^2$ and $\mathcal{N}_3 = \{-1,0,1\}^2$ specify the shape-preserving 5×5 and 3×3 neighborhoods.

A.2.1 CHANNEL-AXIS STANDARDIZATION

At each bottleneck location, AGRL removes the shared channel offset, normalizes the channel dispersion, and restores a learned channel-specific affine calibration. Expanding the vector operation over all channel coordinates gives

$$\mu_{\mathfrak{b}ij} := \frac{1}{C}\sum_{\chi'=1}^{C} z_{\mathfrak{b}\chi' ij},$$

$$\nu_{\mathfrak{b}ij} := \frac{1}{C}\sum_{\chi'=1}^{C}\left(z_{\mathfrak{b}\chi' ij} - \frac{1}{C}\sum_{\chi''=1}^{C} z_{\mathfrak{b}\chi'' ij}\right)^2,$$

$$\widetilde{z}_{\mathfrak{b}\chi ij} := a_\chi \frac{z_{\mathfrak{b}\chi ij} - \mu_{\mathfrak{b}ij}}{\sqrt{\nu_{\mathfrak{b}ij} + \epsilon}} + b_{0,\chi} = a_\chi \frac{z_{\mathfrak{b}\chi ij} - \frac{1}{C}\sum_{\chi'=1}^{C} z_{\mathfrak{b}\chi' ij}}{\sqrt{\frac{1}{C}\sum_{\chi'=1}^{C}\left(z_{\mathfrak{b}\chi' ij} - \frac{1}{C}\sum_{\chi''=1}^{C} z_{\mathfrak{b}\chi'' ij}\right)^2 + \epsilon}} + b_{0,\chi},$$
$$\mu, \nu \in \mathbb{R}^{B\times 1\times h\times w}, \qquad \widetilde{z} \in \mathbb{R}^{B\times C\times h\times w}, \qquad \epsilon = 10^{-6}. \tag{A1}$$

At each location, normalization centers and rescales the channel responses. The learned scale and offset then adapt these responses to screen reconstruction. The spatial aggregation stage operates on the resulting features $\widetilde{z}$.

A.2.2 Axis-guided residual proposal

The normalized tensor is first mixed along the channel axis, then aggregated independently within each hidden channel over $\mathcal{N}_5$, and finally projected back to C channels. We use the exact Gaussian error linear unit, $\mathrm{GELU}(a) = \frac{a}{2}[1 + \mathrm{erf}(a/\sqrt{2})]$. The stages of the residual proposal are

$$\begin{aligned}
s_{\mathfrak{b}kij} &:= \sum_{\chi'=1}^{C} W^{(1)}_{k\chi'} \widetilde{z}_{\mathfrak{b}\chi' ij} + b^{(1)}_k, \\
\tau_{\mathfrak{b}kij} &:= \sum_{(\delta_i,\delta_j)\in\mathcal{N}_5} D^{(5)}_{k\delta_i\delta_j} s_{\mathfrak{b}k,\, i+\delta_i,\, j+\delta_j} + \gamma^{(5)}_k, \\
v_{\mathfrak{b}kij} &:= \mathrm{GELU}(\tau_{\mathfrak{b}kij}), \\
r_{\mathfrak{b}\chi ij} &:= \sum_{k=1}^{K} W^{(2)}_{\chi k} v_{\mathfrak{b}kij} + b^{(2)}_\chi, \\
s, v \in \mathbb{R}^{B\times K\times h\times w}, &\qquad r \in \mathbb{R}^{B\times C\times h\times w}.
\end{aligned} \tag{A2}$$

Here, $W^{(1)}$ and $W^{(2)}$ define the two 1×1 projections. The depthwise kernel $D^{(5)}$ aggregates the 5×5 neighborhood independently within each hidden channel k. This sequence combines channel responses and neighboring evidence into the candidate residual r. The resulting r is the sole candidate residual passed to TraceTrust, and its bottleneck integration occurs after trust estimation.

A.2.3 Deterministic continuity and artifact evidence

TraceTrust receives $\widetilde{z}$ and computes two complementary statistics. The 3×3 deviation measures local continuity, whereas the 5×5 deviation measures artifact suspicion relative to the current image. Their complete indexed derivation is

$$\begin{aligned}
m_{\mathfrak{b}\chi ij} &:= \frac{1}{9} \sum_{(\delta_i,\delta_j)\in\mathcal{N}_3} \widetilde{z}_{\mathfrak{b}\chi,\, i+\delta_i,\, j+\delta_j}, \\
c_{\mathfrak{b}ij} &:= 1 - \mathrm{clip}_{[0,1]}\left[\frac{1}{C}\sum_{\chi=1}^{C} |\widetilde{z}_{\mathfrak{b}\chi ij} - m_{\mathfrak{b}\chi ij}|\right] \\
&= 1 - \mathrm{clip}_{[0,1]}\left[\frac{1}{C}\sum_{\chi=1}^{C}\left|\widetilde{z}_{\mathfrak{b}\chi ij} - \frac{1}{9}\sum_{(\delta_i,\delta_j)\in\mathcal{N}_3} \widetilde{z}_{\mathfrak{b}\chi,\, i+\delta_i,\, j+\delta_j}\right|\right], \\
q_{\mathfrak{b}ij} &:= \frac{1}{C}\sum_{\chi=1}^{C}\left|\widetilde{z}_{\mathfrak{b}\chi ij} - \frac{1}{25}\sum_{(\delta_i,\delta_j)\in\mathcal{N}_5} \widetilde{z}_{\mathfrak{b}\chi,\, i+\delta_i,\, j+\delta_j}\right|, \\
\overline{q}_{\mathfrak{b}} &:= \frac{1}{hw}\sum_{i'=1}^{h}\sum_{j'=1}^{w} q_{\mathfrak{b}i'j'}, \\
e_{\mathfrak{b}ij} &:= \mathrm{clip}_{[0,1]}\left[\frac{q_{\mathfrak{b}ij}}{\overline{q}_{\mathfrak{b}} + \epsilon}\right] = \mathrm{clip}_{[0,1]}\left[\frac{\frac{1}{C}\sum_{\chi=1}^{C}\left|\widetilde{z}_{\mathfrak{b}\chi ij} - \frac{1}{25}\sum_{(\delta_i,\delta_j)\in\mathcal{N}_5}\widetilde{z}_{\mathfrak{b}\chi,\, i+\delta_i,\, j+\delta_j}\right|}{\frac{1}{hw}\sum_{i'=1}^{h}\sum_{j'=1}^{w} q_{\mathfrak{b}i'j'} + \epsilon}\right], \\
c, e \in [0,1]^{B\times 1\times h\times w}, &\qquad q \in \mathbb{R}^{B\times 1\times h\times w}_{\geq 0}.
\end{aligned} \tag{A3}$$

The continuity value increases when a location agrees with its 3×3 neighborhood. The suspicion value first measures the high-frequency deviation from the wider 5×5 local average and then normalizes it by $\overline{q}_{\mathfrak{b}} + \epsilon$, so its scale follows the current wall-derived representation. The learned gate receives $\widetilde{z}$, c, and $1 - e$. The intermediate quantities m, q, and $\overline{q}$ provide the local averages, local deviations, and image-level reference used to construct these inputs.

A.2.4 LEARNED TRACETRUST ARBITRATION

Let $f_{\mathfrak{b}\lambda ij}$ be the $(C+2)$-channel fusion produced by $\widetilde{z} \oplus_a c \oplus_a (1-e)$. The pointwise projection, local trust aggregation, learned gate, and evidence modulation are

$$\begin{aligned}
f_{\mathfrak{b}:ij} &:= \widetilde{z}_{\mathfrak{b}:ij} \oplus_a c_{\mathfrak{b}ij} \oplus_a (1 - e_{\mathfrak{b}ij}) \in \mathbb{R}^{C+2}, \\
\varpi_{\mathfrak{b}kij} &:= \sum_{\lambda=1}^{C+2} W^{(3)}_{k\lambda} f_{\mathfrak{b}\lambda ij} + b^{(3)}_k, \\
\vartheta_{\mathfrak{b}kij} &:= \sum_{(\delta_i,\delta_j)\in\mathcal{N}_3} D^{(3)}_{k\delta_i\delta_j} \varpi_{\mathfrak{b}k,\, i+\delta_i,\, j+\delta_j} + \gamma^{(3)}_k, \\
g_{\mathfrak{b}ij} &:= \sigma\left(\sum_{k=1}^{K} W^{(4)}_k \,\mathrm{GELU}(\vartheta_{\mathfrak{b}kij}) + b^{(4)} \right), \\
t_{\mathfrak{b}ij} &:= g_{\mathfrak{b}ij}\, c_{\mathfrak{b}ij} \big(1 - e_{\mathfrak{b}ij}\big) \in [0,1].
\end{aligned} \tag{A4}$$

The operator $\oplus_a$ concatenates the inputs along the channel axis, and the pointwise projection over λ learns their combination. The 3×3 depthwise coefficients $D^{(3)}$ then aggregate local trust evidence, while g adapts the deterministic factors c and $1-e$ to the current representation. The output is the one-channel trust map t, which is broadcast across the C channels of r.

A.2.5 BOUNDED TRUSTED UPDATE

The learned scalar β is initialized to 10^{-4}. We set $\rho = 1.0$. The coefficient, trusted contribution, bottleneck update, and magnitude bound are

$$\begin{aligned}
\alpha &:= \rho \tanh(\beta), && |\alpha| \le |\rho|, \\
u_{\mathfrak{b}\chi ij} &:= \alpha\, r_{\mathfrak{b}\chi ij}\, t_{\mathfrak{b}ij}, \\
z^{+}_{\mathfrak{b}\chi ij} &:= z_{\mathfrak{b}\chi ij} + u_{\mathfrak{b}\chi ij}, \\
\big| z^{+}_{\mathfrak{b}\chi ij} - z_{\mathfrak{b}\chi ij} \big| &= |u_{\mathfrak{b}\chi ij}| \le |\rho|\, |r_{\mathfrak{b}\chi ij}|, \qquad u, z^{+} \in \mathbb{R}^{B\times C\times h\times w}.
\end{aligned} \tag{A5}$$

The inequality follows from $|\tanh(\beta)| \le 1$ and $0 \le t_{\mathfrak{b}ij} \le 1$. It controls the trusted contribution relative to the candidate update r: t weights the proposal, u is the resulting contribution, and z^{+} passes to the middle transformation.

A.2.6 MIDDLE TRANSFORMATION, SKIP DECODING, AND RGB READOUT

The middle transformation and skip-connected decoder are defined below. Let $\mathcal{N}_{\mathrm{m}}$ and $\mathcal{N}_\ell$ denote their neighborhoods, $M^{(5)}$ and $F^{(\ell)}$ their learned coefficients, and $\Lambda^{(\ell)}_{ijpq}$ the interpolation weights that map decoder level $\ell+1$ to level ℓ. Encoder skips ξ_ℓ enter their matched scales.

$$\begin{aligned}
\omega^{(5)}_{\mathfrak{b}\chi ij} &:= \varphi_5\left(\sum_{\substack{\chi'=1,\dots,C \\ (\delta_i,\delta_j)\in\mathcal{N}_{\mathrm{m}}}} M^{(5)}_{\chi\chi'\delta_i\delta_j} z^{+}_{\mathfrak{b}\chi',\, i+\delta_i,\, j+\delta_j} + \zeta^{(5)}_{\chi} \right), \\
\upsilon^{(\ell)}_{\mathfrak{b}\chi ij} &:= \sum_{i'=1}^{h_{\ell+1}} \sum_{j'=1}^{w_{\ell+1}} \Lambda^{(\ell)}_{iji'j'}\, \omega^{(\ell+1)}_{\mathfrak{b}\chi i'j'}, \\
\varrho^{(\ell)}_{\mathfrak{b}:ij} &:= \upsilon^{(\ell)}_{\mathfrak{b}:ij} \oplus_a \xi_{\ell,\mathfrak{b}:ij}, \\
\omega^{(\ell)}_{\mathfrak{b}\chi ij} &:= \varphi_\ell\left(\sum_{\substack{\lambda=1,\dots,\widetilde{C}_\ell \\ (\delta_i,\delta_j)\in\mathcal{N}_\ell}} F^{(\ell)}_{\chi\lambda\delta_i\delta_j} \varrho^{(\ell)}_{\mathfrak{b}\lambda,\, i+\delta_i,\, j+\delta_j} + \zeta^{(\ell)}_{\chi} \right), \qquad \ell = 4,3,2,1, \\
d_{\mathfrak{b}\kappa ij} &:= \sum_{\substack{\chi=1,\dots,C_1 \\ (\delta_i,\delta_j)\in\mathcal{N}_{\mathrm{rgb}}}} P_{\kappa\chi\delta_i\delta_j} \omega^{(1)}_{\mathfrak{b}\chi,\, i+\delta_i,\, j+\delta_j} + b^{(\mathrm{rgb})}_{\kappa}, \\
\widehat{y}_{\mathfrak{b}\kappa ij} &:= x_{\mathfrak{b}\kappa ij} + d_{\mathfrak{b}\kappa ij}, \qquad d, \widehat{y} \in \mathbb{R}^{B\times 3\times H\times W}, \qquad \kappa = 1,2,3.
\end{aligned} \tag{A6}$$

Here, $\widetilde{C}_\ell$ is the channel count after concatenation with ξ_ℓ, and φ_5 and φ_ℓ denote the nonlinear transformations in the middle block and decoder, respectively. The sequence is explicit: $z^{+} \to$

$\omega^{(5)} \to \omega^{(4)} \to \omega^{(3)} \to \omega^{(2)} \to \omega^{(1)} \to d \to \widehat{y}$. The addition $z^+ = z + u$ in Equation A5 and the RGB addition $\widehat{y} = x + d$ above are two distinct residual operations at different representations.

A.2.7 Paired reconstruction objective

Let $y \in \mathbb{R}^{B\times 3\times H\times W}$ be the paired hidden-screen target and let $\theta_\jmath$ denote one coordinate of the complete learned parameter set. Expanding the mean absolute error and its subgradient through the RGB residual path gives

$$\begin{aligned}
\mathcal{L}(\theta) &:= \frac{1}{3BHW}\sum_{\mathfrak{b}=1}^{B}\sum_{\kappa=1}^{3}\sum_{i=1}^{H}\sum_{j=1}^{W}\left|\widehat{y}_{\mathfrak{b}\kappa ij} - y_{\mathfrak{b}\kappa ij}\right|, \\
\frac{\partial \mathcal{L}}{\partial \theta_\jmath} &= \frac{1}{3BHW}\sum_{\mathfrak{b}=1}^{B}\sum_{\kappa=1}^{3}\sum_{i=1}^{H}\sum_{j=1}^{W}\varsigma_{\mathfrak{b}\kappa ij}\frac{\partial d_{\mathfrak{b}\kappa ij}}{\partial \theta_\jmath}, \\
\varsigma_{\mathfrak{b}\kappa ij} &\in \begin{cases} \{\operatorname{sgn}(\widehat{y}_{\mathfrak{b}\kappa ij} - y_{\mathfrak{b}\kappa ij})\}, & \widehat{y}_{\mathfrak{b}\kappa ij} \neq y_{\mathfrak{b}\kappa ij}, \\ [-1,1], & \widehat{y}_{\mathfrak{b}\kappa ij} = y_{\mathfrak{b}\kappa ij}. \end{cases}
\end{aligned} \tag{A7}$$

We train the complete network with this mean absolute reconstruction objective. Its gradient links the paired RGB error back through d, the descending decoder states, z^+, the trusted update u, the TraceTrust weight t, and the AGRL proposal r. At inference, the forward chain from the wall observation x to $\widehat{y}$ is evaluated with fixed learned parameters.

The encoder produces the wall representation z, AGRL forms the screen-oriented candidate residual r, and TraceTrust computes its spatial weight t. Their combination produces the bounded update u and the updated representation z^+. The middle block and decoder then produce the RGB correction d, yielding $\widehat{y} = x + d$. The paired reconstruction objective trains correspondence learning and trace weighting together to recover the displayed screen.

A.3 Supplementary Experimental Results

The supplementary experiments examine how learned correspondence and trace weighting support screen recovery. We first examine AGRL and TraceTrust, then vary the strength, spatial arrangement, and availability of the wall signal. The final comparisons examine how representation design and training objectives support reconstruction.

A.3.1 Axis-Guided Representation Learning

We compare channel-projection and spatial-aggregation variants of AGRL with the same wall inputs and screen targets.

Supplementary Table 1: Quantitative results for axis-guided representation learning. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	OSLock			PhotoDisplay			Screen2000			ScreenGame		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Channel projection												
w/o input projection	16.579	0.609	0.644	13.849	0.422	0.835	19.186	0.697	0.463	17.572	0.655	0.603
Input projection → low-rank projection	16.587	0.608	0.648	13.834	**0.425**	0.825	19.156	0.683	0.478	17.529	0.646	**0.600**
Input projection → grouped projection	16.482	0.605	0.644	13.776	0.424	0.835	**19.412**	0.687	**0.455**	**17.751**	0.655	0.601
Input projection → pointwise MLP	16.557	0.608	0.643	13.528	0.415	**0.821**	19.345	0.709	0.458	17.563	0.655	0.602
Spatial aggregation and output												
w/o spatial aggregation + output projection	16.560	0.608	**0.640**	13.719	0.422	0.829	19.347	**0.713**	0.466	17.576	0.649	0.610
5×5 depthwise variant	16.520	0.608	0.645	**13.861**	0.421	0.836	19.213	0.686	0.461	17.629	0.653	**0.600**
Spatial aggregation → dilated local convolution	16.445	0.604	0.643	13.839	0.423	0.836	19.209	0.704	0.460	17.707	**0.656**	**0.600**
Spatial aggregation + output projection → multi-scale residual	**16.648**	**0.610**	0.648	13.805	0.423	0.831	19.237	0.699	0.461	17.614	0.653	0.605
AGTRL	**16.710**	**0.613**	**0.620**	**13.925**	**0.426**	**0.806**	**24.706**	**0.837**	**0.259**	**18.639**	**0.684**	**0.533**

The AGRL ablation links feature organization to the recovery of screen layout. The complete design leads all 12 metric columns in Supplementary Table 1, and the Screen2000 examples show clearer separation between the window and the surrounding scene, together with more continuous major boundaries. Channel projection combines complementary feature responses, while spatial aggregation brings neighboring wall evidence into the same candidate update. Their combination helps organize the mixed wall signal into the regions and relationships of the displayed screen.

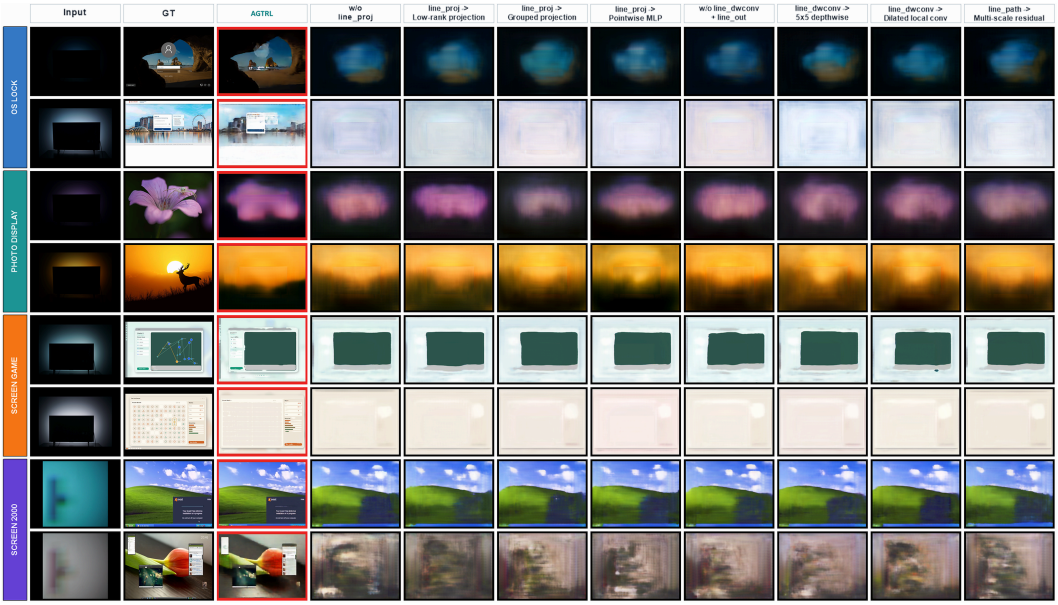

Supplementary Figure 9: Axis-Guided Representation Learning ablation on fixed samples. The upper group compares channel-projection variants; the lower group compares spatial aggregation and output variants. Each group repeats the same wall input, ground-truth screen, and AGTRL result so that the variant outputs can be read against the same reference.

A.3.2 TraceTrust Arbitration

We vary the learned TraceTrust gate and its continuity and artifact-suspicion components while keeping the remaining model fixed.

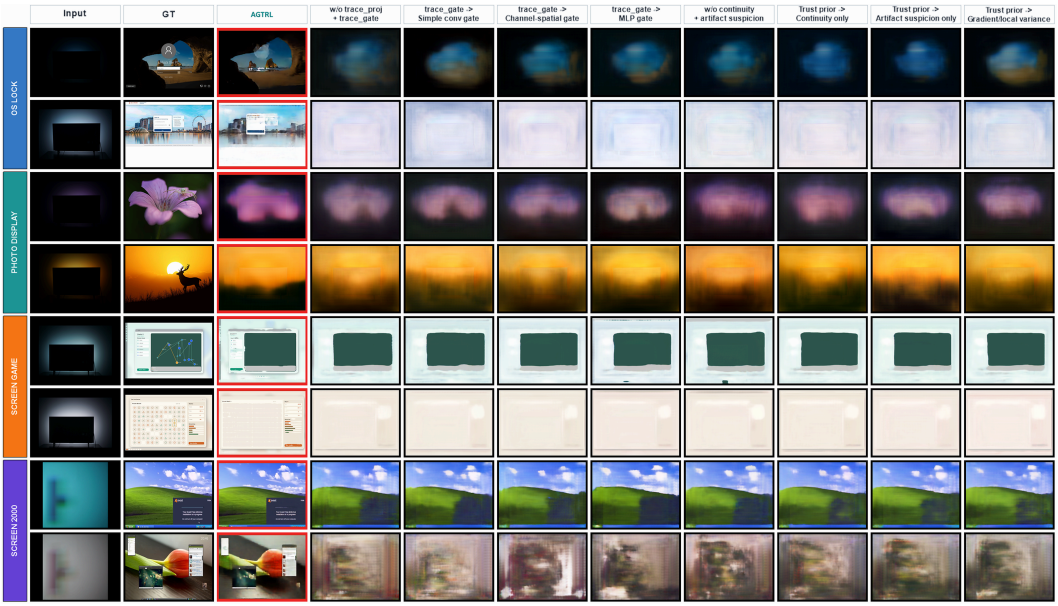

Supplementary Figure 10: TraceTrust arbitration ablation on fixed samples. The upper group compares learned-gate variants; the lower group compares continuity and artifact-suspicion variants. Each group repeats the same wall input, ground-truth screen, and AGTRL result so that the variant outputs can be read against the same reference.

The TraceTrust ablation shows the value of combining complementary evidence when selecting reconstruction updates. The complete design leads all 12 metric columns in Supplementary Table 2.

