EVIDENCE-BUDGETED LIGHT TRANSPORT REPRESENTATION LEARNING: WALL-REFLECTED SCREEN CONTENT RECOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

Can an ordinary camera recover hidden screen content from the diffuse light that the screen casts onto a wall?We test this possibility with Evidence-Budgeted Light Transport Representation Learning (EBLTR).Wall transport preserves screen information unevenly across locations, scales, and frequency bands, so reconstruction must identify which wall cues remain reliable and how much local refinement they support.Our Light Transport Observability Representation (LTOR) combines frequency-separated responses, local wall statistics, and multilevel features to estimate the local observational support for screen reconstruction.Evidence-Budgeted Learning (EBL) turns this estimated support into a local correction budget that bounds the refinement magnitude in each region.A frozen screen-to-wall surrogate further checks observation consistency during training.Across four screen-content categories and varied observation conditions, EBLTR outperforms 19 reconstruction baselines and recovers screen colour, regional layout, and local structural cues.The reconstructions reduce erroneous edges, textures, interface contours, and text strokes that lack wall-observation support.These results demonstrate the preliminary feasibility of recovering hidden screen content from ordinary-camera recordings of wall illumination, with reconstructed detail guided by the evidence that survives light transport.

1 INTRODUCTION

Can the diffuse light on a wall reveal what a hidden screen is displaying? We investigate this question with an ordinary camera that records light cast onto wall by a screen hidden from direct view. Saunders et al. (2019) demonstrated ordinary-camera recovery of hidden pixel-art figures, letters, and colour bars from wall penumbra of an occluder. We ask whether ordinary-camera wall observations support the recovery of screen colour, regional layout, interface boundaries, and text-bearing structure in desktop windows, webpages, lock screens, and game interfaces.

Wall reflection preserves screen information unevenly across space, scale, and frequency. Broad colour regions can remain evident while text strokes, icon boundaries, and texture become dispersed. Recovery therefore requires two linked decisions: which wall cues still support hidden screen, and how much detail each region should receive. Local brightness and texture mix screen evidence with wall and sensor responses. Learned image priors can amplify these mixed responses during refinement, producing spurious edges and strokes. The central challenge is to make the reliability of surviving wall evidence explicitly limit the detail added during reconstruction.



Figure 1: Around-corner acquisition geometry in which the camera records diffuse wall illumination from a hidden screen.

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.

We propose Evidence-Budgeted Light Transport Representation Learning (EBLTR) to recover hidden screen content with local refinement bounded by surviving wall evidence. LTOR organizes frequency responses, local wall statistics, and multilevel context into an observability representation. EBL turns this representation into a local detail budget that explicitly bounds the correction added to the screen estimate. The recovered screen must also account for the wall response: a frozen screen-to-wall surrogate checks their consistency during training. Across four screen-content categories and varied observation conditions, EBLTR outperforms 19 reconstruction baselines, preserves screen colour and layout, and reduces unsupported local detail. Together, these results demonstrate preliminary hidden-screen recovery from ordinary-camera wall observations and support a clear refinement principle: surviving wall evidence should govern the detail added to the screen.

The main contributions are summarized as follows:

(1) We propose EBLTR, a unified framework for passive wall-reflected screen recovery that links wall-evidence representation, locally bounded refinement, and training-time observation consistency.

(2) We introduce LTOR to prioritize screen-relevant evidence within the wall observation, and EBL to align residual refinement with that evidence through a bounded correction budget.

(3) We evaluate EBLTR through reconstruction comparisons, controlled representation and learning studies, and robustness analyses across four screen-content categories, demonstrating preliminary hidden-screen recovery from ordinary-camera wall observations.

2 Related Work

Indirect-light imaging recovers hidden-scene properties by exploiting structure in light measured on visible surfaces.Time-of-flight measurements enabled hidden-shape reconstruction (Velten et al., 2012),