Supplementary Table 2: Quantitative results for TraceTrust arbitration. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	OSLock			PhotoDisplay			Screen2000			ScreenGame		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Learned gate												
w/o trust-feature projection + learned trust gate	**16.678**	**0.611**	0.642	13.877	**0.423**	0.834	19.500	**0.707**	0.464	17.573	0.645	0.606
TraceTrust → convolutional spatial gate	16.598	0.606	0.655	13.897	**0.423**	0.830	19.530	0.705	**0.456**	17.554	0.651	**0.601**
Learned trust gate → channel-spatial gate	16.505	0.603	0.649	13.814	**0.423**	0.831	19.442	0.693	0.467	**17.683**	**0.654**	0.603
Learned trust gate → MLP gate	16.603	**0.611**	**0.641**	13.554	0.418	**0.825**	19.368	0.698	0.459	17.474	0.648	0.605
Continuity and suspicion												
w/o continuity + artifact suspicion	16.457	0.607	0.646	13.791	0.421	0.832	19.491	0.680	0.458	17.545	0.645	0.605
Trust prior → continuity only	16.515	0.597	0.645	13.868	0.421	0.829	**19.534**	0.699	0.467	17.537	0.652	0.610
Trust prior → artifact suspicion only	16.675	0.609	0.645	**13.912**	**0.423**	0.829	19.354	0.682	0.466	17.584	0.651	0.607
Trust prior → gradient/local variance	16.581	0.603	0.645	13.900	**0.423**	0.828	19.528	0.684	0.464	17.594	0.652	0.607
AGTRL	**16.710**	**0.613**	**0.620**	**13.925**	**0.426**	**0.806**	**24.706**	**0.837**	**0.259**	**18.639**	**0.684**	**0.533**

On Screen2000, it achieves 24.706 dB PSNR, compared with 19.534 dB for the continuity-only trust-prior variant and 19.354 dB for the artifact-suspicion-only trust-prior variant. The accompanying reconstructions separate the screen regions more clearly. Continuity supplies neighborhood support, artifact suspicion measures unusually strong local departures, and the learned gate conditions their joint weighting on the current features. Together, these cues weight the candidate update when reconstructing screen structure from mixed wall evidence.

For the input-noise experiments, RGB wall intensities are scaled to $[0, 1]$. Let $v_x = \mathrm{Var}(x)$ and let s denote the nominal input SNR in dB. Gaussian noise has zero mean and variance $v_x 10^{-s/10}$, corresponding to $s = 10 \log_{10}(v_x/v_n)$, where v_n is the input-noise power. For salt-and-pepper noise, the total replacement probability is $p = \mathrm{clip}_{[0,1]}\{v_x 10^{-s/10}/[\mathbb{E}(x^2) + \frac{1}{2}(1 - 2\mathbb{E}(x))]\}$. A spatial mask replaces all RGB channels at a selected pixel by 1 or 0 with equal salt and pepper probabilities. These SNR labels specify the input perturbation before clipping and image quantization; output PSNR remains a reconstruction metric. The separate additive Gaussian-noise sweep uses σ in the same $[0, 1]$ intensity units and clips perturbed inputs to this range.

A.3.3 Gaussian Noise on PhotoDisplay

We evaluate fixed models on PhotoDisplay across the labeled Gaussian input-noise SNR settings.

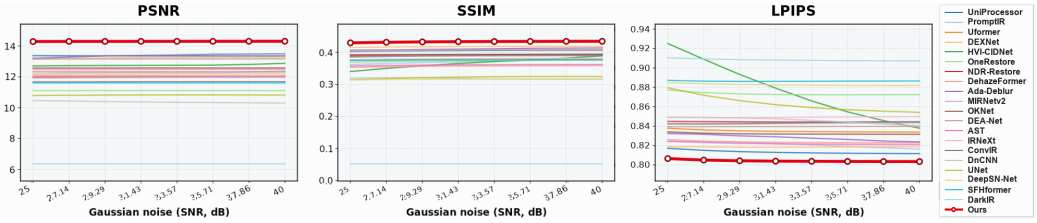

Supplementary Figure 11: Three-metric robustness comparison for Gaussian noise on PhotoDisplay. Gaussian input-noise SNR is varied from 25 to 40 dB. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

PhotoDisplay extends the noise analysis to natural-image composition and broad color transitions. AGTRL leads all three metrics, while its PSNR changes by only 0.017 dB across the sweep. The recovery advantage therefore extends from interface content to natural-image composition and broad color transitions.

A.3.4 Gaussian Noise on ScreenGame

We evaluate fixed models on ScreenGame across the labeled Gaussian input-noise SNR settings.

Game screens include dark regions and weak wall contrast. AGTRL retains the leading scores throughout the noise sweep and reaches 0.689 SSIM at the lowest input SNR. This structural agreement supports recovery of game content from weak wall traces under the tested noise conditions.

Supplementary Table 3: Quantitative results for Gaussian noise on PhotoDisplay. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	25			27.14			29.29			31.43			33.57			35.71			37.86			40		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	13.375	0.403	0.817	13.369	0.404	0.815	13.364	0.405	0.814	13.362	0.406	0.813	13.360	0.406	0.812	13.359	0.406	0.812	13.358	0.407	0.812	13.357	0.407	0.812
OneRestore	11.113	0.369	0.877	11.114	0.372	0.875	11.114	0.373	0.873	11.114	0.374	0.873	11.114	0.374	0.872	11.114	0.375	0.872	11.114	0.375	0.872	11.114	0.375	0.873
DeepSN-Net	12.201	0.314	0.885	12.208	0.315	0.884	12.213	0.316	0.883	12.218	0.317	0.883	12.221	0.317	0.882	12.223	0.317	0.882	12.224	0.317	0.882	12.224	0.317	0.882
PromptIR	12.197	0.362	0.824	12.239	0.365	0.824	12.265	0.368	0.823	12.270	0.370	0.822	12.265	0.372	0.821	12.268	0.375	0.819	12.309	0.378	0.818	12.390	0.381	0.816
Uformer	12.095	0.375	0.838	12.095	0.377	0.836	12.095	0.378	0.835	12.096	0.379	0.835	12.096	0.379	0.834	12.096	0.380	0.834	12.096	0.380	0.834	12.096	0.380	0.834
MIRNetv2	11.924	0.404	0.842	11.939	0.405	0.842	11.950	0.407	0.842	11.958	0.407	0.843	11.965	0.408	0.843	11.971	0.408	0.843	11.975	0.408	0.843	11.977	0.408	0.843
OKNet	12.571	0.389	0.834	12.573	0.391	0.833	12.574	0.392	0.832	12.575	0.393	0.832	12.576	0.394	0.832	12.576	0.394	0.832	12.577	0.394	0.832	12.577	0.394	0.832
AST	12.326	0.357	0.824	12.327	0.359	0.823	12.328	0.360	0.822	12.329	0.361	0.822	12.329	0.361	0.821	12.330	0.362	0.821	12.330	0.362	0.821	12.330	0.362	0.821
ConvIR	11.681	0.386	0.842	11.683	0.387	0.842	11.686	0.388	0.843	11.687	0.389	0.843	11.689	0.390	0.844	11.690	0.390	0.844	11.691	0.390	0.844	11.692	0.390	0.845
DnCNN	10.461	0.320	0.849	10.438	0.321	0.849	10.405	0.321	0.849	10.377	0.319	0.848	10.356	0.317	0.846	10.340	0.316	0.844	10.322	0.316	0.842	10.305	0.316	0.841
U-Net	10.795	0.314	0.880	10.810	0.318	0.872	10.822	0.321	0.866	10.830	0.323	0.862	10.834	0.324	0.859	10.834	0.325	0.857	10.830	0.325	0.855	10.825	0.325	0.854
DExNet	13.154	0.415	0.819	13.199	0.417	0.819	13.228	0.418	0.818	13.246	0.418	0.818	13.257	0.419	0.818	13.263	0.419	0.819	13.266	0.419	0.819	13.268	0.419	0.819
HVI-CIDNet	12.711	0.340	0.925	12.728	0.350	0.909	12.739	0.358	0.893	12.747	0.365	0.879	12.754	0.371	0.866	12.763	0.376	0.855	12.806	0.382	0.846	12.879	0.389	0.838
NDR-Restore	11.989	0.393	0.845	11.994	0.393	0.845	11.997	0.393	0.844	11.999	0.393	0.844	12.001	0.393	0.844	12.002	0.393	0.844	12.002	0.393	0.844	12.003	0.393	0.844
DehazeFormer	12.467	0.353	0.826	12.472	0.355	0.825	12.476	0.356	0.824	12.478	0.357	0.823	12.480	0.358	0.823	12.482	0.358	0.823	12.483	0.359	0.823	12.484	0.359	0.823
Ada-Deblur	13.210	0.405	0.832	13.283	0.407	0.832	13.365	0.409	0.830	13.393	0.410	0.829	13.420	0.411	0.828	13.468	0.413	0.826	13.483	0.413	0.825	13.495	0.414	0.824
DEA-Net	13.157	0.407	0.839	13.160	0.408	0.839	13.162	0.408	0.840	13.164	0.408	0.840	13.164	0.408	0.840	13.165	0.408	0.840	13.164	0.408	0.840	13.164	0.408	0.840
IRNeXt	12.070	0.377	0.849	12.070	0.379	0.849	12.071	0.380	0.849	12.070	0.380	0.849	12.071	0.381	0.849	12.070	0.381	0.850	12.070	0.381	0.850	12.070	0.381	0.850
SFHformer	11.607	0.375	0.887	11.608	0.376	0.886	11.609	0.376	0.886	11.610	0.376	0.886	11.610	0.377	0.886	11.611	0.377	0.886	11.611	0.377	0.887	11.611	0.377	0.887
DarkIR	6.370	0.052	0.911	6.370	0.052	0.909	6.369	0.052	0.909	6.369	0.052	0.908	6.369	0.052	0.908	6.369	0.052	0.908	6.369	0.052	0.907	6.369	0.052	0.907
AGTRL	**14.290**	**0.430**	**0.807**	**14.298**	**0.432**	**0.805**	**14.299**	**0.433**	**0.804**	**14.302**	**0.433**	**0.804**	**14.304**	**0.434**	**0.804**	**14.305**	**0.434**	**0.804**	**14.306**	**0.434**	**0.804**	**14.307**	**0.435**	**0.803**

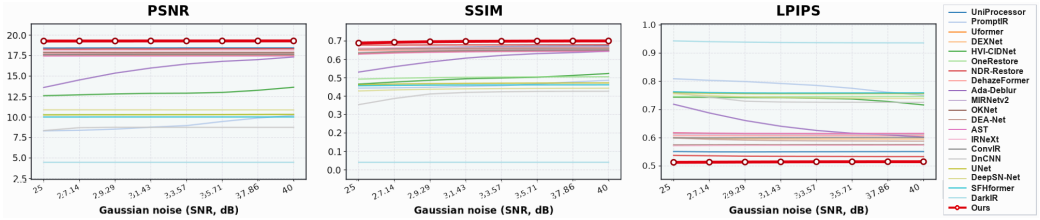

Supplementary Figure 12: Three-metric robustness comparison for Gaussian noise on ScreenGame. Gaussian input-noise SNR is varied from 25 to 40 dB. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 4: Quantitative results for Gaussian noise on ScreenGame. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	25			27.14			29.29			31.43			33.57			35.71			37.86			40		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	18.431	0.655	0.551	18.437	0.661	0.550	18.447	0.665	0.550	18.442	0.668	0.550	18.449	0.670	0.551	18.448	0.671	0.551	18.449	0.671	0.551	18.450	0.672	0.551
OneRestore	10.287	0.493	0.747	10.295	0.498	0.746	10.301	0.501	0.745	10.305	0.503	0.746	10.308	0.504	0.746	10.309	0.505	0.746	10.310	0.505	0.747	10.311	0.506	0.747
DeepSN-Net	10.883	0.428	0.756	10.880	0.434	0.750	10.878	0.438	0.745	10.876	0.440	0.743	10.875	0.441	0.741	10.873	0.442	0.740	10.872	0.443	0.740	10.871	0.443	0.740
PromptIR	8.312	0.444	0.809	8.379	0.446	0.804	8.519	0.449	0.799	8.761	0.454	0.793	8.954	0.457	0.786	9.451	0.467	0.775	9.877	0.477	0.762	10.194	0.487	0.750
Uformer	17.536	0.630	0.617	17.544	0.636	0.616	17.544	0.640	0.614	17.545	0.643	0.614	17.545	0.644	0.614	17.546	0.645	0.614	17.545	0.646	0.614	17.547	0.646	0.614
MIRNetv2	17.663	0.655	0.613	17.671	0.656	0.613	17.671	0.656	0.613	17.680	0.656	0.613	17.679	0.656	0.612	17.681	0.656	0.612	17.680	0.656	0.612	17.680	0.656	0.612
OKNet	17.901	0.657	0.574	17.899	0.659	0.575	17.889	0.660	0.575	17.892	0.661	0.575	17.894	0.662	0.575	17.893	0.662	0.575	17.893	0.662	0.575	17.894	0.662	0.575
AST	17.447	0.628	0.618	17.451	0.635	0.617	17.454	0.638	0.616	17.449	0.641	0.616	17.454	0.643	0.615	17.455	0.643	0.616	17.454	0.644	0.616	17.457	0.645	0.616
ConvIR	17.615	0.636	0.600	17.633	0.641	0.600	17.637	0.645	0.600	17.635	0.647	0.601	17.639	0.649	0.601	17.640	0.650	0.601	17.640	0.650	0.601	17.644	0.651	0.601
DnCNN	8.345	0.354	0.760	8.703	0.388	0.744	8.750	0.412	0.730	8.745	0.420	0.727	8.743	0.424	0.726	8.740	0.426	0.708	8.739	0.427	0.726	8.738	0.427	0.726
U-Net	10.257	0.462	0.760	10.273	0.466	0.758	10.275	0.468	0.757	10.280	0.470	0.756	10.291	0.471	0.756	10.305	0.471	0.757	10.318	0.472	0.757	10.322	0.472	0.757
DExNet	17.512	0.651	0.600	17.587	0.655	0.597	17.617	0.656	0.597	17.627	0.657	0.596	17.632	0.658	0.596	17.633	0.658	0.596	17.633	0.658	0.596	17.631	0.658	0.595
HVI-CIDNet	12.605	0.466	0.744	12.712	0.477	0.744	12.823	0.486	0.743	12.891	0.495	0.742	12.913	0.499	0.740	13.002	0.503	0.738	13.265	0.513	0.729	13.638	0.523	0.716
NDR-Restore	18.282	0.678	0.538	18.280	0.678	0.536	18.312	0.678	0.535	18.301	0.679	0.534	18.320	0.679	0.534	18.330	0.679	0.534	18.332	0.679	0.534	18.336	0.679	0.533
DehazeFormer	17.762	0.648	0.608	17.771	0.648	0.607	17.769	0.649	0.607	17.771	0.649	0.607	17.772	0.649	0.607	17.772	0.649	0.607	17.774	0.649	0.607	17.774	0.649	0.607
Ada-Deblur	13.613	0.531	0.719	14.530	0.560	0.688	15.367	0.586	0.661	15.988	0.607	0.641	16.475	0.622	0.626	16.799	0.631	0.616	17.012	0.637	0.610	17.338	0.645	0.602
DEA-Net	17.601	0.654	0.600	17.654	0.656	0.595	17.686	0.656	0.592	17.704	0.657	0.590	17.718	0.658	0.589	17.720	0.658	0.589	17.730	0.658	0.589	17.733	0.658	0.588
IRNeXt	18.122	0.661	0.572	18.128	0.664	0.572	18.120	0.665	0.573	18.142	0.666	0.573	18.138	0.667	0.573	18.146	0.668	0.573	18.141	0.668	0.573	18.135	0.668	0.574
SFHformer	10.004	0.457	0.763	10.006	0.459	0.760	10.008	0.460	0.759	10.009	0.461	0.759	10.010	0.461	0.759	10.010	0.462	0.759	10.011	0.462	0.759	10.011	0.462	0.759
DarkIR	4.467	0.041	0.943	4.466	0.041	0.941	4.466	0.041	0.939	4.465	0.041	0.938	4.465	0.041	0.937	4.465	0.041	0.937	4.465	0.041	0.937	4.465	0.041	0.936
AGTRL	**19.281**	**0.689**	**0.513**	**19.287**	**0.693**	**0.514**	**19.287**	**0.696**	**0.514**	**19.294**	**0.698**	**0.515**	**19.292**	**0.699**	**0.515**	**19.296**	**0.699**	**0.515**	**19.297**	**0.700**	**0.515**	**19.297**	**0.700**	**0.515**

A.3.5 Salt-and-Pepper Noise on OSLock

We vary salt-and-pepper input noise on OSLock while keeping the learned model fixed.

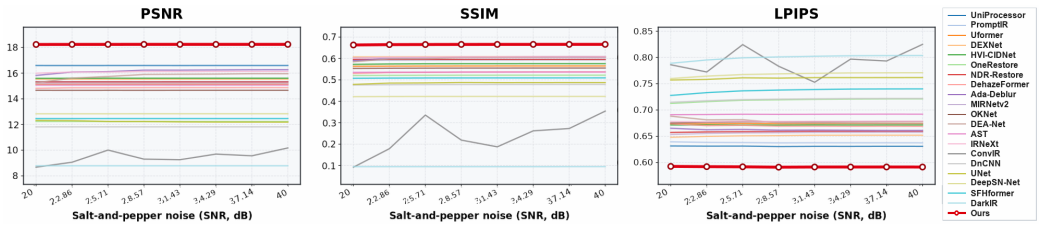

Supplementary Figure 13: Three-metric robustness comparison for salt-and-pepper noise on OSLock. Salt-and-pepper input-noise SNR is varied from 20 to 40 dB. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 5: Quantitative results for salt-and-pepper noise on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	20			22.86			25.71			28.57			31.43			34.29			37.14			40		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**16.595**	0.603	**0.631**	**16.593**	0.604	**0.631**	**16.595**	0.605	**0.631**	**16.596**	0.605	**0.630**	**16.595**	0.605	**0.631**	**16.596**	0.606	**0.631**	**16.595**	0.606	**0.631**	**16.595**	0.606	**0.631**
OneRestore	12.239	0.520	0.712	12.239	0.521	0.716	12.238	0.522	0.719	12.238	0.522	0.719	12.237	0.522	0.720	12.237	0.522	0.721	12.237	0.522	0.721	12.237	0.522	0.721
DeepSN-Net	12.834	0.422	0.760	12.835	0.422	0.765	12.836	0.422	0.767	12.836	0.423	0.769	12.836	0.423	0.770	12.836	0.423	0.771	12.836	0.423	0.771	12.836	0.423	0.771
PromptIR	15.344	0.591	0.639	15.344	0.592	0.638	15.343	0.593	0.638	15.342	0.593	0.637	15.342	0.593	0.637	15.342	0.593	0.637	15.342	0.594	0.637	15.342	0.594	0.637
Uformer	15.147	0.564	0.674	15.147	0.564	0.674	15.146	0.565	0.674	15.150	0.565	0.674	15.150	0.565	0.674	15.149	0.565	0.674	15.149	0.565	0.674	15.149	0.565	0.674
MIRNetv2	15.210	0.602	0.653	15.210	0.603	0.655	15.211	0.604	0.656	15.211	0.604	0.657	15.211	0.604	0.657	15.211	0.604	0.657	15.211	0.604	0.657	15.211	0.604	0.657
OKNet	14.665	0.554	0.675	14.662	0.554	0.677	14.664	0.555	0.678	14.663	0.555	0.677	14.664	0.555	0.678	14.664	0.555	0.678	14.664	0.555	0.678	14.664	0.555	0.678
AST	15.071	0.535	0.691	15.072	0.536	0.691	15.072	0.536	0.692	15.072	0.536	0.692	15.072	0.536	0.692	15.072	0.537	0.692	15.072	0.537	0.692	15.072	0.537	0.692
ConvIR	8.678	0.092	0.786	9.064	0.179	0.772	10.009	0.336	0.825	9.304	0.219	0.783	9.259	0.187	0.753	9.698	0.262	0.797	9.565	0.273	0.793	10.175	0.354	0.825
DnCNN	11.816	0.476	0.715	11.817	0.477	0.718	11.817	0.478	0.720	11.817	0.478	0.721	11.818	0.478	0.721	11.818	0.479	0.722	11.818	0.479	0.722	11.818	0.479	0.722
U-Net	12.323	0.479	0.757	12.300	0.485	0.758	12.239	0.486	0.761	12.236	0.488	0.761	12.203	0.487	0.762	12.190	0.487	0.762	12.186	0.487	0.762	12.185	0.487	0.762
DExNet	15.518	**0.607**	0.648	15.518	**0.608**	0.650	15.519	**0.608**	0.651	15.518	**0.608**	0.651	15.518	**0.608**	0.652	15.518	**0.608**	0.652	15.518	**0.608**	0.652	15.518	**0.608**	0.652
HVI-CIDNet	15.597	0.573	0.672	15.603	0.574	0.672	15.608	0.575	0.671	15.606	0.575	0.671	15.606	0.576	0.672	15.605	0.576	0.671	15.606	0.576	0.672	15.606	0.576	0.671
NDR-Restore	15.319	0.595	0.657	15.319	0.595	0.658	15.320	0.595	0.659	15.319	0.595	0.659	15.319	0.595	0.659	15.319	0.595	0.659	15.319	0.595	0.659	15.319	0.595	0.659
DehazeFormer	14.818	0.530	0.674	14.862	0.534	0.670	14.862	0.535	0.671	14.871	0.537	0.669	14.871	0.537	0.669	14.872	0.537	0.669	14.874	0.537	0.669	14.873	0.537	0.669
Ada-Deblur	15.822	0.589	0.665	16.088	0.600	0.662	16.119	0.603	0.663	16.235	0.606	0.662	16.234	0.607	0.662	16.252	**0.608**	0.661	16.268	**0.608**	0.661	16.271	**0.608**	0.661
DEA-Net	15.257	0.585	0.688	15.610	0.598	0.681	15.741	0.599	0.681	15.896	0.605	0.676	15.914	0.605	0.675	15.928	0.606	0.675	15.957	0.607	0.675	15.956	0.607	0.674
IRNeXt	15.995	0.602	0.678	16.089	0.606	0.677	16.085	0.606	0.679	16.131	**0.608**	0.678	16.124	**0.608**	0.678	16.132	**0.608**	0.678	16.132	**0.608**	0.679	16.132	**0.608**	0.678
SFHformer	12.460	0.508	0.728	12.460	0.508	0.733	12.461	0.509	0.737	12.461	0.509	0.738	12.461	0.509	0.739	12.461	0.509	0.740	12.461	0.509	0.740	12.461	0.509	0.740
DarkIR	8.788	0.095	0.789	8.788	0.095	0.795	8.787	0.095	0.799	8.787	0.095	0.801	8.787	0.095	0.802	8.787	0.095	0.803	8.787	0.095	0.804	8.787	0.095	0.804
AGTRL	**18.233**	**0.663**	**0.593**	**18.238**	**0.665**	**0.592**	**18.240**	**0.665**	**0.592**	**18.241**	**0.666**	**0.591**	**18.241**	**0.666**	**0.591**	**18.240**	**0.666**	**0.591**	**18.242**	**0.666**	**0.591**	**18.242**	**0.666**	**0.591**

Salt-and-pepper noise adds isolated, high-amplitude disturbances to the wall image. AGTRL retains the leading scores across the tested levels, extending the stable recovery observed under distributed Gaussian fluctuations to sparse impulses.

A.3.6 Salt-and-Pepper Noise on PhotoDisplay

We vary salt-and-pepper input noise on PhotoDisplay while keeping the learned model fixed.

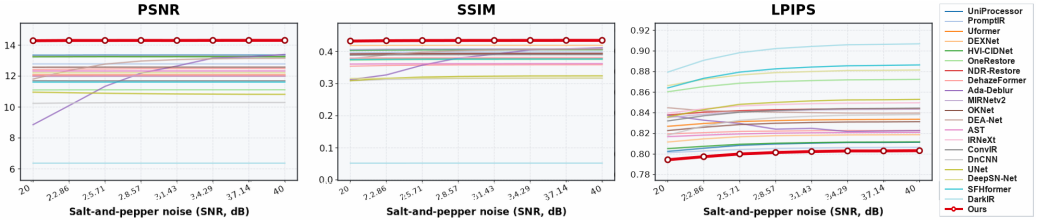

Supplementary Figure 14: Three-metric robustness comparison for salt-and-pepper noise on PhotoDisplay. Salt-and-pepper input-noise SNR is varied from 20 to 40 dB. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 6: Quantitative results for salt-and-pepper noise on PhotoDisplay. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	20			22.86			25.71			28.57			31.43			34.29			37.14			40		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**13.350**	0.405	0.803	**13.352**	0.406	0.806	**13.357**	0.406	0.808	**13.355**	0.407	0.810	**13.358**	0.407	0.811	**13.357**	0.407	0.811	**13.357**	0.407	0.811	13.357	0.407	0.811
OneRestore	11.104	0.373	0.860	11.108	0.374	0.865	11.111	0.375	0.869	11.113	0.375	0.870	11.113	0.375	0.871	11.114	0.375	0.872	11.114	0.376	0.872	11.114	0.376	0.872
DeepSN-Net	12.223	0.316	0.866	12.223	0.317	0.873	12.224	0.317	0.877	12.224	0.317	0.879	12.224	0.317	0.880	12.224	0.317	0.881	12.224	0.317	0.881	12.224	0.317	0.882
PromptIR	12.787	0.396	**0.801**	12.786	0.398	**0.803**	12.786	0.398	**0.804**	12.786	0.399	**0.805**	12.785	0.399	**0.806**	12.785	0.399	**0.806**	12.785	0.399	**0.806**	12.785	0.399	**0.807**
Uformer	12.095	0.379	0.827	12.096	0.380	0.830	12.096	0.380	0.832	12.096	0.380	0.832	12.096	0.380	0.833	12.096	0.380	0.833	12.096	0.380	0.834	12.096	0.380	0.834
MIRNetv2	11.971	0.406	0.835	11.975	0.407	0.839	11.977	0.408	0.841	11.978	0.408	0.842	11.979	0.408	0.843	11.979	0.408	0.843	11.979	0.409	0.843	11.980	0.409	0.843
OKNet	12.572	0.392	0.823	12.573	0.393	0.826	12.576	0.394	0.829	12.576	0.394	0.830	12.576	0.394	0.831	12.577	0.394	0.831	12.577	0.394	0.831	12.577	0.394	0.831
AST	12.328	0.361	0.817	12.330	0.362	0.818	12.329	0.362	0.820	12.330	0.362	0.820	12.330	0.362	0.821	12.330	0.362	0.821	12.330	0.362	0.821	12.330	0.362	0.821
ConvIR	11.698	0.388	0.832	11.694	0.389	0.837	11.694	0.390	0.841	11.693	0.390	0.842	11.693	0.390	0.843	11.693	0.390	0.844	11.693	0.390	0.844	11.692	0.390	0.845
DnCNN	10.234	0.309	0.818	10.263	0.313	0.827	10.275	0.315	0.833	10.282	0.316	0.835	10.286	0.316	0.837	10.288	0.317	0.838	10.289	0.317	0.838	10.290	0.317	0.838
U-Net	10.959	0.309	0.836	10.923	0.317	0.843	10.889	0.320	0.848	10.863	0.322	0.850	10.844	0.323	0.852	10.833	0.324	0.852	10.822	0.324	0.853	10.820	0.324	0.853
DExNet	13.267	**0.419**	0.812	13.268	**0.419**	0.815	13.268	**0.419**	0.817	13.269	**0.419**	0.818	13.270	**0.419**	0.818	13.269	**0.419**	0.819	13.269	**0.419**	0.819	13.269	**0.419**	0.819
HVI-CIDNet	13.242	0.404	0.805	13.247	0.405	0.808	13.251	0.405	0.810	13.249	0.405	0.811	13.250	0.405	0.811	13.248	0.405	0.812	13.249	0.405	0.812	13.250	0.406	0.812
NDR-Restore	11.999	0.392	0.838	12.001	0.393	0.841	12.002	0.393	0.842	12.002	0.393	0.843	12.003	0.393	0.843	12.003	0.393	0.844	12.004	0.393	0.844	12.004	0.393	0.844
DehazeFormer	12.424	0.354	0.820	12.446	0.357	0.821	12.457	0.357	0.822	12.475	0.358	0.822	12.475	0.358	0.823	12.483	0.359	0.823	12.483	0.359	0.823	12.484	0.359	0.823
Ada-Deblur	8.858	0.312	0.838	10.089	0.327	0.833	11.336	0.358	0.830	12.189	0.380	0.824	12.653	0.392	0.825	13.162	0.405	0.822	13.243	0.408	0.822	**13.401**	0.412	0.823
DEA-Net	11.806	0.378	0.845	12.362	0.389	0.842	12.775	0.398	0.841	12.972	0.403	0.841	13.060	0.405	0.840	13.122	0.406	0.840	13.137	0.407	0.840	13.153	0.407	0.840
IRNeXt	12.070	0.380	0.840	12.070	0.381	0.844	12.070	0.381	0.847	12.070	0.381	0.848	12.070	0.381	0.849	12.070	0.381	0.849	12.070	0.382	0.850	12.070	0.382	0.850
SFHformer	11.610	0.375	0.864	11.610	0.376	0.874	11.611	0.376	0.880	11.611	0.377	0.883	11.611	0.377	0.884	11.611	0.377	0.886	11.611	0.377	0.886	11.611	0.377	0.886
DarkIR	6.369	0.052	0.880	6.369	0.052	0.891	6.369	0.052	0.898	6.369	0.052	0.902	6.369	0.052	0.904	6.369	0.052	0.906	6.369	0.052	0.906	6.369	0.052	0.907
AGTRL	**14.285**	**0.433**	**0.794**	**14.300**	**0.434**	**0.797**	**14.302**	**0.434**	**0.800**	**14.303**	**0.434**	**0.802**	**14.302**	**0.435**	**0.802**	**14.306**	**0.435**	**0.803**	**14.307**	**0.435**	**0.803**	**14.308**	**0.435**	**0.803**

On PhotoDisplay, the reconstruction advantage persists in the presence of isolated impulses. AGTRL's structural similarity stays between 0.433 and 0.435 at the displayed precision, showing stable structural recovery of natural-image content under impulse noise.

A.3.7 Salt-and-Pepper Noise on ScreenGame

We vary salt-and-pepper input noise on ScreenGame while keeping the learned model fixed.

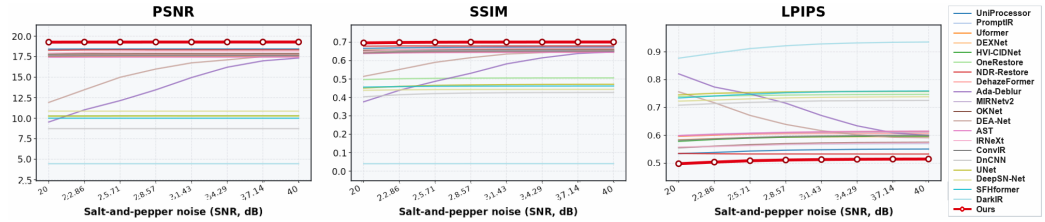

Supplementary Figure 15: Three-metric robustness comparison for salt-and-pepper noise on ScreenGame. Salt-and-pepper input-noise SNR is varied from 20 to 40 dB. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 7: Quantitative results for salt-and-pepper noise on ScreenGame. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	20			22.86			25.71			28.57			31.43			34.29			37.14			40		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	18.430	0.666	0.535	18.445	0.669	0.539	18.442	0.671	0.544	18.446	0.671	0.547	18.447	0.672	0.549	18.446	0.672	0.550	18.451	0.672	0.551	18.450	0.672	0.551
OneRestore	10.291	0.497	0.739	10.302	0.501	0.741	10.307	0.504	0.743	10.310	0.505	0.745	10.312	0.506	0.746	10.312	0.506	0.747	10.313	0.506	0.747	10.313	0.506	0.747
DeepSN-Net	10.876	0.439	0.723	10.871	0.441	0.727	10.870	0.443	0.731	10.869	0.443	0.735	10.868	0.444	0.737	10.868	0.444	0.738	10.868	0.444	0.738	10.868	0.444	0.739
PromptIR	17.773	0.653	0.556	17.782	0.656	0.561	17.772	0.658	0.564	17.782	0.659	0.567	17.782	0.659	0.569	17.787	0.660	0.570	17.783	0.660	0.570	17.784	0.660	0.571
Uformer	17.536	0.640	0.597	17.541	0.644	0.601	17.542	0.645	0.606	17.546	0.646	0.609	17.547	0.647	0.611	17.546	0.647	0.613	17.548	0.647	0.613	17.546	0.647	0.614
MIRNetv2	17.676	0.655	0.600	17.679	0.655	0.605	17.678	0.656	0.608	17.679	0.656	0.610	17.678	0.656	0.611	17.679	0.656	0.611	17.679	0.656	0.612	17.679	0.656	0.612
OKNet	17.861	0.654	0.556	17.880	0.658	0.561	17.883	0.660	0.567	17.893	0.661	0.570	17.890	0.662	0.573	17.889	0.662	0.574	17.891	0.662	0.575	17.891	0.662	0.575
AST	17.435	0.638	0.599	17.447	0.642	0.604	17.453	0.643	0.608	17.453	0.644	0.611	17.453	0.645	0.613	17.455	0.645	0.614	17.455	0.645	0.615	17.453	0.645	0.615
ConvIR	17.604	0.641	0.583	17.629	0.646	0.588	17.638	0.649	0.593	17.639	0.650	0.596	17.638	0.651	0.598	17.642	0.651	0.600	17.642	0.651	0.600	17.641	0.651	0.601
DnCNN	8.737	0.404	0.708	8.740	0.415	0.714	8.739	0.421	0.718	8.738	0.424	0.721	8.738	0.426	0.724	8.738	0.427	0.725	8.738	0.427	0.725	8.738	0.427	0.725
U-Net	10.262	0.451	0.745	10.263	0.460	0.751	10.275	0.466	0.754	10.281	0.468	0.756	10.280	0.470	0.757	10.285	0.471	0.757	10.287	0.471	0.758	10.285	0.471	0.758
DExNet	17.625	0.657	0.585	17.627	0.658	0.589	17.628	0.658	0.592	17.629	0.658	0.594	17.629	0.658	0.594	17.629	0.659	0.595	17.629	0.659	0.595	17.630	0.659	0.595
HVI-CIDNet	17.766	0.650	0.579	17.772	0.652	0.585	17.774	0.654	0.591	17.771	0.655	0.594	17.772	0.655	0.596	17.771	0.655	0.597	17.774	0.655	0.597	17.774	0.655	0.598
NDR-Restore	18.293	0.678	0.535	18.335	0.679	0.534	18.351	0.679	0.533	18.338	0.679	0.533	18.340	0.679	0.533	18.342	0.679	0.533	18.345	0.680	0.533	18.344	0.679	0.533
DehazeFormer	17.661	0.644	0.598	17.719	0.647	0.602	17.741	0.648	0.604	17.763	0.649	0.606	17.757	0.649	0.607	17.767	0.649	0.607	17.773	0.649	0.607	17.773	0.649	0.607
Ada-Deblur	9.546	0.376	0.821	11.033	0.439	0.774	12.153	0.487	0.749	13.468	0.531	0.716	14.964	0.582	0.671	16.226	0.615	0.634	17.001	0.639	0.609	17.345	0.648	0.600
DEA-Net	11.933	0.514	0.756	13.479	0.552	0.717	14.999	0.590	0.672	16.002	0.616	0.639	16.755	0.636	0.617	17.119	0.644	0.602	17.522	0.653	0.593	17.619	0.656	0.592
IRNeXt	17.939	0.649	0.558	18.044	0.658	0.560	18.098	0.663	0.563	18.120	0.666	0.568	18.125	0.667	0.570	18.132	0.667	0.572	18.131	0.668	0.572	18.124	0.668	0.573
SFHformer	10.007	0.456	0.734	10.008	0.459	0.742	10.009	0.460	0.748	10.010	0.461	0.753	10.010	0.462	0.756	10.010	0.462	0.758	10.010	0.462	0.759	10.010	0.462	0.759
DarkIR	4.466	0.041	0.877	4.465	0.041	0.895	4.465	0.041	0.911	4.465	0.041	0.922	4.465	0.041	0.928	4.465	0.041	0.932	4.465	0.041	0.934	4.465	0.041	0.935
AGTRL	19.279	0.696	0.499	19.288	0.698	0.504	19.289	0.700	0.509	19.292	0.700	0.512	19.291	0.700	0.513	19.296	0.701	0.514	19.295	0.701	0.515	19.297	0.701	0.515

AGTRL retains the leading reconstruction quality for game screens under impulse noise. The comparison extends its recovery advantage to wall observations containing both weak screen traces and locally strong disturbances.

A.3.8 Distance Robustness: Acquisition Setup

Supplementary Figure 16 documents the physical setup for varying the camera's observation distance in the display-self-occluded configuration. The photographs show different camera placements relative to the display and relay wall. This geometry provides context for the distance-robustness analysis on OSLock, PhotoDisplay, and ScreenGame in Supplementary Figures 17–19. The reported distance values denote actual acquisition distances in meters (m). The corresponding plots and tables report recovery at the specified distance settings.



Supplementary Figure 16: Real acquisition photographs for the distance-robustness setting. The six panels show camera placements and setup adjustments at different observation distances relative to the display and relay wall. The screen faces the wall, and the camera observes the reflected illumination from behind the display. The photographs illustrate the acquisition geometry; the numerical distance settings, in m, are given in the corresponding result plots and tables.

A.3.9 Distance Degradation on OSLock

We compare fixed models over the labeled distance settings from 1.5 to 2.0 m on OSLock.

AGTRL leads throughout the distance sweep. Its PSNR stays within 18.139–18.161 dB, showing stable interface recovery as the distance setting changes.

A.3.10 Distance Degradation on PhotoDisplay

We compare fixed models over the labeled distance settings from 1.5 to 2.0 m on PhotoDisplay.

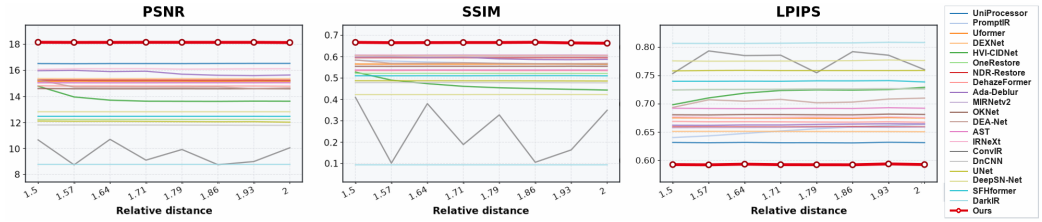

Supplementary Figure 17: Three-metric robustness comparison for distance degradation on OSLock. The labeled distance setting spans 1.5 to 2.0 m. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 8: Quantitative results for distance degradation on OSLock. Distance is reported in m; PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	1.5			1.57			1.64			1.71			1.79			1.86			1.93			2		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	16.525	0.607	0.632	16.515	0.606	0.632	16.526	0.607	0.632	16.527	0.608	0.632	16.523	0.606	0.632	16.532	0.607	0.631	16.536	0.607	0.632	16.535	0.606	0.632
OneRestore	12.237	0.522	0.725	12.237	0.522	0.725	12.236	0.522	0.726	12.236	0.522	0.726	12.236	0.522	0.726	12.236	0.522	0.726	12.235	0.522	0.727	12.235	0.521	0.727
DeepSN-Net	12.827	0.423	0.776	12.827	0.423	0.775	12.828	0.423	0.776	12.827	0.423	0.776	12.828	0.423	0.776	12.827	0.423	0.776	12.828	0.423	0.778	12.827	0.423	0.777
PromptIR	15.267	0.584	0.641	15.240	0.579	0.644	15.190	0.575	0.648	15.143	0.572	0.652	15.119	0.569	0.656	15.126	0.567	0.659	15.139	0.565	0.662	15.162	0.563	0.664
Uformer	15.160	0.566	0.676	15.171	0.565	0.675	15.168	0.565	0.675	15.175	0.565	0.675	15.172	0.566	0.675	15.161	0.566	0.675	15.177	0.566	0.676	15.168	0.565	0.675
MIRNetv2	15.062	0.600	0.658	15.053	0.599	0.658	15.056	0.599	0.658	15.056	0.599	0.659	15.058	0.599	0.659	15.061	0.600	0.659	15.065	0.600	0.659	15.067	0.600	0.659
OKNet	14.608	0.555	0.681	14.591	0.555	0.681	14.589	0.555	0.681	14.589	0.555	0.681	14.594	0.555	0.681	14.590	0.555	0.681	14.597	0.555	0.682	14.591	0.555	0.681
AST	15.052	0.538	0.692	15.050	0.538	0.692	15.047	0.538	0.692	15.048	0.538	0.692	15.053	0.538	0.692	15.053	0.538	0.692	15.052	0.538	0.693	15.056	0.537	0.692
ConvIR	10.664	0.410	0.754	8.751	0.102	0.794	10.708	0.380	0.785	9.112	0.189	0.786	9.928	0.328	0.755	8.760	0.106	0.792	9.005	0.164	0.786	10.060	0.349	0.761
DnCNN	11.810	0.479	0.725	11.810	0.479	0.725	11.808	0.479	0.725	11.806	0.479	0.726	11.801	0.479	0.726	11.796	0.479	0.727	11.789	0.478	0.727	11.781	0.478	0.727
U-Net	12.112	0.488	0.758	12.095	0.488	0.759	12.099	0.488	0.759	12.083	0.487	0.759	12.068	0.487	0.759	12.056	0.487	0.759	12.062	0.487	0.759	12.023	0.486	0.759
DExNet	15.370	0.604	0.651	15.362	0.604	0.652	15.364	0.604	0.652	15.364	0.604	0.651	15.363	0.604	0.651	15.364	0.604	0.651	15.364	0.604	0.651	15.365	0.604	0.651
HVI-CIDNet	14.791	0.528	0.699	13.961	0.490	0.711	13.716	0.474	0.719	13.637	0.461	0.724	13.629	0.455	0.725	13.624	0.451	0.724	13.642	0.447	0.725	13.635	0.444	0.729
NDR-Restore	15.260	0.595	0.659	15.257	0.595	0.659	15.258	0.595	0.659	15.257	0.595	0.659	15.253	0.595	0.660	15.252	0.595	0.660	15.254	0.595	0.660	15.252	0.595	0.660
DehazeFormer	14.815	0.535	0.669	14.803	0.535	0.668	14.805	0.535	0.669	14.799	0.535	0.668	14.805	0.535	0.668	14.801	0.535	0.668	14.802	0.535	0.669	14.787	0.534	0.668
Ada-Deblur	15.978	0.599	0.662	16.006	0.600	0.662	15.911	0.596	0.663	15.938	0.598	0.663	15.710	0.590	0.664	15.627	0.587	0.664	15.600	0.586	0.665	15.647	0.587	0.665
DEA-Net	15.265	0.584	0.694	14.720	0.569	0.707	14.723	0.570	0.705	14.719	0.572	0.708	14.715	0.569	0.702	14.727	0.567	0.703	14.600	0.567	0.709	14.657	0.568	0.710
IRNeXt	16.080	0.608	0.677	16.123	0.608	0.676	16.130	0.608	0.676	16.133	0.608	0.676	16.099	0.608	0.677	16.110	0.608	0.677	16.123	0.608	0.677	16.131	0.608	0.676
SFHformer	12.473	0.511	0.740	12.473	0.511	0.740	12.473	0.511	0.740	12.473	0.511	0.740	12.473	0.511	0.741	12.473	0.511	0.741	12.473	0.511	0.741	12.472	0.510	0.739
DarkIR	8.782	0.094	0.807	8.782	0.094	0.807	8.782	0.094	0.807	8.782	0.094	0.807	8.782	0.094	0.808	8.782	0.094	0.808	8.782	0.094	0.809	8.782	0.094	0.808
AGTRL	18.161	0.667	0.593	18.151	0.665	0.593	18.153	0.666	0.594	18.160	0.666	0.593	18.157	0.666	0.593	18.157	0.667	0.593	18.157	0.664	0.594	18.139	0.663	0.593

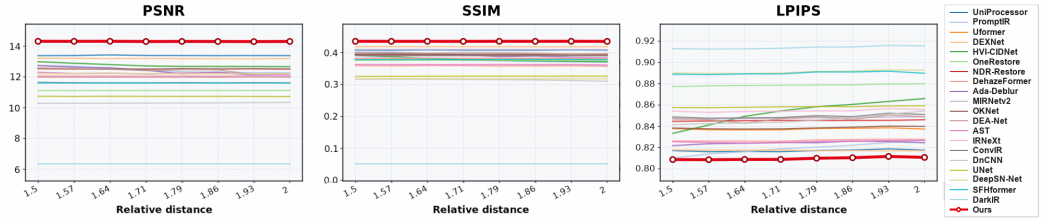

Supplementary Figure 18: Three-metric robustness comparison for distance degradation on PhotoDisplay. The labeled distance setting spans 1.5 to 2.0 m. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 9: Quantitative results for distance degradation on PhotoDisplay. Distance is reported in m; PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	1.5			1.57			1.64			1.71			1.79			1.86			1.93			2		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	13.380	0.407	0.817	13.390	0.407	0.816	13.428	0.408	0.816	13.393	0.407	0.816	13.391	0.407	0.817	13.383	0.407	0.818	13.388	0.408	0.819	13.384	0.407	0.818
OneRestore	11.122	0.376	0.877	11.122	0.376	0.878	11.123	0.376	0.878	11.123	0.376	0.879	11.123	0.376	0.879	11.123	0.376	0.879	11.124	0.376	0.880	11.124	0.375	0.880
DeepSN-Net	12.207	0.318	0.890	12.211	0.318	0.890	12.219	0.318	0.890	12.216	0.318	0.890	12.207	0.318	0.892	12.208	0.318	0.892	12.207	0.318	0.893	12.215	0.318	0.893
PromptIR	12.749	0.390	0.810	12.704	0.383	0.814	12.551	0.378	0.816	12.440	0.376	0.819	12.344	0.374	0.820	12.312	0.372	0.822	12.278	0.371	0.825	12.279	0.370	0.827
Uformer	12.056	0.381	0.838	12.043	0.381	0.837	12.050	0.381	0.837	12.037	0.380	0.837	12.054	0.381	0.838	12.054	0.381	0.838	12.037	0.381	0.839	12.026	0.380	0.837
MIRNetv2	12.034	0.409	0.846	12.034	0.409	0.847	12.033	0.409	0.847	12.033	0.409	0.848	12.033	0.409	0.848	12.032	0.409	0.848	12.032	0.409	0.849	12.030	0.409	0.849
OKNet	12.554	0.395	0.838	12.523	0.395	0.838	12.510	0.394	0.837	12.508	0.395	0.838	12.542	0.395	0.838	12.526	0.395	0.839	12.536	0.395	0.840	12.524	0.394	0.840
AST	12.303	0.363	0.825	12.220	0.362	0.825	12.248	0.362	0.824	12.185	0.362	0.825	12.260	0.362	0.826	12.252	0.362	0.826	12.209	0.362	0.827	12.109	0.361	0.827
ConvIR	11.650	0.390	0.849	11.622	0.390	0.848	11.616	0.390	0.848	11.615	0.390	0.848	11.631	0.390	0.850	11.614	0.390	0.849	11.630	0.390	0.852	11.614	0.390	0.851
DnCNN	10.291	0.316	0.842	10.290	0.316	0.843	10.296	0.316	0.843	10.300	0.316	0.844	10.305	0.315	0.846	10.318	0.314	0.848	10.331	0.313	0.851	10.349	0.310	0.854
U-Net	10.746	0.325	0.858	10.734	0.325	0.858	10.744	0.325	0.858	10.738	0.326	0.858	10.735	0.326	0.859	10.739	0.326	0.858	10.732	0.326	0.859	10.725	0.326	0.859
DExNet	13.212	0.419	0.818	13.202	0.418	0.818	13.195	0.418	0.818	13.199	0.418	0.818	13.193	0.418	0.817	13.178	0.418	0.817	13.180	0.418	0.817	13.199	0.418	0.817
HVI-CIDNet	12.988	0.393	0.833	12.889	0.387	0.841	12.803	0.382	0.849	12.714	0.377	0.854	12.687	0.375	0.858	12.685	0.374	0.861	12.681	0.372	0.863	12.668	0.371	0.866
NDR-Restore	12.006	0.392	0.845	12.006	0.392	0.845	12.005	0.392	0.845	12.007	0.392	0.846	12.005	0.392	0.845	12.001	0.392	0.845	12.014	0.392	0.846	12.008	0.392	0.846
DehazeFormer	12.513	0.359	0.826	12.502	0.358	0.826	12.506	0.359	0.826	12.504	0.359	0.826	12.505	0.358	0.827	12.501	0.358	0.828	12.513	0.359	0.828	12.480	0.357	0.828
Ada-Deblur	12.737	0.400	0.822	12.633	0.398	0.824	12.592	0.397	0.824	12.439	0.393	0.825	12.202	0.387	0.825	12.306	0.389	0.826	12.096	0.385	0.826	12.015	0.384	0.825
DEA-Net	12.587	0.399	0.848	12.540	0.400	0.846	12.498	0.398	0.843	12.410	0.399	0.847	12.475	0.399	0.849	12.461	0.398	0.847	12.149	0.395	0.851	12.213	0.396	0.849
IRNeXt	12.067	0.383	0.854	12.063	0.383	0.853	12.063	0.383	0.854	12.062	0.383	0.854	12.065	0.383	0.854	12.064	0.383	0.855	12.066	0.383	0.857	12.064	0.383	0.856
SFHformer	11.606	0.378	0.889	11.606	0.378	0.889	11.605	0.378	0.889	11.606	0.378	0.889	11.605	0.378	0.891	11.604	0.378	0.891	11.605	0.378	0.892	11.606	0.378	0.890
DarkIR	6.362	0.052	0.913	6.362	0.052	0.913	6.362	0.052	0.913	6.362	0.052	0.913	6.362	0.052	0.915	6.362	0.052	0.915	6.362	0.052	0.916	6.362	0.052	0.916
AGTRL	14.311	0.435	0.809	14.311	0.435	0.809	14.314	0.435	0.809	14.300	0.435	0.809	14.304	0.435	0.810	14.301	0.435	0.810	14.294	0.435	0.812	14.304	0.435	0.811

Natural-image recovery follows the same trend across the distance sweep. AGTRL leads all three metrics, and SSIM remains 0.435 at the displayed precision. The wall-to-screen representation continues to preserve the spatial organization of this content.

A.3.11 Distance Degradation on ScreenGame

We compare fixed models over the labeled distance settings from 1.5 to 2.0 m on ScreenGame.

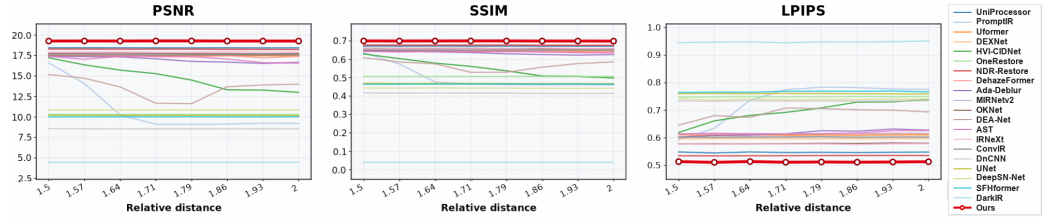

Supplementary Figure 19: Three-metric robustness comparison for distance degradation on ScreenGame. The labeled distance setting spans 1.5 to 2.0 m. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 10: Quantitative results for distance degradation on ScreenGame. Distance is reported in m; PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	1.5			1.57			1.64			1.71			1.79			1.86			1.93			2		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	18.454	0.672	0.549	18.454	0.671	0.545	18.451	0.672	0.549	18.446	0.671	0.547	18.446	0.672	0.548	18.451	0.671	0.546	18.441	0.671	0.548	18.449	0.671	0.548
OneRestore	10.314	0.507	0.748	10.314	0.507	0.749	10.313	0.507	0.750	10.313	0.507	0.749	10.315	0.507	0.750	10.312	0.507	0.750	10.313	0.507	0.750	10.310	0.507	0.750
DeepSN-Net	10.849	0.444	0.742	10.853	0.444	0.738	10.850	0.445	0.741	10.859	0.444	0.738	10.860	0.444	0.739	10.866	0.443	0.739	10.869	0.442	0.740	10.860	0.443	0.741
PromptIR	16.614	0.636	0.592	14.081	0.575	0.635	10.232	0.477	0.737	9.120	0.466	0.775	9.101	0.463	0.784	9.134	0.461	0.783	9.245	0.462	0.777	9.212	0.461	0.776
Uformer	17.497	0.645	0.612	17.483	0.644	0.608	17.477	0.644	0.612	17.457	0.643	0.609	17.468	0.643	0.610	17.472	0.643	0.610	17.478	0.643	0.611	17.449	0.642	0.611
MIRNetv2	17.665	0.656	0.615	17.646	0.656	0.616	17.646	0.656	0.616	17.655	0.656	0.615	17.653	0.656	0.615	17.638	0.655	0.616	17.646	0.656	0.616	17.645	0.656	0.616
OKNet	17.798	0.659	0.578	17.802	0.659	0.578	17.815	0.659	0.578	17.795	0.659	0.579	17.745	0.658	0.580	17.760	0.658	0.579	17.702	0.657	0.582	17.747	0.658	0.580
AST	17.407	0.643	0.614	17.047	0.640	0.617	17.402	0.643	0.614	17.404	0.643	0.613	17.376	0.643	0.615	17.078	0.640	0.618	16.640	0.636	0.626	16.568	0.635	0.626
ConvIR	17.595	0.650	0.603	17.586	0.650	0.602	17.592	0.650	0.603	17.576	0.650	0.603	17.578	0.650	0.604	17.587	0.650	0.602	17.577	0.650	0.603	17.580	0.649	0.603
DnCNN	8.583	0.418	0.734	8.552	0.417	0.733	8.547	0.417	0.735	8.543	0.417	0.734	8.543	0.416	0.734	8.538	0.415	0.733	8.541	0.415	0.734	8.539	0.415	0.735
U-Net	10.227	0.470	0.760	10.216	0.469	0.761	10.202	0.469	0.760	10.212	0.469	0.761	10.198	0.469	0.761	10.185	0.469	0.760	10.178	0.468	0.760	10.182	0.468	0.759
DEANet	17.602	0.658	0.596	17.597	0.658	0.597	17.595	0.658	0.597	17.597	0.658	0.597	17.599	0.658	0.597	17.596	0.658	0.597	17.603	0.658	0.597	17.598	0.658	0.597
HVI-CIDNet	17.232	0.629	0.620	16.369	0.604	0.662	15.727	0.579	0.682	15.293	0.562	0.693	14.508	0.538	0.709	13.321	0.509	0.730	13.287	0.507	0.731	13.011	0.499	0.739
NDR-Restore	18.300	0.678	0.535	18.299	0.678	0.534	18.278	0.677	0.535	18.291	0.678	0.535	18.290	0.678	0.536	18.264	0.677	0.536	18.254	0.677	0.536	18.246	0.676	0.536
DehazeFormer	17.604	0.645	0.611	17.565	0.644	0.612	17.525	0.644	0.613	17.365	0.641	0.614	17.357	0.641	0.614	17.519	0.644	0.611	17.218	0.638	0.616	17.400	0.642	0.615
Ada-Deblur	17.408	0.648	0.602	17.381	0.643	0.609	17.345	0.641	0.610	17.106	0.637	0.615	16.790	0.630	0.626	16.705	0.626	0.624	16.544	0.623	0.632	16.683	0.626	0.628
DEA-Net	15.191	0.609	0.646	14.735	0.586	0.681	13.672	0.575	0.674	11.678	0.530	0.709	11.628	0.530	0.706	13.691	0.558	0.702	13.915	0.577	0.701	14.004	0.586	0.694
IRNeXt	17.915	0.658	0.579	17.923	0.658	0.577	17.942	0.660	0.578	17.899	0.657	0.578	17.912	0.658	0.578	18.001	0.662	0.575	17.918	0.658	0.578	17.931	0.659	0.579
SFHformer	10.019	0.464	0.765	10.019	0.465	0.766	10.021	0.465	0.766	10.020	0.466	0.769	10.022	0.466	0.770	10.014	0.465	0.769	10.015	0.466	0.770	10.018	0.465	0.767
DarkIR	4.457	0.040	0.946	4.457	0.040	0.948	4.457	0.040	0.949	4.457	0.040	0.945	4.457	0.040	0.949	4.457	0.040	0.948	4.457	0.040	0.949	4.457	0.040	0.952
AGTRL	19.279	0.700	0.514	19.285	0.700	0.511	19.282	0.700	0.514	19.294	0.700	0.511	19.293	0.700	0.512	19.278	0.699	0.511	19.278	0.699	0.512	19.274	0.699	0.513

The distance sweep also preserves AGTRL's advantage on game content. Its SSIM values remain closely grouped as the labeled condition changes, extending the stability analysis to a third content family.

A.3.12 Ambient-Light Conditions

Supplementary Figure 20 shows the display-self-occluded setup under different ambient-light conditions. Changes in room illumination alter the background brightness and color against which the screen-induced wall pattern is observed. These photographs document the physical lighting conditions relevant to acquisition. The following brightness-reduction experiments quantify recovery under their specified wall-image brightness adjustments, with the results reported in Supplementary Figures 21–23.



Supplementary Figure 20: Real acquisition photographs under different ambient-light conditions. The four panels show the display, tripod-mounted camera, and illuminated relay wall as the room's background illumination changes. Screen-induced wall light is visible together with the ambient illumination. These views document lighting variability in the physical setup; the subsequent brightness-reduction plots and tables specify the numerical input adjustments used for quantitative evaluation.

A.3.13 BRIGHTNESS REDUCTION ON OSLOCK

We evaluate OSLock recovery across the labeled wall-brightness reduction settings.

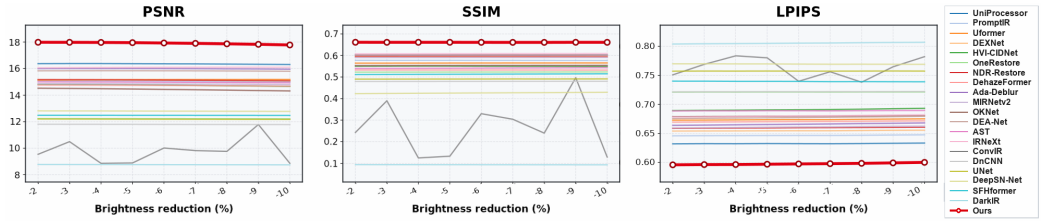

Supplementary Figure 21: Three-metric robustness comparison for brightness reduction on OSLock. The plotted brightness-adjustment setting ranges from -2% to -10%. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 11: Quantitative results for brightness reduction on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	-2			-3			-4			-5			-6			-7			-8			-9			-10		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	16.372	0.602	0.632	16.378	0.602	0.632	16.378	0.602	0.632	16.369	0.602	0.633	16.362	0.602	0.632	16.356	0.602	0.632	16.340	0.602	0.633	16.325	0.602	0.633	16.306	0.602	0.633
OneRestore	12.235	0.523	0.721	12.234	0.523	0.721	12.232	0.523	0.721	12.230	0.523	0.721	12.229	0.523	0.721	12.227	0.523	0.721	12.225	0.523	0.721	12.223	0.523	0.721	12.222	0.524	0.721
DeepSN-Net	12.814	0.423	0.770	12.811	0.424	0.770	12.807	0.425	0.770	12.802	0.425	0.770	12.797	0.426	0.769	12.792	0.427	0.769	12.786	0.428	0.769	12.781	0.429	0.769	12.775	0.429	0.769
PromptIR	15.086	0.577	0.646	15.082	0.577	0.646	15.076	0.577	0.646	15.068	0.577	0.646	15.055	0.577	0.647	15.042	0.577	0.647	15.028	0.577	0.647	15.012	0.577	0.647	14.995	0.577	0.648
Uformer	15.133	0.564	0.673	15.143	0.564	0.673	15.155	0.565	0.673	15.168	0.565	0.674	15.176	0.565	0.674	15.179	0.565	0.674	15.181	0.565	0.674	15.184	0.565	0.675	15.183	0.565	0.675
MIRNetv2	14.860	0.591	0.659	14.858	0.591	0.660	14.856	0.591	0.660	14.851	0.591	0.661	14.843	0.591	0.661	14.833	0.591	0.661	14.822	0.591	0.662	14.808	0.591	0.662	14.792	0.591	0.663
OKNet	14.523	0.553	0.679	14.504	0.553	0.679	14.484	0.553	0.679	14.461	0.553	0.680	14.434	0.553	0.680	14.406	0.554	0.680	14.378	0.554	0.680	14.347	0.554	0.680	14.317	0.554	0.681
AST	14.997	0.540	0.689	14.992	0.540	0.689	14.987	0.541	0.689	14.981	0.541	0.689	14.972	0.541	0.689	14.963	0.542	0.689	14.953	0.542	0.689	14.941	0.543	0.689	14.929	0.543	0.689
ConvIR	9.542	0.243	0.751	10.494	0.390	0.769	8.861	0.125	0.784	8.890	0.133	0.780	10.016	0.330	0.740	9.822	0.304	0.756	9.761	0.240	0.738	11.786	0.497	0.765	8.866	0.129	0.782
DnCNN	11.808	0.480	0.722	11.804	0.480	0.722	11.802	0.480	0.722	11.800	0.480	0.722	11.796	0.480	0.722	11.793	0.480	0.722	11.789	0.480	0.722	11.786	0.480	0.722	11.783	0.480	0.722
U-Net	12.202	0.490	0.757	12.197	0.490	0.757	12.193	0.491	0.757	12.189	0.491	0.757	12.185	0.491	0.757	12.180	0.491	0.757	12.176	0.491	0.757	12.173	0.491	0.758	12.169	0.491	0.757
DExNet	15.159	0.597	0.654	15.157	0.597	0.654	15.153	0.597	0.654	15.146	0.597	0.655	15.137	0.597	0.655	15.128	0.597	0.655	15.117	0.597	0.655	15.105	0.597	0.655	15.092	0.597	0.656
HVI-CIDNet	14.809	0.550	0.689	14.807	0.551	0.690	14.799	0.551	0.690	14.783	0.551	0.691	14.755	0.550	0.691	14.729	0.550	0.691	14.708	0.550	0.692	14.683	0.550	0.693	14.660	0.549	0.693
NDR-Restore	15.172	0.594	0.659	15.167	0.594	0.659	15.160	0.594	0.659	15.150	0.594	0.659	15.137	0.594	0.660	15.125	0.594	0.660	15.111	0.593	0.660	15.096	0.593	0.660	15.079	0.593	0.661
DehazeFormer	14.760	0.533	0.669	14.757	0.533	0.670	14.753	0.533	0.670	14.745	0.533	0.670	14.732	0.533	0.670	14.721	0.533	0.671	14.707	0.533	0.671	14.691	0.532	0.671	14.674	0.532	0.672
Ada-Deblur	16.003	0.602	0.664	16.009	0.602	0.664	16.006	0.602	0.665	16.008	0.602	0.665	16.005	0.602	0.666	15.995	0.602	0.666	15.982	0.601	0.667	15.974	0.601	0.668	15.957	0.601	0.668
DEA-Net	15.833	0.603	0.676	15.837	0.603	0.676	15.840	0.603	0.677	15.839	0.603	0.677	15.835	0.603	0.677	15.831	0.603	0.678	15.823	0.603	0.678	15.813	0.603	0.679	15.799	0.602	0.680
IRNeXt	16.062	0.608	0.678	16.075	0.608	0.679	16.086	0.608	0.679	16.093	0.608	0.680	16.099	0.608	0.680	16.103	0.609	0.680	16.106	0.609	0.681	16.105	0.609	0.682	16.102	0.609	0.682
SFHformer	12.482	0.511	0.740	12.481	0.512	0.740	12.480	0.512	0.740	12.479	0.513	0.740	12.478	0.513	0.740	12.477	0.514	0.739	12.476	0.514	0.739	12.474	0.515	0.739	12.473	0.515	0.739
DarkIR	8.774	0.093	0.804	8.771	0.093	0.804	8.768	0.093	0.805	8.765	0.093	0.805	8.761	0.093	0.806	8.758	0.093	0.806	8.755	0.093	0.806	8.751	0.093	0.807	8.748	0.093	0.807
AGTRL	**17.983**	**0.661**	**0.596**	**17.980**	**0.661**	**0.596**	**17.971**	**0.661**	**0.596**	**17.955**	**0.661**	**0.597**	**17.929**	**0.661**	**0.597**	**17.902**	**0.661**	**0.598**	**17.866**	**0.661**	**0.598**	**17.829**	**0.661**	**0.599**	**17.781**	**0.661**	**0.600**

AGTRL preserves its reconstruction advantage on interface content as wall-image brightness decreases. The method leads the comparison throughout the sweep, with SSIM remaining 0.661 at the displayed precision. The stable structural agreement indicates that the recovered screen organization persists as the wall observation becomes darker.

A.3.14 BRIGHTNESS REDUCTION ON PHOTODISPLAY

We evaluate PhotoDisplay recovery across the labeled wall-brightness reduction settings.

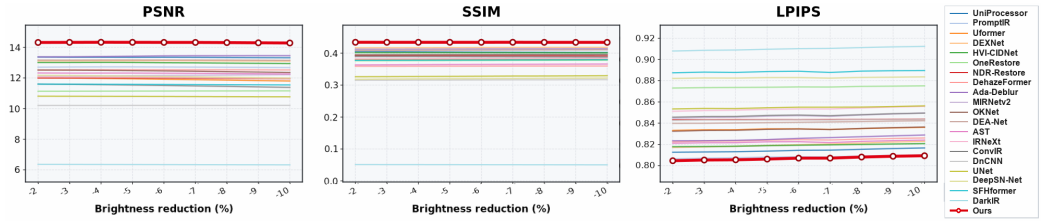

Supplementary Figure 22: Three-metric robustness comparison for brightness reduction on PhotoDisplay. The plotted brightness-adjustment setting ranges from -2% to -10%. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

The same recovery advantage extends from interface organization to natural-image composition. AGTRL maintains the highest SSIM throughout the brightness sweep, with a displayed value of 0.434. Broad image structure therefore remains recoverable from the reduced-brightness wall observations.

A.3.15 BRIGHTNESS REDUCTION ON SCREENGAME

We evaluate ScreenGame recovery across the labeled wall-brightness reduction settings.

ScreenGame brings this analysis to content with dark regions and weak wall contrast. AGTRL leads all three metrics at every brightness level and retains 0.664 SSIM at the -10% setting. This

Supplementary Table 12: Quantitative results for brightness reduction on PhotoDisplay. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	-2			-3			-4			-5			-6			-7			-8			-9			-10		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	13.367	0.405	0.813	13.356	0.405	0.813	13.353	0.404	0.813	13.348	0.404	0.814	13.351	0.403	0.814	13.346	0.403	0.815	13.337	0.403	0.815	13.326	0.402	0.816	13.317	0.402	0.817
OneRestore	11.129	0.376	0.873	11.132	0.376	0.873	11.137	0.377	0.874	11.140	0.377	0.874	11.144	0.377	0.874	11.147	0.377	0.874	11.150	0.377	0.875	11.154	0.378	0.875	11.156	0.378	0.875
DeepSN-Net	12.168	0.318	0.882	12.155	0.319	0.883	12.137	0.320	0.882	12.115	0.320	0.883	12.090	0.320	0.883	12.067	0.321	0.882	12.043	0.321	0.883	12.016	0.322	0.883	11.991	0.322	0.884
PromptIR	12.703	0.394	**0.806**	12.715	0.394	**0.807**	12.720	0.394	**0.807**	12.721	0.395	**0.808**	12.713	0.395	**0.808**	12.707	0.395	**0.808**	12.699	0.395	**0.809**	12.690	0.395	**0.810**	12.679	0.395	**0.810**
Uformer	12.012	0.380	0.833	11.992	0.381	0.834	11.969	0.381	0.834	11.944	0.381	0.835	11.916	0.381	0.835	11.888	0.381	0.834	11.860	0.381	0.835	11.829	0.381	0.836	11.799	0.381	0.836
MIRNetv2	12.092	0.411	0.845	12.108	0.411	0.845	12.124	0.411	0.846	12.141	0.411	0.847	12.159	0.411	0.847	12.175	0.411	0.846	12.188	0.411	0.848	12.202	0.412	0.849	12.212	0.412	0.850
OKNet	12.517	0.395	0.832	12.509	0.395	0.833	12.497	0.395	0.833	12.480	0.395	0.834	12.458	0.395	0.835	12.436	0.395	0.834	12.411	0.396	0.835	12.383	0.396	0.836	12.355	0.396	0.836
AST	12.323	0.363	0.821	12.315	0.364	0.822	12.315	0.364	0.822	12.316	0.364	0.822	12.307	0.365	0.823	12.298	0.365	0.822	12.285	0.365	0.823	12.273	0.366	0.824	12.259	0.366	0.824
ConvIR	11.599	0.390	0.846	11.580	0.389	0.846	11.559	0.389	0.846	11.534	0.389	0.847	11.506	0.389	0.848	11.477	0.389	0.847	11.447	0.389	0.848	11.414	0.388	0.849	11.382	0.388	0.849
DnCNN	10.202	0.315	0.843	10.204	0.315	0.843	10.204	0.316	0.843	10.205	0.316	0.843	10.206	0.316	0.843	10.206	0.317	0.843	10.206	0.317	0.843	10.208	0.317	0.843	10.208	0.317	0.844
U-Net	10.806	0.326	0.853	10.801	0.327	0.854	10.796	0.327	0.854	10.792	0.327	0.855	10.787	0.328	0.855	10.781	0.328	0.855	10.776	0.329	0.855	10.771	0.329	0.856	10.767	0.329	0.856
DEXNet	13.114	**0.417**	0.817	13.115	**0.417**	0.818	13.119	**0.417**	0.818	13.123	**0.418**	0.819	13.126	**0.418**	0.820	13.121	**0.418**	0.820	13.113	**0.418**	0.821	13.105	**0.418**	0.822	13.096	**0.418**	0.823
HVI-CIDNet	13.006	0.402	0.818	13.007	0.402	0.818	13.014	0.402	0.818	13.008	0.402	0.819	12.994	0.402	0.819	12.984	0.401	0.820	12.967	0.401	0.820	12.953	0.401	0.820	12.937	0.401	0.821
NDR-Restore	11.994	0.392	0.843	11.999	0.392	0.843	12.000	0.392	0.843	11.997	0.392	0.843	11.991	0.392	0.843	11.985	0.392	0.843	11.976	0.392	0.844	11.967	0.391	0.844	11.957	0.391	0.844
DehazeFormer	12.518	0.359	0.822	12.526	0.359	0.823	12.532	0.359	0.823	12.535	0.359	0.824	12.536	0.359	0.824	12.535	0.360	0.824	12.533	0.360	0.825	12.529	0.360	0.826	12.523	0.360	0.826
Ada-Deblur	**13.377**	0.412	0.823	**13.380**	0.412	0.824	**13.399**	0.412	0.824	**13.419**	0.413	0.825	**13.443**	0.413	0.826	**13.451**	0.413	0.826	**13.450**	0.413	0.827	**13.455**	0.414	0.828	**13.455**	0.414	0.829
DEA-Net	13.168	0.408	0.840	13.165	0.408	0.840	13.174	0.409	0.840	13.174	0.409	0.840	13.172	0.409	0.841	13.170	0.410	0.841	13.166	0.410	0.842	13.157	0.410	0.842	13.149	0.411	0.842
IRNeXt	12.062	0.383	0.851	12.060	0.383	0.852	12.055	0.383	0.852	12.047	0.383	0.853	12.036	0.383	0.853	12.024	0.384	0.853	12.011	0.384	0.854	11.994	0.384	0.855	11.978	0.384	0.856
SFHformer	11.596	0.378	0.887	11.591	0.378	0.888	11.585	0.378	0.888	11.580	0.378	0.889	11.573	0.379	0.889	11.566	0.379	0.888	11.559	0.379	0.889	11.551	0.380	0.889	11.543	0.380	0.890
DarkIR	6.353	0.051	0.908	6.350	0.051	0.909	6.347	0.051	0.909	6.343	0.051	0.910	6.339	0.051	0.910	6.335	0.051	0.910	6.331	0.051	0.911	6.328	0.050	0.912	6.324	0.050	0.912
AGTRL	**14.316**	**0.434**	**0.805**	**14.323**	**0.434**	**0.805**	**14.329**	**0.434**	**0.806**	**14.326**	**0.434**	**0.806**	**14.321**	**0.434**	**0.807**	**14.319**	**0.434**	**0.807**	**14.309**	**0.434**	**0.808**	**14.295**	**0.434**	**0.809**	**14.282**	**0.434**	**0.809**

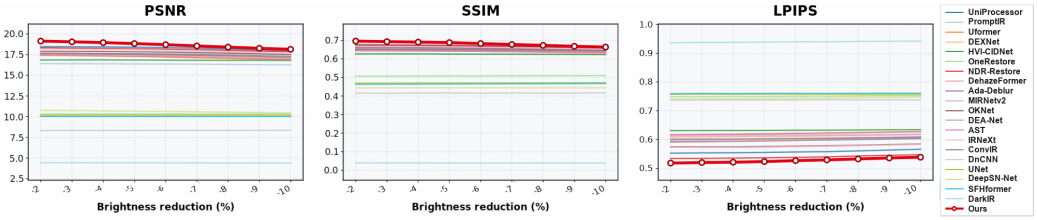

Supplementary Figure 23: Three-metric robustness comparison for brightness reduction on ScreenGame. The plotted brightness-adjustment setting ranges from -2% to -10%. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 13: Quantitative results for brightness reduction on ScreenGame. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	-2			-3			-4			-5			-6			-7			-8			-9			-10		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**18.466**	0.673	0.553	**18.451**	0.672	0.554	**18.436**	0.672	0.554	**18.396**	0.671	0.555	**18.349**	0.669	0.557	**18.287**	0.667	0.558	**18.186**	**0.665**	0.560	**18.074**	0.661	0.563	**17.960**	0.658	0.566
OneRestore	10.318	0.507	0.748	10.320	0.507	0.749	10.322	0.508	0.749	10.325	0.508	0.749	10.327	0.509	0.751	10.330	0.509	0.751	10.333	0.509	0.752	10.336	0.510	0.752	10.339	0.510	0.753
DeepSN-Net	10.769	0.444	0.740	10.738	0.444	0.741	10.705	0.444	0.741	10.670	0.444	0.742	10.633	0.444	0.743	10.595	0.444	0.744	10.555	0.444	0.745	10.510	0.444	0.746	10.470	0.444	0.748
PromptIR	16.414	0.629	0.600	16.402	0.629	0.601	16.393	0.629	0.601	16.372	0.628	0.602	16.364	0.628	0.602	16.351	0.627	0.603	16.321	0.626	0.604	16.293	0.625	0.604	16.262	0.625	0.605
Uformer	17.448	0.645	0.616	17.409	0.644	0.617	17.344	0.643	0.618	17.292	0.642	0.619	17.203	0.641	0.621	17.121	0.639	0.623	17.044	0.638	0.625	16.948	0.636	0.626	16.858	0.634	0.628
MIRNetv2	17.619	0.655	0.614	17.619	0.655	0.615	17.621	0.654	0.615	17.612	0.654	0.616	17.608	0.653	0.616	17.595	0.653	0.616	17.577	0.652	0.617	17.555	0.651	0.618	17.533	0.650	0.619
OKNet	17.899	0.662	0.575	17.889	0.662	0.575	17.883	0.661	0.575	17.843	0.659	0.576	17.803	0.658	0.577	17.749	0.656	0.578	17.684	0.653	0.580	17.597	0.650	0.582	17.508	0.647	0.584
AST	17.397	0.644	0.617	17.375	0.643	0.618	17.330	0.642	0.619	17.290	0.642	0.620	17.249	0.641	0.621	17.184	0.640	0.622	17.116	0.638	0.624	17.021	0.637	0.626	16.928	0.635	0.627
ConvIR	17.633	0.652	0.601	17.624	0.652	0.602	17.606	0.652	0.602	17.584	0.651	0.602	17.537	0.650	0.604	17.485	0.649	0.605	17.430	0.648	0.606	17.362	0.647	0.608	17.299	0.645	0.609
DnCNN	8.325	0.416	0.737	8.336	0.415	0.738	8.334	0.416	0.738	8.340	0.417	0.737	8.342	0.415	0.738	8.341	0.417	0.738	8.346	0.416	0.738	8.353	0.416	0.738	8.363	0.417	0.738
U-Net	10.190	0.470	0.757	10.195	0.469	0.758	10.194	0.470	0.757	10.195	0.470	0.758	10.193	0.471	0.758	10.189	0.471	0.758	10.186	0.471	0.757	10.185	0.472	0.757	10.184	0.472	0.758
DEXNet	17.474	0.655	0.600	17.451	0.654	0.601	17.430	0.654	0.601	17.403	0.653	0.601	17.367	0.652	0.602	17.330	0.650	0.602	17.282	0.649	0.603	17.201	0.647	0.604	17.139	0.646	0.605
HVI-CIDNet	16.853	0.626	0.631	16.867	0.626	0.631	16.865	0.626	0.632	16.850	0.625	0.632	16.832	0.625	0.632	16.820	0.625	0.633	16.816	0.624	0.633	16.811	0.624	0.634	16.809	0.623	0.634
NDR-Restore	18.323	**0.677**	**0.534**	18.277	**0.675**	**0.535**	18.257	**0.675**	**0.536**	18.203	**0.672**	**0.537**	18.160	**0.671**	**0.538**	18.090	**0.668**	**0.540**	18.010	**0.665**	**0.543**	17.922	**0.662**	**0.545**	17.852	**0.661**	**0.547**
DehazeFormer	17.711	0.647	0.609	17.687	0.646	0.610	17.642	0.645	0.610	17.589	0.643	0.611	17.522	0.641	0.612	17.438	0.638	0.613	17.337	0.636	0.614	17.212	0.633	0.616	17.078	0.630	0.617
Ada-Deblur	17.600	0.652	0.594	17.576	0.651	0.594	17.550	0.650	0.595	17.512	0.649	0.596	17.460	0.647	0.597	17.407	0.646	0.598	17.344	0.644	0.600	17.250	0.642	0.602	17.149	0.639	0.603
DEA-Net	17.614	0.655	0.590	17.594	0.654	0.592	17.563	0.653	0.593	17.527	0.651	0.595	17.469	0.650	0.597	17.394	0.647	0.600	17.330	0.645	0.603	17.252	0.643	0.605	17.158	0.641	0.608
IRNeXt	18.159	0.668	0.573	18.148	0.667	0.573	18.126	0.666	0.573	18.090	0.665	0.573	18.028	0.663	0.575	17.948	0.660	0.577	17.860	0.658	0.579	17.749	0.655	0.581	17.622	0.651	0.583
SFHformer	10.033	0.463	0.760	10.033	0.464	0.760	10.033	0.464	0.760	10.032	0.465	0.760	10.031	0.465	0.761	10.030	0.466	0.760	10.028	0.466	0.761	10.027	0.467	0.761	10.025	0.468	0.761
DarkIR	4.445	0.040	0.937	4.439	0.040	0.938	4.434	0.039	0.938	4.428	0.039	0.939	4.421	0.039	0.940	4.416	0.039	0.940	4.410	0.039	0.941	4.404	0.039	0.942	4.398	0.039	0.942
AGTRL	**19.134**	**0.696**	**0.518**	**19.046**	**0.693**	**0.520**	**18.951**	**0.691**	**0.521**	**18.835**	**0.688**	**0.523**	**18.699**	**0.683**	**0.527**	**18.550**	**0.678**	**0.529**	**18.403**	**0.673**	**0.533**	**18.265**	**0.668**	**0.536**	**18.142**	**0.664**	**0.539**

result shows that the wall observation continues to support recovery of useful screen structure as its brightness decreases.

A.3.16 Observation-Angle Robustness: Acquisition Setup

Supplementary Figure 24 documents camera placements for varying the observation angle relative to the display and relay wall. Such changes alter the visible wall region and its projected spatial arrangement. These acquisition photographs document the physical setup independently of the subsequent image-transformation experiments. The perspective, pitch, and yaw tests specify their own transformations and numerical settings below.



Supplementary Figure 24: Real acquisition photographs for observation-angle robustness. The six panels show different camera positions and viewing directions relative to the display and relay wall. The camera remains on the rear side of the display and observes the indirect wall illumination. These setup views document physical angular variation independently of the subsequent image-transformation tests.

A.3.17 Perspective Change on OSLock

We compare OSLock recovery under image transformations specified by $(\text{pitch}, \text{yaw}, \text{roll})$, in degrees. The components denote rotation about the horizontal, vertical, and viewing axes, respectively. They parameterize a rotated image plane projected back to the input coordinates; the transformed wall image retains its original dimensions.

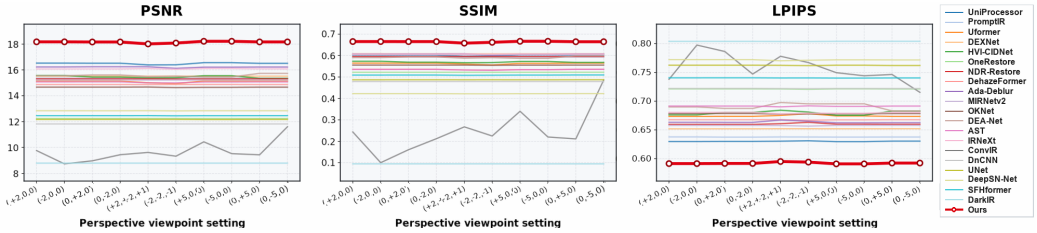

Supplementary Figure 25: Three-metric robustness comparison for perspective change on OSLock. Ten three-component $(\text{pitch}, \text{yaw}, \text{roll})$ settings in degrees are compared, including single-component ± 2 and ± 5 changes and two combined $(\pm 2, \pm 2, \pm 1)$ settings. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 14: Quantitative results for perspective change on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	(+2,0,0)			(-2,0,0)			(0,+2,0)			(0,-2,0)			(+2,+2,+1)			(-2,-2,-1)			(+5,0,0)			(-5,0,0)			(0,+5,0)			(0,-5,0)		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	16.526	0.602	0.630	16.526	0.602	0.630	16.517	0.602	0.630	16.517	0.602	0.630	16.389	0.596	0.630	16.394	0.598	0.631	16.560	0.603	0.630	16.560	0.603	0.630	16.509	0.602	0.631	16.509	0.602	0.631
OneRestore	12.238	0.522	0.721	12.238	0.522	0.721	12.238	0.522	0.721	12.238	0.522	0.721	12.238	0.522	0.721	12.239	0.522	0.721	12.237	0.522	0.721	12.237	0.522	0.721	12.237	0.522	0.721	12.237	0.522	0.721
DeepSN-Net	12.840	0.423	0.772	12.840	0.423	0.772	12.840	0.423	0.772	12.840	0.423	0.772	12.846	0.422	0.770	12.828	0.422	0.771	12.843	0.422	0.772	12.843	0.422	0.772	12.841	0.423	0.772	12.841	0.423	0.772
PromptIR	15.354	0.593	0.638	15.354	0.593	0.638	15.344	0.593	0.637	15.344	0.593	0.637	15.350	0.593	0.638	15.350	0.593	0.638	15.352	0.593	0.638	15.352	0.593	0.638	15.344	0.593	0.638	15.344	0.593	0.638
Uformer	15.141	0.565	0.674	15.141	0.565	0.674	15.136	0.565	0.674	15.136	0.565	0.674	15.006	0.560	0.675	14.937	0.556	0.678	15.147	0.564	0.674	15.147	0.564	0.674	15.140	0.565	0.673	15.140	0.565	0.673
MIRNetv2	15.216	0.604	0.658	15.216	0.604	0.658	15.216	0.604	0.658	15.216	0.604	0.658	15.213	0.604	0.658	15.213	0.604	0.657	15.212	0.604	0.658	15.212	0.604	0.658	15.212	0.604	0.658	15.212	0.604	0.658
OKNet	14.666	0.556	0.679	14.666	0.556	0.679	14.669	0.556	0.679	14.669	0.556	0.679	14.675	0.556	0.677	14.631	0.555	0.678	14.653	0.555	0.678	14.653	0.555	0.678	14.666	0.556	0.678	14.666	0.556	0.678
AST	15.073	0.535	0.691	15.073	0.535	0.691	15.074	0.535	0.691	15.074	0.535	0.691	15.000	0.533	0.690	14.995	0.531	0.692	15.063	0.534	0.691	15.063	0.534	0.691	15.072	0.535	0.691	15.072	0.535	0.691
ConvIR	9.758	0.244	0.738	8.730	0.101	0.797	8.968	0.161	0.786	9.437	0.212	0.747	9.616	0.268	0.778	9.324	0.225	0.767	10.434	0.340	0.749	9.521	0.220	0.744	9.430	0.211	0.746	11.603	0.486	0.715
DnCNN	11.819	0.479	0.722	11.819	0.479	0.722	11.820	0.479	0.722	11.820	0.479	0.722	11.819	0.479	0.722	11.818	0.479	0.722	11.818	0.479	0.722	11.819	0.479	0.722	11.818	0.479	0.722	11.818	0.479	0.722
U-Net	12.166	0.487	0.762	12.166	0.487	0.762	12.170	0.487	0.763	12.170	0.487	0.763	12.171	0.487	0.762	12.172	0.487	0.762	12.159	0.487	0.762	12.159	0.487	0.762	12.168	0.487	0.762	12.168	0.487	0.762
DEANet	15.520	0.608	0.652	15.520	0.608	0.652	15.520	0.608	0.652	15.520	0.608	0.652	15.516	0.608	0.652	15.517	0.608	0.652	15.518	0.608	0.652	15.518	0.608	0.652	15.516	0.608	0.652	15.516	0.608	0.652
HVI-CIDNet	15.564	0.573	0.676	15.564	0.573	0.676	15.448	0.569	0.680	15.448	0.569	0.680	15.389	0.566	0.684	15.403	0.567	0.681	15.550	0.572	0.675	15.550	0.572	0.675	15.351	0.568	0.683	15.351	0.568	0.683
NDR-Restore	15.320	0.596	0.660	15.320	0.596	0.660	15.320	0.596	0.660	15.320	0.596	0.660	15.318	0.596	0.661	15.302	0.596	0.663	15.310	0.596	0.660	15.310	0.596	0.660	15.316	0.596	0.660	15.316	0.596	0.660
DehazeFormer	14.868	0.536	0.668	14.868	0.536	0.668	14.867	0.536	0.668	14.867	0.536	0.668	14.864	0.533	0.668	14.874	0.533	0.667	14.873	0.534	0.668	14.873	0.534	0.668	14.867	0.536	0.668	14.867	0.536	0.668
Ada-Deblur	16.230	0.607	0.663	16.230	0.607	0.663	16.245	0.608	0.663	16.245	0.608	0.663	16.245	0.608	0.667	16.124	0.602	0.665	16.214	0.606	0.662	16.214	0.606	0.662	16.224	0.608	0.663	16.224	0.608	0.663
DEA-Net	15.548	0.590	0.690	15.548	0.590	0.690	15.610	0.594	0.687	15.610	0.594	0.687	15.486	0.588	0.698	15.518	0.590	0.695	15.453	0.587	0.695	15.453	0.587	0.695	15.734	0.599	0.683	15.734	0.599	0.683
IRNeXt	16.076	0.607	0.681	16.076	0.607	0.681	16.078	0.607	0.681	16.078	0.607	0.681	16.113	0.607	0.678	16.046	0.606	0.679	16.110	0.607	0.679	16.110	0.607	0.679	16.082	0.607	0.681	16.082	0.607	0.681
SFHformer	12.460	0.509	0.741	12.460	0.509	0.741	12.460	0.509	0.741	12.460	0.509	0.741	12.465	0.508	0.740	12.448	0.508	0.740	12.459	0.509	0.740	12.459	0.509	0.740	12.461	0.510	0.740	12.461	0.510	0.740
DarkIR	8.789	0.095	0.804	8.789	0.095	0.804	8.789	0.095	0.804	8.789	0.095	0.804	8.788	0.095	0.804	8.788	0.095	0.804	8.788	0.095	0.804	8.788	0.095	0.804	8.788	0.095	0.804	8.788	0.095	0.804
AGTRL	18.173	0.665	0.591	18.173	0.665	0.591	18.160	0.665	0.592	18.160	0.665	0.592	18.011	0.658	0.595	18.079	0.662	0.594	18.214	0.667	0.591	18.214	0.667	0.591	18.165	0.664	0.592	18.165	0.664	0.592

The perspective sweep moves the analysis from signal strength to spatial correspondence. AGTRL retains the leading reconstruction scores across the plotted settings, with SSIM between 0.658 and 0.667.

A.3.18 Perspective Change on PhotoDisplay

We compare PhotoDisplay recovery under the labeled three-component perspective settings.

AGTRL maintains the leading results when natural-image wall patterns undergo the plotted perspective changes. The advantage therefore extends to changes in the spatial arrangement of observed evidence.

A.3.19 Perspective Change on ScreenGame

We compare ScreenGame recovery under the labeled three-component perspective settings.

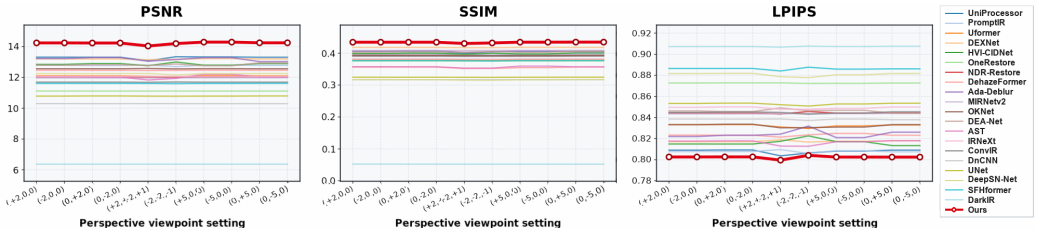

Supplementary Figure 26: Three-metric robustness comparison for perspective change on PhotoDisplay. Ten three-component $(\text{pitch}, \text{yaw}, \text{roll})$ settings in degrees are compared, including single-component ± 2 and ± 5 changes and two combined $(\pm 2, \pm 2, \pm 1)$ settings. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 15: Quantitative results for perspective change on PhotoDisplay. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	(+2,0,0)			(-2,0,0)			(0,+2,0)			(0,-2,0)			(+2,+2,+1)			(-2,-2,-1)			(+5,0,0)			(-5,0,0)			(0,+5,0)			(0,-5,0)		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	13.312	0.406	0.809	13.312	0.406	0.809	13.308	0.406	0.809	13.308	0.406	0.809	13.077	0.399	0.804	13.341	0.405	0.806	13.327	0.405	0.808	13.327	0.405	0.808	13.306	0.406	0.809	13.306	0.406	0.809
OneRestore	11.114	0.376	0.873	11.114	0.376	0.873	11.114	0.376	0.873	11.114	0.376	0.873	11.113	0.375	0.873	11.115	0.375	0.872	11.114	0.375	0.872	11.114	0.375	0.872	11.114	0.376	0.873	11.114	0.376	0.873
DeepSN-Net	12.250	0.317	0.882	12.250	0.317	0.882	12.251	0.317	0.882	12.251	0.317	0.882	12.254	0.315	0.879	12.158	0.315	0.878	12.243	0.317	0.880	12.243	0.317	0.880	12.254	0.317	0.882	12.254	0.317	0.882
PromptIR	12.783	0.397	0.808	12.783	0.397	0.808	12.772	0.398	0.808	12.772	0.398	0.808	12.805	0.397	0.810	12.699	0.394	0.806	12.787	0.396	0.809	12.787	0.396	0.809	12.776	0.398	0.808	12.776	0.398	0.808
Uformer	12.105	0.380	0.833	12.105	0.380	0.833	12.101	0.380	0.833	12.101	0.380	0.833	12.069	0.379	0.831	12.027	0.378	0.830	12.121	0.380	0.832	12.121	0.380	0.832	12.100	0.380	0.833	12.100	0.380	0.833
MIRNetv2	11.960	0.408	0.843	11.960	0.408	0.843	11.960	0.408	0.843	11.960	0.408	0.843	11.964	0.408	0.844	11.965	0.408	0.844	11.974	0.409	0.844	11.974	0.409	0.844	11.974	0.409	0.844	11.974	0.409	0.844
OKNet	12.561	0.394	0.833	12.561	0.394	0.833	12.560	0.394	0.834	12.560	0.394	0.834	12.579	0.393	0.830	12.552	0.393	0.830	12.565	0.393	0.831	12.565	0.393	0.831	12.559	0.394	0.833	12.559	0.394	0.833
AST	12.052	0.358	0.818	12.052	0.358	0.818	12.024	0.357	0.818	12.024	0.357	0.818	11.830	0.353	0.813	11.937	0.353	0.813	12.206	0.359	0.817	12.206	0.359	0.817	12.023	0.357	0.818	12.023	0.357	0.818
ConvIR	11.692	0.390	0.845	11.692	0.390	0.845	11.692	0.390	0.845	11.692	0.390	0.845	11.682	0.390	0.845	11.703	0.390	0.843	11.690	0.390	0.844	11.690	0.390	0.844	11.689	0.390	0.845	11.689	0.390	0.845
DnCNN	10.298	0.317	0.839	10.298	0.317	0.839	10.295	0.317	0.838	10.295	0.317	0.838	10.298	0.317	0.839	10.296	0.317	0.837	10.295	0.317	0.839	10.295	0.317	0.839	10.294	0.317	0.838	10.294	0.317	0.838
U-Net	10.787	0.325	0.853	10.787	0.325	0.853	10.795	0.325	0.854	10.795	0.325	0.854	10.783	0.324	0.852	10.780	0.324	0.851	10.787	0.325	0.853	10.787	0.325	0.853	10.795	0.325	0.854	10.795	0.325	0.854
DEsNet	13.199	0.418	0.817	13.199	0.418	0.817	13.193	0.418	0.817	13.193	0.418	0.817	13.147	0.417	0.819	13.176	0.417	0.817	13.218	0.418	0.818	13.218	0.418	0.818	13.194	0.418	0.818	13.194	0.418	0.818
HVI-CIDNet	12.845	0.399	0.815	12.845	0.399	0.815	12.899	0.400	0.815	12.899	0.400	0.815	12.758	0.397	0.817	13.001	0.400	0.823	12.792	0.398	0.817	12.792	0.398	0.817	12.895	0.401	0.813	12.895	0.401	0.813
NDR-Restore	12.034	0.392	0.844	12.034	0.392	0.844	12.033	0.392	0.844	12.033	0.392	0.844	12.006	0.392	0.843	11.992	0.391	0.846	12.039	0.392	0.844	12.039	0.392	0.844	12.037	0.392	0.844	12.037	0.392	0.844
DehazeFormer	12.461	0.356	0.823	12.461	0.356	0.823	12.461	0.356	0.823	12.461	0.356	0.823	12.448	0.352	0.821	12.456	0.352	0.824	12.470	0.354	0.825	12.470	0.354	0.825	12.462	0.357	0.823	12.462	0.357	0.823
Ada-Deblur	13.220	0.406	0.822	13.220	0.406	0.822	13.286	0.408	0.823	13.286	0.408	0.823	13.051	0.402	0.824	13.158	0.404	0.832	13.279	0.407	0.821	13.279	0.407	0.821	13.024	0.405	0.826	13.024	0.405	0.826
DEA-Net	12.788	0.403	0.846	12.788	0.403	0.846	12.824	0.404	0.846	12.824	0.404	0.846	12.767	0.403	0.849	12.871	0.405	0.847	12.754	0.403	0.847	12.754	0.403	0.847	12.945	0.406	0.844	12.945	0.406	0.844
IRNeXt	12.067	0.382	0.850	12.067	0.382	0.850	12.067	0.382	0.850	12.067	0.382	0.850	12.068	0.382	0.848	12.070	0.381	0.848	12.065	0.382	0.849	12.065	0.382	0.849	12.066	0.382	0.850	12.066	0.382	0.850
SFHformer	11.608	0.377	0.886	11.608	0.377	0.886	11.608	0.377	0.886	11.608	0.377	0.886	11.604	0.376	0.884	11.589	0.376	0.888	11.606	0.377	0.886	11.606	0.377	0.886	11.610	0.377	0.886	11.610	0.377	0.886
DarkIR	6.370	0.052	0.907	6.370	0.052	0.907	6.370	0.052	0.907	6.370	0.052	0.907	6.370	0.052	0.907	6.370	0.052	0.908	6.369	0.052	0.907	6.369	0.052	0.907	6.369	0.052	0.908	6.369	0.052	0.908
AGTRL	14.238	0.434	0.803	14.238	0.434	0.803	14.230	0.434	0.803	14.230	0.434	0.803	14.036	0.430	0.800	14.194	0.432	0.804	14.289	0.434	0.803	14.289	0.434	0.803	14.244	0.435	0.802	14.244	0.435	0.802

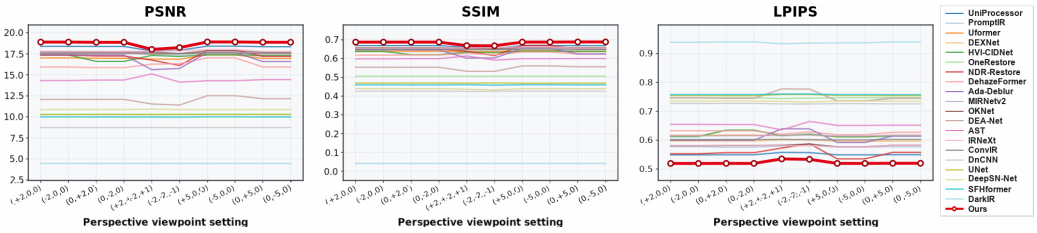

Supplementary Figure 27: Three-metric robustness comparison for perspective change on ScreenGame. Ten three-component $(\text{pitch}, \text{yaw}, \text{roll})$ settings in degrees are compared, including single-component ± 2 and ± 5 changes and two combined $(\pm 2, \pm 2, \pm 1)$ settings. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 16: Quantitative results for perspective change on ScreenGame. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	(+2,0,0)			(-2,0,0)			(0,+2,0)			(0,-2,0)			(+2,+2,+1)			(-2,-2,-1)			(+5,0,0)			(-5,0,0)			(0,+5,0)			(0,-5,0)		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	18.387	0.669	0.549	18.387	0.669	0.549	18.374	0.669	0.549	18.374	0.669	0.549	17.838	0.661	0.558	17.921	0.654	0.557	18.415	0.670	0.549	18.415	0.670	0.549	18.348	0.669	0.550	18.348	0.669	0.550
OneRestore	10.315	0.506	0.747	10.315	0.506	0.747	10.315	0.506	0.747	10.315	0.506	0.747	10.311	0.506	0.744	10.310	0.506	0.745	10.313	0.506	0.747	10.313	0.506	0.747	10.315	0.506	0.747	10.315	0.506	0.747
DeepSN-Net	10.879	0.440	0.736	10.879	0.440	0.736	10.879	0.440	0.735	10.879	0.440	0.735	10.926	0.436	0.734	10.870	0.433	0.732	10.867	0.440	0.736	10.867	0.440	0.736	10.874	0.440	0.735	10.874	0.440	0.735
PromptIR	17.494	0.650	0.578	17.494	0.650	0.578	17.612	0.654	0.575	17.612	0.654	0.575	17.480	0.648	0.580	17.362	0.646	0.583	17.528	0.652	0.576	17.528	0.652	0.576	17.549	0.653	0.577	17.549	0.653	0.577
Uformer	17.014	0.639	0.616	17.014	0.639	0.616	16.968	0.638	0.617	16.968	0.638	0.617	16.880	0.638	0.616	16.854	0.637	0.618	17.382	0.642	0.612	17.382	0.642	0.612	16.950	0.638	0.617	16.950	0.638	0.617
MIRNetv2	17.634	0.653	0.615	17.634	0.653	0.615	17.629	0.653	0.615	17.629	0.653	0.615	17.552	0.649	0.617	17.522	0.648	0.618	17.637	0.655	0.614	17.637	0.655	0.614	17.605	0.652	0.617	17.605	0.652	0.617
OKNet	17.755	0.658	0.581	17.755	0.658	0.581	17.739	0.657	0.581	17.739	0.657	0.581	17.707	0.656	0.582	17.534	0.652	0.587	17.860	0.661	0.577	17.860	0.661	0.577	17.706	0.657	0.582	17.706	0.657	0.582
AST	14.322	0.598	0.655	14.322	0.598	0.655	14.382	0.600	0.654	14.382	0.600	0.654	15.122	0.614	0.636	14.135	0.592	0.665	14.300	0.600	0.651	14.300	0.600	0.651	14.434	0.600	0.652	14.434	0.600	0.652
ConvIR	17.568	0.649	0.602	17.568	0.649	0.602	17.555	0.649	0.602	17.555	0.649	0.602	17.479	0.649	0.603	17.543	0.648	0.602	17.608	0.650	0.601	17.608	0.650	0.601	17.525	0.648	0.603	17.525	0.648	0.603
DnCNN	8.735	0.426	0.727	8.735	0.426	0.727	8.733	0.426	0.727	8.733	0.426	0.727	8.745	0.425	0.726	8.719	0.425	0.725	8.741	0.427	0.726	8.741	0.427	0.726	8.735	0.426	0.726	8.735	0.426	0.726
U-Net	10.280	0.470	0.755	10.280	0.470	0.755	10.283	0.470	0.754	10.283	0.470	0.754	10.273	0.470	0.758	10.279	0.470	0.758	10.288	0.469	0.755	10.288	0.469	0.755	10.286	0.469	0.754	10.286	0.469	0.754
DEsNet	17.572	0.656	0.597	17.572	0.656	0.597	17.571	0.656	0.597	17.571	0.656	0.597	17.477	0.652	0.601	17.507	0.653	0.601	17.585	0.657	0.596	17.585	0.657	0.596	17.573	0.657	0.596	17.573	0.657	0.596
HVI-CIDNet	17.360	0.638	0.612	17.360	0.638	0.612	16.602	0.621	0.635	16.602	0.621	0.635	17.315	0.635	0.615	17.188	0.631	0.621	17.379	0.640	0.611	17.379	0.640	0.611	17.306	0.638	0.615	17.306	0.638	0.615
NDR-Restore	17.345	0.653	0.553	17.345	0.653	0.553	17.244	0.650	0.557	17.244	0.650	0.557	16.751	0.639	0.573	16.116	0.615	0.588	17.920	0.670	0.535	17.920	0.670	0.535	17.196	0.649	0.558	17.196	0.649	0.558
DehazeFormer	15.935	0.618	0.633	15.935	0.618	0.633	15.864	0.619	0.632	15.864	0.619	0.632	16.272	0.627	0.620	16.323	0.620	0.629	17.023	0.635	0.619	17.023	0.635	0.619	15.940	0.624	0.627	15.940	0.624	0.627
Ada-Deblur	17.391	0.647	0.600	17.391	0.647	0.600	17.363	0.646	0.600	17.363	0.646	0.600	15.620	0.602	0.639	15.761	0.603	0.639	17.666	0.654	0.592	17.666	0.654	0.592	16.589	0.625	0.614	16.589	0.625	0.614
DEA-Net	12.082	0.554	0.747	12.082	0.554	0.747	12.088	0.554	0.745	12.088	0.554	0.745	11.535	0.532	0.778	11.406	0.532	0.777	12.525	0.561	0.736	12.525	0.561	0.736	12.165	0.556	0.745	12.165	0.556	0.745
IRNeXt	17.802	0.655	0.582	17.802	0.655	0.582	17.791	0.655	0.582	17.791	0.655	0.582	17.810	0.657	0.580	17.873	0.658	0.580	18.014	0.662	0.576	18.014	0.662	0.576	17.775	0.654	0.582	17.775	0.654	0.582
SFHformer	10.000	0.460	0.758	10.000	0.460	0.758	9.998	0.460	0.758	9.998	0.460	0.758	9.988	0.461	0.760	9.978	0.458	0.760	10.010	0.460	0.758	10.010	0.460	0.758	9.998	0.460	0.757	9.998	0.460	0.757
DarkIR	4.467	0.041	0.939	4.467	0.041	0.939	4.467	0.041	0.940	4.467	0.041	0.940	4.466	0.041	0.934	4.466	0.041	0.936	4.465	0.041	0.937	4.465	0.041	0.937	4.464	0.041	0.940	4.464	0.041	0.940
AGTRL	18.898	0.688	0.519	18.898	0.688	0.519	18.875	0.688	0.519	18.875	0.688	0.519	18.029	0.669	0.535	18.241	0.668	0.533	18.914	0.688	0.519	18.914	0.688	0.519	18.878	0.689	0.520	18.878	0.689	0.520

Combined perspective settings produce a broader range of reconstruction scores on ScreenGame. AGTRL remains the leading method, extending its recovery advantage to altered arrangements of game-screen evidence.

A.3.20 Radial Lens Distortion on OSLock

We evaluate OSLock recovery with the labeled radial-distortion coefficients.

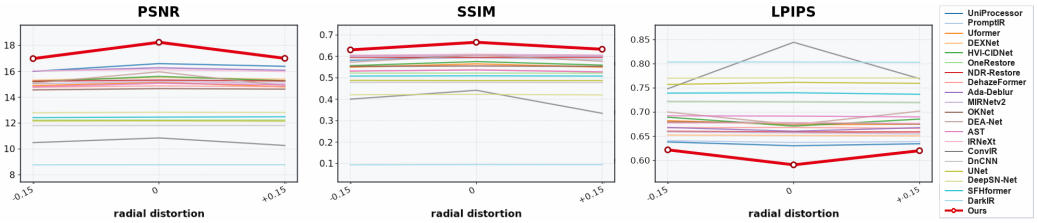

Supplementary Figure 28: Three-metric robustness comparison for radial lens distortion on OSLock. The radial-distortion coefficient is -0.15, 0, or $+0.15$. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 17: Quantitative results for radial lens distortion on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	-0.15			0			+0.15		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	15.997	0.581	**0.639**	**16.595**	0.606	**0.631**	**16.385**	0.600	**0.635**
OneRestore	12.233	0.522	0.722	12.237	0.522	0.721	12.239	0.522	0.720
DeepSN-Net	12.796	0.421	0.770	12.836	0.423	0.771	12.823	0.420	0.771
PromptIR	15.229	0.585	0.641	15.342	0.594	0.637	15.248	0.585	0.639
Uformer	14.892	0.550	0.683	15.149	0.565	0.674	14.796	0.550	0.676
MIRNetv2	15.037	0.599	0.659	15.211	0.604	0.658	15.078	0.600	0.655
OKNet	14.566	0.553	0.680	14.664	0.555	0.678	14.623	0.553	0.676
AST	14.793	0.532	0.693	15.072	0.537	0.692	14.955	0.528	0.690
ConvIR	10.493	0.401	0.749	10.854	0.442	0.845	10.271	0.334	0.769
DnCNN	11.807	0.479	0.723	11.818	0.479	0.722	11.818	0.478	0.721
U-Net	12.169	0.488	0.757	12.180	0.487	0.762	12.158	0.487	0.760
DExNet	15.318	0.603	0.653	15.518	0.608	0.652	15.397	**0.605**	0.651
HVI-CIDNet	15.210	0.557	0.690	15.606	0.576	0.671	15.272	0.559	0.686
NDR-Restore	15.238	0.595	0.661	15.319	0.595	0.659	15.254	0.595	0.659
DehazeFormer	14.769	0.531	0.668	14.873	0.537	0.669	14.805	0.527	0.667
Ada-Deblur	16.005	0.604	0.669	16.273	**0.609**	0.661	16.077	**0.605**	0.668
DEA-Net	15.087	0.573	0.701	15.962	0.607	0.674	14.886	0.577	0.702
IRNeXt	**16.036**	**0.606**	0.677	16.134	0.608	0.678	15.979	**0.605**	0.677
SFHformer	12.425	0.508	0.740	12.461	0.509	0.740	12.488	0.508	0.737
DarkIR	8.778	0.094	0.804	8.787	0.095	0.804	8.786	0.095	0.803
AGTRL	**16.982**	**0.630**	**0.622**	**18.242**	**0.666**	**0.591**	**17.000**	**0.633**	**0.621**

Radial distortion introduces nonlinear warping into the wall observation. AGTRL retains the best PSNR, SSIM, and LPIPS at radial-distortion coefficients of -0.15 and $+0.15$, extending the geometric recovery analysis beyond the perspective sweep.

A.3.21 Vertical Camera-Viewpoint Shift on OSLock

We compare OSLock recovery across the plotted pitch settings from 0 to +15 degrees.

The pitch sweep examines directional change along the image height. AGTRL maintains the leading reconstruction quality and retains 17.347 dB PSNR at $+15°$ pitch. This extends the geometric analysis to a sustained vertical change in the observation.

A.3.22 Horizontal Camera-Viewpoint Shift on OSLock

We compare OSLock recovery across the plotted yaw settings from -15 to +15 degrees.

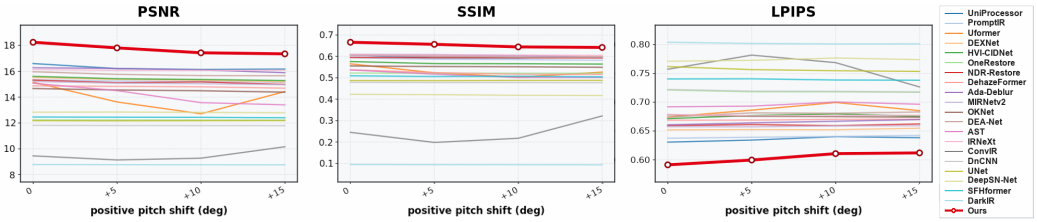

Supplementary Figure 29: Three-metric robustness comparison for pitch-shift setting on OSLock. The plotted positive pitch settings are $0°$, $+5°$, $+10°$, and $+15°$. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 18: Quantitative results for vertical camera-viewpoint shift on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	0			+5			+10			+15		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**16.595**	0.606	**0.631**	16.208	0.593	**0.634**	**16.140**	0.593	**0.640**	**16.171**	0.599	**0.638**
OneRestore	12.237	0.522	0.721	12.233	0.521	0.718	12.228	0.521	0.718	12.221	0.521	0.717
DeepSN-Net	12.836	0.423	0.771	12.847	0.421	0.773	12.822	0.418	0.777	12.788	0.417	0.774
PromptIR	15.342	0.594	0.637	15.244	0.587	0.639	15.206	0.585	**0.640**	15.135	0.582	0.642
Uformer	15.149	0.565	0.674	13.626	0.523	0.686	12.713	0.503	0.700	14.429	0.527	0.685
MIRNetv2	15.211	0.604	0.658	15.077	0.600	0.657	15.036	0.599	0.659	14.960	0.598	0.659
OKNet	14.664	0.555	0.678	14.552	0.553	0.675	14.494	0.551	0.675	14.384	0.549	0.674
AST	15.072	0.537	0.692	14.503	0.517	0.693	13.574	0.501	0.700	13.404	0.501	0.696
ConvIR	9.458	0.245	0.757	9.143	0.198	0.782	9.280	0.218	0.769	10.163	0.322	0.726
DnCNN	11.818	0.479	0.722	11.806	0.478	0.719	11.798	0.478	0.719	11.789	0.477	0.718
U-Net	12.180	0.487	0.762	12.167	0.487	0.757	12.179	0.488	0.755	12.188	0.488	0.753
DExNet	15.518	0.608	0.652	15.372	0.604	0.652	15.313	0.603	0.652	15.184	0.601	0.655
HVI-CIDNet	15.606	0.576	0.671	15.423	0.566	0.677	15.364	0.566	0.680	15.283	0.564	0.676
NDR-Restore	15.319	0.595	0.659	15.117	0.594	0.661	15.176	0.593	0.659	14.979	0.593	0.662
DehazeFormer	14.873	0.537	0.669	14.831	0.524	0.669	14.798	0.517	0.670	14.693	0.515	0.670
Ada-Deblur	16.273	**0.609**	0.661	**16.216**	**0.608**	0.664	16.102	0.604	0.667	15.888	0.595	0.670
DEA-Net	15.962	0.607	0.674	15.761	0.601	0.681	15.666	0.596	0.681	15.653	0.596	0.683
IRNeXt	16.134	0.608	0.678	16.070	0.607	0.676	16.066	**0.606**	0.677	16.076	**0.605**	0.678
SFHformer	12.461	0.509	0.740	12.448	0.506	0.741	12.438	0.507	0.739	12.407	0.504	0.738
DarkIR	8.787	0.095	0.804	8.778	0.094	0.802	8.772	0.094	0.801	8.765	0.093	0.801
AGTRL	**18.242**	**0.666**	**0.591**	**17.807**	**0.656**	**0.600**	**17.429**	**0.644**	**0.611**	**17.347**	**0.642**	**0.612**

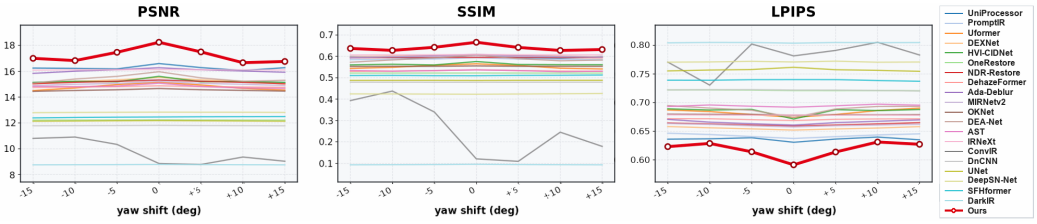

Supplementary Figure 30: Three-metric robustness comparison for yaw-shift setting on OSLock. The plotted yaw settings range from $-15°$ to $+15°$. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

The yaw sweep complements the pitch analysis by varying the labeled observation angle in both horizontal directions. AGTRL leads all three reconstruction metrics across the sweep.

A.3.23 Horizontal Image Translation on OSLock

We compare OSLock recovery across six horizontal shifts of the observed wall image, reported in pixels.

Supplementary Table 19: Quantitative results for horizontal camera-viewpoint shift on OSLock. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	-15			-10			-5			0			+5			+10			+15		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Screen-Oriented Restoration																					
UniProcessor	**16.247**	0.600	**0.636**	**16.223**	0.595	**0.637**	**16.182**	0.593	**0.639**	**16.595**	0.606	**0.631**	**16.283**	0.594	**0.636**	16.060	0.591	**0.640**	**16.276**	0.599	**0.635**
DeepSN-Net	12.792	0.424	0.771	12.814	0.424	0.771	12.825	0.423	0.772	12.836	0.423	0.771	12.830	0.424	0.773	12.828	0.425	0.771	12.816	0.426	0.771
General-Purpose End-to-End Image Restoration																					
PromptIR	15.131	0.585	0.647	15.202	0.586	0.644	15.251	0.588	0.641	15.342	0.594	0.637	15.251	0.588	0.641	15.201	0.586	0.644	15.128	0.584	0.645
Uformer	14.504	0.544	0.687	14.706	0.549	0.684	14.969	0.558	0.679	15.149	0.565	0.674	14.938	0.557	0.679	14.676	0.545	0.686	14.539	0.541	0.690
OneRestore	12.214	0.522	0.722	12.222	0.522	0.722	12.229	0.522	0.722	12.237	0.522	0.721	12.229	0.522	0.721	12.221	0.522	0.721	12.214	0.522	0.721
MIRNetv2	14.895	0.597	0.664	14.993	0.598	0.662	15.062	0.600	0.660	15.211	0.604	0.658	15.062	0.600	0.660	14.993	0.598	0.661	14.895	0.597	0.662
OKNet	14.447	0.553	0.680	14.514	0.553	0.680	14.566	0.554	0.680	14.664	0.555	0.678	14.569	0.554	0.679	14.521	0.554	0.679	14.464	0.554	0.679
AST	14.827	0.534	0.693	14.804	0.531	0.696	14.855	0.534	0.694	15.072	0.537	0.692	14.819	0.534	0.695	14.707	0.528	0.697	14.674	0.531	0.696
IRNeXt	16.085	**0.607**	0.679	16.113	**0.607**	0.678	16.120	**0.608**	0.678	16.134	0.608	0.678	16.118	**0.608**	0.678	**16.107**	**0.607**	0.678	16.100	**0.607**	0.678
ConvIR	10.805	0.393	0.771	10.909	0.438	0.731	10.338	0.340	0.802	8.863	0.120	0.782	8.786	0.109	0.791	9.362	0.245	0.806	9.036	0.179	0.783
U-Net	12.135	0.488	0.755	12.151	0.488	0.757	12.169	0.488	0.758	12.180	0.487	0.762	12.167	0.488	0.758	12.153	0.488	0.756	12.136	0.488	0.755
SFHformer	12.391	0.510	0.739	12.425	0.511	0.739	12.445	0.510	0.740	12.461	0.509	0.740	12.480	0.511	0.740	12.489	0.512	0.739	12.496	0.513	0.737
Models for Related Degradations																					
DExNet	15.031	0.599	0.658	15.213	0.602	0.656	15.338	0.604	0.654	15.518	0.608	0.652	15.337	0.604	0.654	15.206	0.601	0.656	15.017	0.598	0.658
HVI-CIDNet	15.137	0.560	0.689	15.224	0.563	0.687	15.188	0.560	0.688	15.606	0.576	0.671	15.189	0.562	0.688	15.207	0.561	0.686	15.139	0.561	0.688
NDR-Restore	15.035	0.594	0.664	15.142	0.595	0.663	15.222	0.595	0.661	15.319	0.595	0.659	15.221	0.595	0.661	15.140	0.595	0.663	15.032	0.594	0.665
DehazeFormer	14.752	0.530	0.672	14.778	0.531	0.672	14.808	0.533	0.670	14.873	0.537	0.669	14.812	0.534	0.671	14.784	0.531	0.672	14.751	0.530	0.672
Ada-Deblur	15.835	0.593	0.671	16.004	0.599	0.666	16.119	0.603	0.663	16.273	**0.609**	0.661	16.115	0.603	0.665	16.019	0.599	0.667	15.922	0.597	0.670
DEA-Net	15.084	0.573	0.695	15.399	0.583	0.690	15.603	0.593	0.687	15.962	0.607	0.674	15.487	0.589	0.689	15.226	0.579	0.693	15.304	0.583	0.694
DnCNN	11.783	0.479	0.722	11.793	0.479	0.723	11.803	0.479	0.723	11.818	0.479	0.722	11.803	0.479	0.722	11.792	0.479	0.721	11.781	0.479	0.720
DarkIR	8.759	0.093	0.805	8.767	0.093	0.805	8.775	0.094	0.805	8.787	0.095	0.804	8.775	0.094	0.805	8.767	0.093	0.805	8.759	0.093	0.805
AGTRL	**16.999**	**0.637**	**0.623**	**16.826**	**0.628**	**0.629**	**17.466**	**0.643**	**0.614**	**18.242**	**0.666**	**0.591**	**17.497**	**0.642**	**0.614**	**16.663**	**0.628**	**0.631**	**16.757**	**0.632**	**0.627**

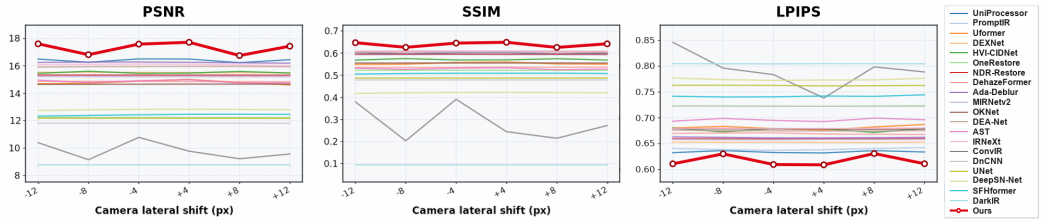

Supplementary Figure 31: Three-metric robustness comparison for horizontal image translation on OSLock. The plotted settings are -12, -8, -4, $+4$, $+8$, and $+12$ pixels. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 20: Quantitative results for horizontal image translation on OSLock (pixels). PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	-12			-8			-4			+4			+8			+12		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Screen-Oriented Restoration																		
UniProcessor	**16.491**	0.602	**0.633**	16.259	0.596	**0.637**	**16.503**	0.602	**0.633**	**16.494**	0.603	**0.632**	**16.231**	0.596	**0.637**	**16.447**	0.601	**0.634**
OneRestore	12.237	0.522	0.723	12.237	0.522	0.723	12.237	0.522	0.722	12.236	0.522	0.722	12.235	0.522	0.723	12.234	0.522	0.724
DeepSN-Net	12.745	0.418	0.777	12.796	0.420	0.774	12.827	0.422	0.773	12.833	0.422	0.773	12.821	0.422	0.774	12.797	0.420	0.777
General-Purpose End-to-End Image Restoration																		
PromptIR	15.336	0.592	0.641	15.339	0.593	0.639	15.341	0.593	0.637	15.340	0.593	0.638	15.336	0.593	0.641	15.332	0.592	0.642
Uformer	14.695	0.549	0.680	14.732	0.550	0.684	14.881	0.558	0.679	15.002	0.560	0.676	14.787	0.551	0.683	14.606	0.550	0.687
MIRNetv2	15.213	0.604	0.658	15.212	0.604	0.658	15.212	0.604	0.658	15.210	0.604	0.658	15.208	0.604	0.658	15.206	0.604	0.658
OKNet	14.645	0.555	0.679	14.653	0.555	0.679	14.659	0.555	0.679	14.667	0.555	0.678	14.669	0.555	0.678	14.668	0.555	0.678
AST	14.959	0.532	0.694	14.778	0.526	0.699	14.913	0.532	0.695	14.993	0.533	0.693	14.745	0.525	0.700	14.757	0.528	0.696
ConvIR	10.402	0.379	0.846	9.159	0.203	0.796	10.786	0.391	0.784	9.780	0.245	0.739	9.228	0.215	0.799	9.577	0.273	0.789
DnCNN	11.818	0.479	0.723	11.818	0.479	0.723	11.818	0.479	0.723	11.817	0.479	0.722	11.816	0.479	0.722	11.814	0.479	0.722
U-Net	12.183	0.487	0.763	12.183	0.487	0.763	12.181	0.487	0.763	12.177	0.487	0.762	12.175	0.487	0.762	12.170	0.487	0.763
Models for Related Degradations																		
DExNet	15.524	**0.608**	0.653	15.525	**0.609**	0.653	15.520	**0.608**	0.653	15.510	**0.608**	0.652	15.505	**0.608**	0.652	15.492	**0.608**	0.653
HVI-CIDNet	15.453	0.569	0.681	15.604	0.575	0.673	15.458	0.569	0.679	15.472	0.569	0.679	15.591	0.575	0.672	15.474	0.569	0.679
NDR-Restore	15.322	0.595	0.660	15.322	0.595	0.660	15.320	0.595	0.660	15.317	0.595	0.659	15.315	0.595	0.660	15.312	0.595	0.660
DehazeFormer	14.867	0.535	0.670	14.871	0.535	0.671	14.868	0.536	0.671	14.857	0.537	0.669	14.842	0.536	0.668	14.840	0.537	0.668
Ada-Deblur	16.246	0.607	0.663	**16.269**	0.608	0.662	16.279	**0.608**	0.661	16.253	**0.608**	0.661	16.220	0.607	0.662	16.190	0.606	0.662
DEA-Net	15.905	0.605	0.676	15.937	0.606	0.675	15.950	0.606	0.675	15.959	0.607	0.675	15.954	0.607	0.675	15.934	0.606	0.675
IRNeXt	16.076	0.607	0.680	16.101	0.607	0.680	16.127	**0.608**	0.679	16.106	0.607	0.679	16.065	0.606	0.680	16.009	0.606	0.681
SFHformer	12.337	0.505	0.742	12.387	0.508	0.740	12.433	0.509	0.741	12.472	0.509	0.742	12.473	0.509	0.742	12.468	0.508	0.745
DarkIR	8.787	0.095	0.805	8.788	0.095	0.805	8.788	0.095	0.804	8.787	0.095	0.805	8.787	0.095	0.805	8.786	0.095	0.805
AGTRL	**17.604**	**0.647**	**0.611**	**16.809**	**0.626**	**0.630**	**17.580**	**0.645**	**0.610**	**17.715**	**0.649**	**0.609**	**16.749**	**0.626**	**0.631**	**17.431**	**0.642**	**0.611**

Horizontal displacement changes where wall evidence appears in the input. AGTRL retains the leading scores across all six settings, showing that its recovery advantage persists when the observed pattern is shifted.

A.3.24 Horizontal Image Translation on PhotoDisplay

We compare PhotoDisplay recovery across the same six horizontal image translations in pixels.

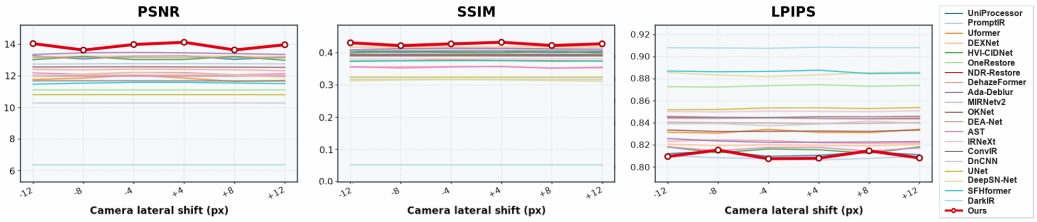

Supplementary Figure 32: Three-metric robustness comparison for horizontal image translation on PhotoDisplay. The plotted settings are -12, -8, -4, $+4$, $+8$, and $+12$ pixels. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

The PhotoDisplay curves extend the horizontal-shift result to natural-image content. AGTRL remains the leading method across the displayed settings, showing that broad screen composition remains recoverable from the displaced wall pattern.

A.3.25 Motion Blur on OSLock with Kernel Length 3 Pixels

On OSLock, we evaluate directional motion blur with kernel length 3 pixels at four labeled angles.

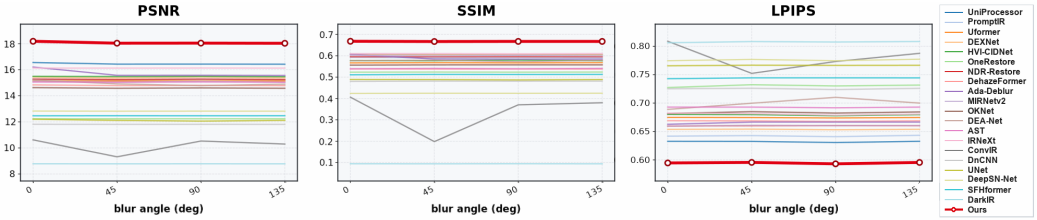

Supplementary Figure 33: Three-metric robustness comparison for motion blur on OSLock with kernel length 3 pixels. At kernel length 3, the labeled blur directions are $0°$, $45°$, $90°$, and $135°$. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 21: Quantitative results for motion blur on OSLock with kernel length 3 pixels. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	0			45			90			135		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**16.555**	0.608	**0.633**	**16.434**	0.605	**0.632**	**16.436**	0.604	**0.630**	**16.428**	0.605	**0.633**
OneRestore	12.237	0.523	0.727	12.238	0.524	0.732	12.239	0.523	0.730	12.238	0.524	0.732
DeepSN-Net	12.831	0.424	0.774	12.822	0.425	0.777	12.826	0.424	0.774	12.821	0.425	0.777
PromptIR	15.269	0.592	0.642	15.219	0.591	0.643	15.271	0.592	0.641	15.219	0.591	0.643
Uformer	15.143	0.567	0.675	15.149	0.568	0.674	15.152	0.567	0.674	15.148	0.568	0.675
MIRNetv2	15.120	0.602	0.659	15.048	0.600	0.661	15.120	0.602	0.660	15.048	0.600	0.660
OKNet	14.629	0.556	0.682	14.581	0.557	0.685	14.607	0.557	0.682	14.580	0.557	0.684
AST	15.058	0.539	0.693	15.047	0.541	0.693	15.062	0.539	0.692	15.047	0.541	0.693
ConvIR	10.617	0.406	0.809	9.318	0.198	0.752	10.529	0.370	0.773	10.297	0.380	0.787
DnCNN	11.817	0.480	0.725	11.812	0.481	0.726	11.814	0.481	0.724	11.812	0.481	0.726
U-Net	12.202	0.489	0.766	12.114	0.489	0.767	12.059	0.487	0.766	12.114	0.489	0.766
DExNet	15.440	0.606	0.654	15.373	0.605	0.654	15.444	0.607	0.653	15.373	0.605	0.654
HVI-CIDNet	15.496	0.577	0.680	15.479	0.577	0.680	15.498	0.580	0.678	15.469	0.577	0.680
NDR-Restore	15.287	0.595	0.659	15.253	0.595	0.660	15.285	0.596	0.660	15.253	0.595	0.660
DehazeFormer	14.840	0.538	0.669	14.785	0.538	0.669	14.801	0.538	0.668	14.784	0.538	0.669
Ada-Deblur	16.210	0.606	0.663	15.573	0.585	0.666	15.577	0.584	0.667	15.562	0.584	0.666
DEA-Net	15.214	0.577	0.689	14.919	0.579	0.700	14.781	0.572	0.710	14.897	0.579	0.700
IRNeXt	16.133	**0.610**	0.682	16.129	**0.610**	0.681	16.135	**0.609**	0.679	16.129	**0.610**	0.681
SFHformer	12.467	0.511	0.743	12.474	0.513	0.744	12.469	0.512	0.744	12.473	0.513	0.744
DarkIR	8.784	0.094	0.806	8.780	0.094	0.808	8.783	0.095	0.807	8.780	0.094	0.808
AGTRL	**18.188**	**0.668**	**0.595**	**18.041**	**0.667**	**0.596**	**18.046**	**0.667**	**0.593**	**18.034**	**0.667**	**0.595**

The length-3 kernel blends nearby wall responses along each tested blur direction. AGTRL retains the best scores at all four angles, showing a recovery advantage under this initial level of directional spreading.

A.3.26 Motion Blur on OSLock with Kernel Length 5 Pixels

On OSLock, we evaluate directional motion blur with kernel length 5 pixels at four labeled angles.

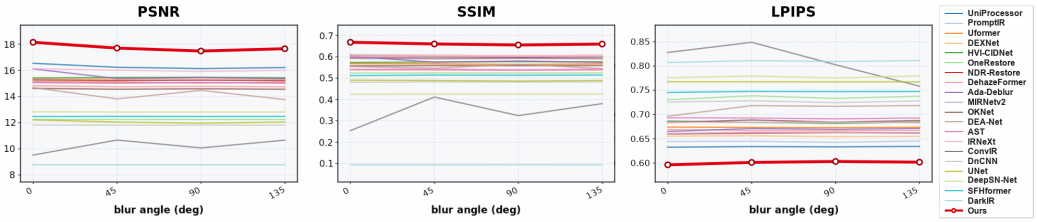

Supplementary Figure 34: Three-metric robustness comparison for motion blur on OSLock with kernel length 5 pixels. At kernel length 5, the labeled blur directions are $0°$, $45°$, $90°$, and $135°$. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 22: Quantitative results for motion blur on OSLock with kernel length 5 pixels. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	0			45			90			135		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**16.534**	0.608	**0.633**	**16.231**	0.602	**0.634**	**16.133**	0.597	**0.633**	**16.209**	0.601	**0.634**
OneRestore	12.238	0.524	0.731	12.240	0.524	0.739	12.241	0.524	0.733	12.240	0.525	0.738
DeepSN-Net	12.830	0.424	0.776	12.815	0.426	0.779	12.816	0.425	0.776	12.814	0.426	0.780
PromptIR	15.252	0.592	0.644	15.193	0.589	0.645	15.247	0.592	0.643	15.192	0.589	0.646
Uformer	15.130	0.568	0.674	15.090	0.567	0.673	15.032	0.564	0.673	15.088	0.567	0.674
MIRNetv2	15.084	0.601	0.661	14.997	0.598	0.664	15.084	0.601	0.663	14.997	0.599	0.663
OKNet	14.617	0.557	0.684	14.555	0.559	0.689	14.586	0.559	0.685	14.554	0.559	0.688
AST	15.054	0.541	0.693	15.044	0.542	0.693	15.057	0.539	0.691	15.043	0.542	0.693
ConvIR	9.527	0.254	0.828	10.669	0.412	0.849	10.072	0.325	0.803	10.662	0.381	0.759
DnCNN	11.817	0.481	0.726	11.812	0.482	0.728	11.814	0.482	0.725	11.811	0.483	0.728
U-Net	12.211	0.491	0.768	12.044	0.489	0.768	11.962	0.486	0.768	12.045	0.490	0.767
DExNet	15.408	0.606	0.656	15.330	0.604	0.656	15.418	**0.606**	0.654	15.329	0.604	0.655
HVI-CIDNet	15.410	0.573	0.687	15.451	0.576	0.684	15.465	0.578	0.682	15.429	0.576	0.684
NDR-Restore	15.274	0.595	0.660	15.225	0.595	0.662	15.253	0.595	0.663	15.225	0.595	0.662
DehazeFormer	14.826	0.539	0.669	14.739	0.538	0.667	14.769	0.537	0.667	14.737	0.538	0.667
Ada-Deblur	16.104	0.603	0.665	15.367	0.574	0.671	15.447	0.581	0.670	15.365	0.574	0.671
DEA-Net	14.676	0.554	0.697	13.814	0.547	0.719	14.466	0.565	0.717	13.765	0.545	0.719
IRNeXt	16.129	**0.611**	0.684	16.008	**0.608**	0.684	15.911	0.605	0.684	15.997	**0.608**	0.683
SFHformer	12.469	0.512	0.746	12.478	0.515	0.748	12.473	0.514	0.747	12.476	0.515	0.748
DarkIR	8.782	0.094	0.807	8.778	0.094	0.811	8.781	0.094	0.809	8.778	0.094	0.811
AGTRL	**18.141**	**0.668**	**0.596**	**17.695**	**0.660**	**0.601**	**17.473**	**0.656**	**0.603**	**17.650**	**0.660**	**0.602**

Increasing the kernel length strengthens directional mixing in the wall image. AGTRL continues to lead at every angle, extending the advantage as the observed evidence becomes more spatially blended.

A.3.27 Motion Blur on OSLock with Kernel Length 9 Pixels

On OSLock, we evaluate directional motion blur with kernel length 9 pixels at four labeled angles.

The length-9 kernel extends the analysis to stronger directional mixing of wall evidence. AGTRL leads all three metrics at every angle and exceeds UniProcessor by 0.734 dB PSNR at $90°$ (Supplementary Table 23). Directional spreading blends neighboring wall responses, intensifying the correspondence problem addressed by AGRL. Read together with the AGRL ablation, this comparison supports the value of spatially organizing mixed evidence for hidden-screen recovery.

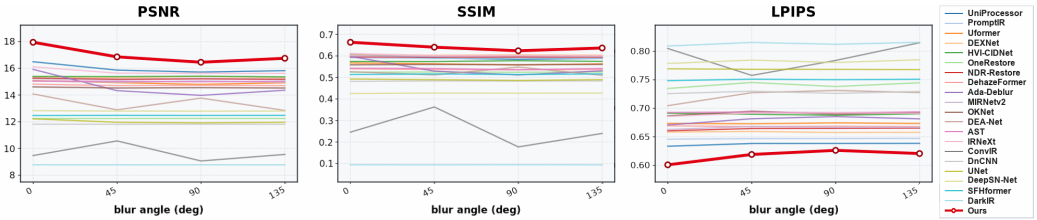

Supplementary Figure 35: Three-metric robustness comparison for motion blur on OSLock with kernel length 9 pixels. At kernel length 9, the labeled blur directions are $0°$, $45°$, $90°$, and $135°$. The three panels report PSNR (dB), SSIM, and LPIPS on the same labeled settings; higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Supplementary Table 23: Quantitative results for motion blur on OSLock with kernel length 9 pixels. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Method	0			45			90			135		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
UniProcessor	**16.491**	0.608	**0.633**	**15.863**	0.591	**0.638**	**15.720**	0.584	**0.638**	**15.824**	0.589	**0.638**
OneRestore	12.239	0.525	0.735	12.244	0.525	0.745	12.246	0.524	0.738	12.244	0.525	0.745
DeepSN-Net	12.828	0.425	0.779	12.801	0.427	0.784	12.795	0.427	0.780	12.800	0.427	0.785
PromptIR	15.229	0.592	0.646	15.154	0.588	0.647	15.226	0.592	0.647	15.153	0.588	0.647
Uformer	15.080	0.568	0.673	14.927	0.563	0.673	14.796	0.558	0.674	14.919	0.563	0.673
MIRNetv2	15.050	0.600	0.663	14.956	0.598	0.668	15.049	0.601	0.666	14.956	0.598	0.667
OKNet	14.608	0.559	0.687	14.526	0.560	0.695	14.551	0.559	0.690	14.524	0.561	0.693
AST	15.049	0.543	0.693	15.030	0.542	0.693	14.997	0.536	0.692	15.028	0.542	0.694
ConvIR	9.470	0.246	0.805	10.565	0.363	0.758	9.067	0.177	0.784	9.554	0.240	0.815
DnCNN	11.818	0.482	0.726	11.813	0.484	0.730	11.816	0.483	0.726	11.813	0.484	0.730
U-Net	12.214	0.493	0.769	11.963	0.490	0.768	11.916	0.486	0.768	11.964	0.490	0.768
DExNet	15.378	0.605	0.658	15.302	**0.603**	0.659	15.393	**0.605**	0.657	15.302	**0.603**	0.658
HVI-CIDNet	15.385	0.574	0.691	15.377	0.574	0.690	15.409	0.578	0.688	15.356	0.575	0.690
NDR-Restore	15.263	0.595	0.660	15.179	0.595	0.665	15.210	0.595	0.665	15.180	0.595	0.665
DehazeFormer	14.810	0.540	0.669	14.705	0.537	0.668	14.730	0.536	0.669	14.703	0.537	0.668
Ada-Deblur	15.918	0.598	0.671	14.327	0.529	0.682	13.973	0.512	0.685	14.359	0.531	0.681
DEA-Net	14.079	0.527	0.705	12.891	0.512	0.727	13.769	0.549	0.731	12.846	0.510	0.728
IRNeXt	16.105	**0.611**	0.686	15.636	0.601	0.692	15.621	0.602	0.692	15.616	0.601	0.691
SFHformer	12.469	0.514	0.748	12.479	0.516	0.751	12.475	0.514	0.750	12.478	0.517	0.751
DarkIR	8.781	0.094	0.809	8.775	0.094	0.815	8.779	0.094	0.812	8.775	0.094	0.816
AGTRL	**17.947**	**0.664**	**0.601**	**16.856**	**0.641**	**0.619**	**16.454**	**0.624**	**0.626**	**16.750**	**0.637**	**0.621**

A.3.28 Fixed U-Net Representation Module

We compare representation modules within a shared U-Net backbone to examine how each module supports the mapping from wall observations to screen content. We compare the proposed module with Identity, SE (Hu et al., 2018), CBAM (Woo et al., 2018), Non-local blocks (Wang et al., 2018), FFC (Chi et al., 2020), and a parameter-matched convolutional module.

SE recalibrates feature channels, while CBAM applies channel and spatial attention sequentially (Hu et al., 2018; Woo et al., 2018). Within the shared U-Net backbone, AGRL and TraceTrust combine screen-oriented feature organization with update weighting guided by local continuity and artifact suspicion, leading all 12 metric columns and achieving 0.627 SSIM on ScreenFigure, compared with 0.559 for SE and 0.547 for CBAM (Supplementary Table 24). The corresponding reconstructions preserve screen regions and major boundaries. Together, these results show the contribution of learning both how wall features map to screen structure and which traces should guide that mapping.

A.3.29 Training Objective

We compare training objectives in the U-Net-based reconstruction setting to examine how supervision shapes hidden-screen recovery. The alternatives include pixelwise, structural, perceptual, and frequency-based objectives.

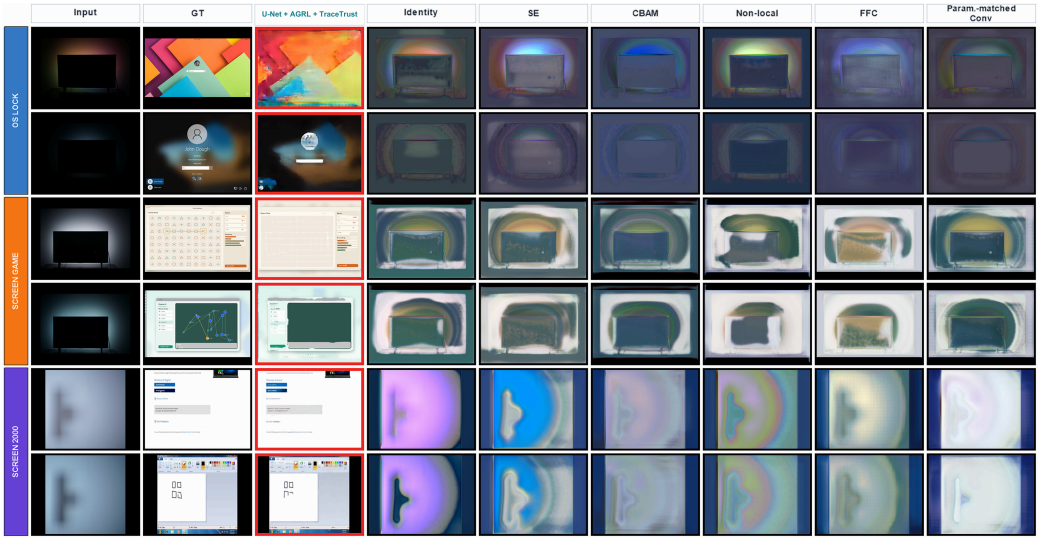

Supplementary Figure 36: Fixed-U-Net representation-module comparison on six samples. Columns show the wall input, ground-truth screen, U-Net with AGRL + TraceTrust, and the labeled alternative modules. The reconstruction columns show the shared U-Net architecture equipped with the indicated representation modules. Column 3 uses AGRL and TraceTrust.

Supplementary Table 24: Quantitative results for fixed U-Net representation module. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	ScreenGame			Screen2000		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Identity	8.275	0.367	0.767	**8.962**	0.268	0.777
SE	9.376	0.411	0.775	7.641	0.268	0.787
CBAM	8.659	0.376	0.765	7.985	0.357	0.826
Non-local	9.565	**0.435**	0.758	8.557	0.346	0.808
FFC	**10.139**	0.409	**0.741**	8.602	**0.385**	0.813
Parameter-matched convolution	9.170	0.379	0.759	8.272	0.324	**0.772**
AGRL + TraceTrust	**11.153**	**0.450**	**0.667**	**9.097**	**0.423**	**0.695**

Representation module	OSLock			ScreenFigure		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Identity	**12.109**	0.448	0.776	11.436	0.551	0.684
SE	11.609	0.429	0.766	10.349	0.559	0.692
CBAM	11.403	0.464	0.752	10.922	0.547	0.714
Non-local	12.054	0.449	**0.746**	11.361	**0.570**	0.695
FFC	12.038	**0.503**	0.759	**11.667**	0.564	**0.678**
Parameter-matched convolution	11.607	0.480	0.761	10.472	0.566	0.695
AGRL + TraceTrust	**13.320**	**0.553**	**0.672**	**11.986**	**0.627**	**0.610**

Focal frequency loss emphasizes spectral components that are difficult to reconstruct (Jiang et al., 2021); the Focal Frequency + L1 objective achieves 8.744 dB PSNR and 0.353 SSIM on Screen2000 in our U-Net objective comparison (Supplementary Table 25). Direct RGB L1 supervision reaches 9.973 dB PSNR and 0.487 SSIM on Screen2000 in this comparison and leads all 12 metric columns across the compared pixelwise, structural, perceptual, and frequency-based objectives. This result supports direct screen-space supervision: the objective compares reconstructed and target RGB values at corresponding screen locations, providing a common learning signal for screen appearance and spatial organization.

Supplementary Table 25: Comparison of training objectives. PSNR is reported in dB. Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Training objective	ScreenGame			Screen2000		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Charbonnier	9.880	0.434	0.756	8.227	0.393	**0.770**
SSIM Loss	9.953	**0.515**	0.771	**9.066**	**0.443**	0.773
Perceptual + L1	9.940	0.406	0.782	8.732	0.267	0.831
Focal Frequency + L1 (Jiang et al., 2021)	**10.550**	0.460	**0.743**	8.744	0.353	0.814
L1 (ours)	**11.606**	**0.566**	**0.669**	**9.973**	**0.487**	**0.693**

Training objective	OSLock			ScreenFigure		
	PSNR (dB)↑	SSIM↑	LPIPS↓	PSNR (dB)↑	SSIM↑	LPIPS↓
Charbonnier	11.772	0.469	0.746	10.488	0.559	0.709
SSIM Loss	12.051	**0.528**	0.736	10.623	**0.617**	0.759
Perceptual + L1	11.995	0.369	0.866	10.554	0.515	0.786
Focal Frequency + L1 (Jiang et al., 2021)	**12.162**	0.490	**0.735**	**11.024**	0.555	**0.701**
L1 (ours)	**13.378**	**0.581**	**0.661**	**11.986**	**0.678**	**0.693**

A.3.30 Weak-Contrast Robustness

We evaluate screen reconstruction as the contrast of the wall observation decreases. The following robustness tests compare representation modules within the shared U-Net backbone; AGRL + TraceTrust denotes the proposed module. These sweeps and their unmasked reference use fixed checkpoints and form a separate evaluation series from the grouped comparison in Supplementary Table 24. Each metric is averaged over test samples within each dataset, then equally across ScreenGame, Screen2000, OSLock, and ScreenFigure. The contrast summary additionally averages over its tested levels.

Supplementary Table 26: Quantitative results for weak-contrast robustness. Higher SSIM and PSNR indicate better reconstruction; PSNR is reported in dB. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	0.25		0.50		0.75		1.00	
	SSIM↑	PSNR (dB)↑	SSIM↑	PSNR (dB)↑	SSIM↑	PSNR (dB)↑	SSIM↑	PSNR (dB)↑
Identity	0.458	10.643	0.441	10.800	0.413	10.137	0.408	10.196
SE	0.489	**11.069**	0.468	10.795	0.433	10.071	0.417	9.744
CBAM	0.493	10.805	0.472	10.711	0.447	9.993	0.436	9.742
Non-local	0.478	10.676	0.473	10.736	0.460	10.492	0.450	10.384
FFC	**0.505**	10.914	**0.488**	**10.959**	**0.465**	**10.503**	**0.465**	**10.611**
Parameter-matched convolution	0.473	9.994	0.461	10.267	0.437	9.914	0.437	9.880
AGRL + TraceTrust	**0.508**	**11.147**	**0.495**	**11.144**	**0.473**	**11.002**	**0.466**	**10.654**

Lower wall contrast makes the spatial variations carrying screen content fainter. The proposed module achieves the best SSIM and PSNR at every contrast level, including the weakest setting of 0.25. Its advantage persists when the wall pattern contains only faint spatial variation.

A.3.31 Spatial Downsampling Robustness

We compare spatial downsampling factors $s = 1, 2, 4, 8$. For $s > 1$, bicubic interpolation with antialiasing reduces the wall observation to $\max(1, \mathrm{round}(H/s)) \times \max(1, \mathrm{round}(W/s))$, then restores it to $H \times W$ using the same interpolation settings. For the 384×512 wall inputs, factor $s = 8$ maps the input to 48×64 before restoring it to 384×512. Both resampling steps use `align_corners=False`; the restored RGB intensities are clipped to $[0, 1]$. Factor $s = 1$ leaves the input unchanged.

Spatial downsampling varies the scale of structure available in the wall observation. The proposed module retains the best LPIPS and PSNR throughout the sweep, with PSNR changing by only 0.050

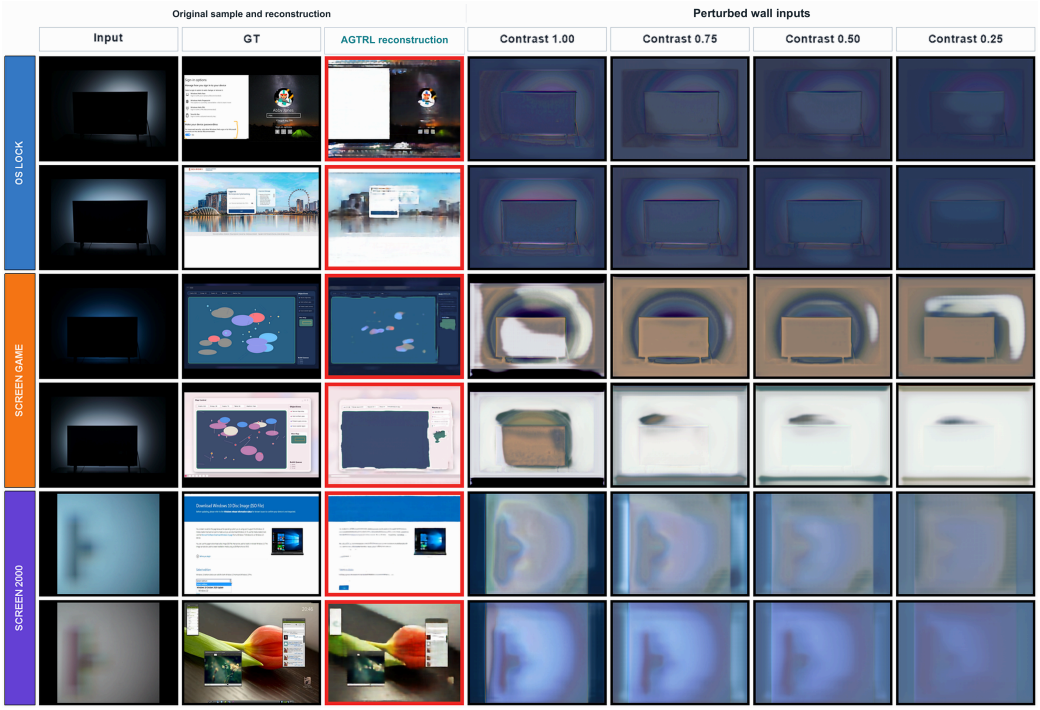

Supplementary Figure 37: Weak-contrast examples from six samples. Columns 1–3 show the original wall observation, ground-truth screen, and reconstruction from the shared U-Net with AGRL + TraceTrust. Columns 4–7 show the wall observations at contrast factors of 1.00, 0.75, 0.50, and 0.25, respectively.

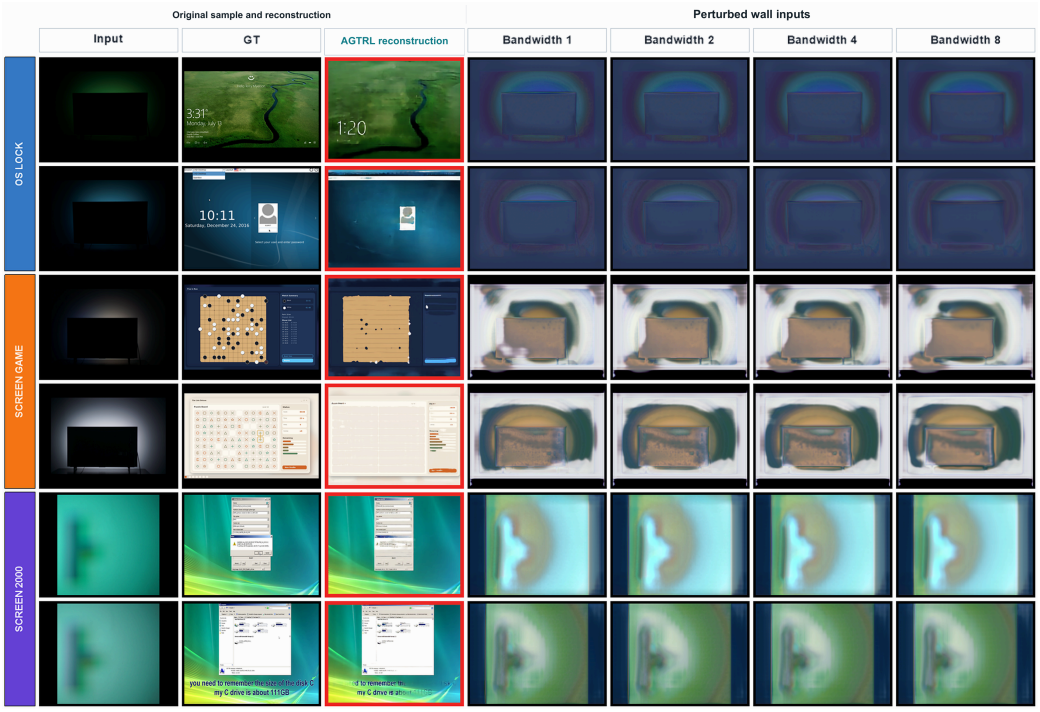

Supplementary Figure 38: Spatial-downsampling examples from six samples. Columns 1–3 show the original wall observation, ground-truth screen, and reconstruction from the shared U-Net with AGRL + TraceTrust. Columns 4–7 show the wall observations at spatial downsampling factors of 1, 2, 4, and 8, respectively.

Supplementary Table 27: Quantitative results for spatial-downsampling robustness. Lower LPIPS and higher PSNR indicate better reconstruction; PSNR is reported in dB. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	1		2		4		8	
	LPIPS↓	PSNR (dB)↑	LPIPS↓	PSNR (dB)↑	LPIPS↓	PSNR (dB)↑	LPIPS↓	PSNR (dB)↑
Identity	0.751	10.196	0.756	10.113	0.765	10.059	0.774	10.041
SE	0.755	9.744	0.755	9.729	0.762	9.722	0.772	9.734
CBAM	0.764	9.742	0.766	9.706	0.776	9.717	0.788	9.758
Non-local	0.752	10.384	0.756	10.366	0.764	10.361	0.774	10.334
FFC	0.748	**10.611**	**0.751**	**10.537**	**0.759**	**10.533**	0.769	**10.492**
Parameter-matched convolution	**0.747**	9.880	**0.751**	9.867	**0.759**	9.829	**0.768**	9.796
AGRL + TraceTrust	**0.744**	**10.654**	**0.750**	**10.650**	**0.756**	**10.629**	**0.765**	**10.604**

dB. This response measures screen reconstruction when wall evidence is represented on a coarser spatial grid, complementing the geometric experiments that vary its arrangement.

A.3.32 Additive Gaussian-Noise Robustness

We evaluate the representation modules as the labeled additive Gaussian-noise level increases.

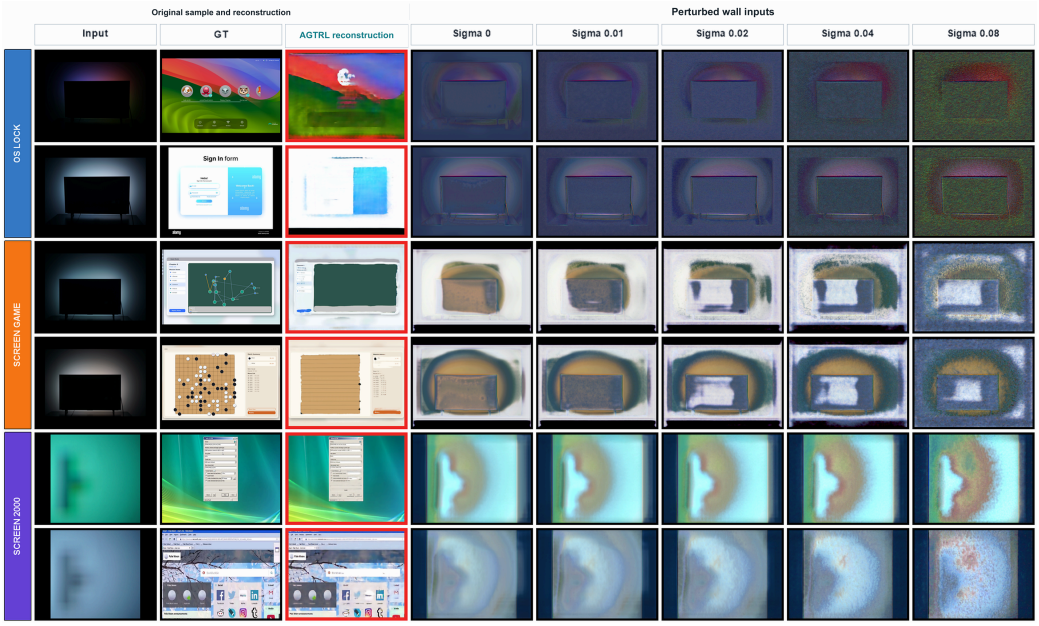

Supplementary Figure 39: Gaussian-noise examples from six samples. Columns 1–3 show the original wall observation, ground-truth screen, and reconstruction from the shared U-Net with AGRL + TraceTrust. Columns 4–8 show wall observations with additive Gaussian noise at $\sigma = 0, 0.01, 0.02, 0.04, 0.08$, respectively.

The representation comparison remains favorable as Gaussian noise increases. The proposed module leads both LPIPS and SSIM throughout the sweep. At $\sigma = 0.08$, its LPIPS is 0.859, compared with 0.920 for SE, the strongest alternative on this metric. This result extends the module's advantage to stronger fluctuations in the wall signal.

A.3.33 Fixed-Pattern Occlusion Robustness

We evaluate screen reconstruction as a fixed-pattern mask covers increasing portions of the wall observation.

The occlusion comparison asks how much recovery the remaining wall evidence can support. FFC combines local processing with image-wide spectral interactions (Chi et al., 2020); at 75% masking, it achieves 0.303 SSIM and the best alternative LPIPS of 0.836 in our representation-module comparison

Supplementary Table 28: Quantitative results for additive Gaussian-noise robustness. The listed σ levels give the standard deviation of additive noise on the RGB wall input, in the input image's intensity units; $\sigma = 0$ adds no noise. Lower LPIPS and higher SSIM indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	0		0.005		0.01		0.02		0.04		0.08	
	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑
Identity	0.751	0.408	0.766	0.394	0.796	0.368	0.830	0.311	0.911	0.240	0.972	0.202
SE	0.755	0.417	0.767	0.403	0.786	0.376	0.820	0.332	0.883	0.260	**0.920**	0.174
CBAM	0.764	0.436	0.776	0.430	0.798	0.411	0.831	0.358	0.896	0.281	0.985	0.214
Non-local	0.752	0.450	0.761	0.437	0.779	0.412	0.827	0.365	0.868	**0.309**	0.942	**0.226**
FFC	0.748	**0.465**	0.762	**0.452**	0.779	**0.434**	0.806	**0.389**	**0.857**	0.305	0.945	0.223
Parameter-matched convolution	**0.747**	0.437	**0.749**	0.430	**0.768**	0.410	**0.802**	0.361	0.868	0.282	0.941	0.206
AGRL + TraceTrust	**0.744**	**0.466**	**0.744**	**0.463**	**0.749**	**0.448**	**0.764**	**0.412**	**0.823**	**0.351**	**0.859**	**0.273**

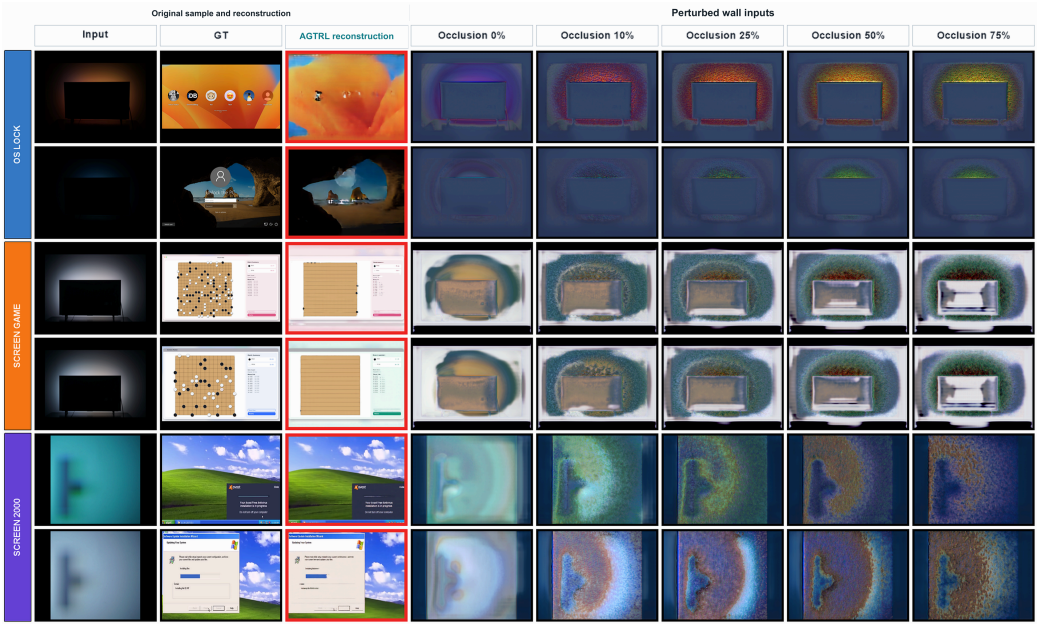

Supplementary Figure 40: Fixed-pattern occlusion examples from six samples. Columns 1–3 show the original wall observation, ground-truth screen, and reconstruction from the shared U-Net with AGRL + TraceTrust. Columns 4–8 show wall observations with masking ratios of 0%, 10%, 25%, 50%, and 75%, respectively.

Supplementary Table 29: Quantitative results for fixed-pattern occlusion robustness. Lower LPIPS and higher SSIM indicate better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	0%		10%		25%		50%		75%	
	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑	LPIPS↓	SSIM↑
Identity	0.751	0.408	0.820	0.333	0.842	**0.326**	0.854	**0.325**	0.860	**0.317**
SE	0.755	0.417	0.818	0.324	0.839	0.307	0.855	0.290	0.857	0.284
CBAM	0.764	0.436	0.825	0.332	0.863	0.309	0.879	0.296	0.872	0.298
Non-local	0.752	0.450	**0.807**	0.340	**0.830**	0.314	**0.845**	0.295	0.845	0.292
FFC	0.748	**0.465**	0.812	**0.344**	0.835	0.312	0.846	0.298	**0.836**	0.303
Parameter-matched convolution	**0.747**	0.437	0.811	0.333	0.832	0.315	0.849	0.301	0.847	0.309
AGRL + TraceTrust	**0.744**	**0.466**	**0.773**	**0.383**	**0.790**	**0.362**	**0.806**	**0.354**	**0.807**	**0.358**

(Supplementary Table 29). With AGRL organizing the remaining wall features and TraceTrust weighting their local support, the proposed module leads both metrics at every masking ratio and reaches 0.358 SSIM and 0.807 LPIPS at 75%, where Identity provides the strongest alternative SSIM of 0.317.

A.3.34 Wall Materials: Acquisition Photographs

Supplementary Figure 41 shows physical acquisition examples with different wall surfaces. Their texture and finish affect the appearance of the reflected screen light and the background against which that light is recorded. The photographs document the real surfaces and viewing arrangements. The following quantitative analysis evaluates the eight labeled wall-material appearance proxies in Supplementary Figure 42 and Supplementary Table 30.



Supplementary Figure 41: Real acquisition photographs illustrating wall-material variation. The panels show the display next to relay surfaces with different textures and finishes, viewed from behind the display and from the side. The reflected illumination varies with the surface appearance. These photographs provide physical examples of the acquisition conditions; the subsequent quantitative comparison uses the eight labeled wall-material appearance proxies.

A.3.35 Wall-Appearance Proxy Stability

We evaluate screen reconstruction across eight wall-appearance proxies.

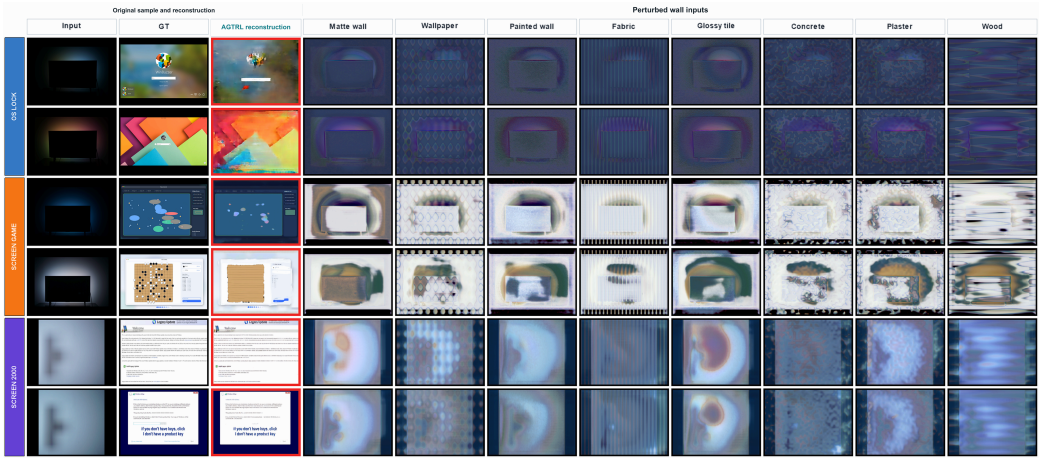

Supplementary Figure 42: Wall-appearance proxy examples from six samples. Columns 1–3 show the original wall observation, ground-truth screen, and reconstruction from the shared U-Net with AGRL + TraceTrust. The remaining columns show wall observations transformed using the labeled appearance proxies.

Wall appearance changes the background against which screen traces are observed. The proposed module achieves the highest SSIM for every labeled proxy, extending the measured recovery advantage across the tested wall-appearance variations.

A.3.36 Reference Reconstruction Quality for the Representation-Module Tests

We establish the representation-module reference from unmasked wall inputs, denoted Full. Metrics are computed within the valid bounding box of each paired sample.

Supplementary Table 30: Quantitative results for wall-appearance proxy stability. Higher SSIM indicates better reconstruction. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	Matte	Wallpaper	Painted	Fabric	Glossy	Concrete	Plaster	Wood
	SSIM↑	SSIM↑	SSIM↑	SSIM↑	SSIM↑	SSIM↑	SSIM↑	SSIM↑
Identity	0.390	0.359	0.298	0.303	0.322	0.356	0.344	0.358
SE	0.401	0.373	0.323	0.294	0.325	0.363	0.363	0.388
CBAM	0.428	0.387	0.345	0.320	0.366	0.372	0.385	0.394
Non-local	0.431	0.395	0.365	0.323	0.369	0.375	0.379	0.386
FFC	**0.450**	**0.411**	**0.384**	**0.349**	**0.403**	**0.392**	**0.397**	**0.398**
Parameter-matched convolution	0.429	0.373	0.345	0.325	0.367	0.375	0.374	0.378
AGRL + TraceTrust	**0.461**	**0.415**	**0.408**	**0.378**	**0.417**	**0.410**	**0.411**	**0.407**

Supplementary Table 31: Reference reconstruction quality for the representation-module tests. Lower Full LPIPS and higher Full SSIM and Full PSNR indicate better reconstruction; PSNR is reported in dB. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	Full LPIPS↓	Full SSIM↑	Full PSNR (dB)↑
Identity	0.7507	0.4084	10.1955
SE	0.7547	0.4167	9.7437
CBAM	0.7640	0.4360	9.7424
Non-local	0.7519	0.4498	10.3843
FFC	0.7476	**0.4651**	**10.6114**
Parameter-matched convolution	**0.7466**	0.4372	9.8801
AGRL + TraceTrust	**0.7444**	**0.4658**	**10.6544**

The proposed module achieves the best reference LPIPS, SSIM, and PSNR. These results establish its reconstruction advantage before the wall evidence is varied in contrast, spatial scale, noise, and visibility.

A.3.37 Reconstruction Quality Across Contrast Variation

We examine how changes in wall contrast relate to screen reconstruction. Supplementary Table 32 reports the mean reconstruction quality across the evaluated contrast levels.

Supplementary Table 32: Reconstruction quality across contrast variation. Higher mean contrast SSIM and PSNR indicate better reconstruction, and PSNR is reported in dB. Best and second-best values in each column are shown in bold red and bold blue, respectively. A pale red background highlights the proposed method, and up to three alternative rows with the most second-best entries are shaded pale blue.

Representation module	Mean contrast SSIM↑	Mean contrast PSNR (dB)↑
Identity	0.4300	10.4440
SE	0.4518	10.4197
CBAM	0.4617	10.3128
Non-local	0.4652	10.5722
FFC	**0.4809**	**10.7469**
Parameter-matched convolution	0.4520	10.0138
AGRL + TraceTrust	**0.4856**	**10.9868**

The proposed representation achieves the highest mean SSIM and PSNR across the contrast sweep. These mean scores complement the per-level results, showing that the proposed representation supports hidden-screen recovery across the tested contrast range, including faint wall patterns.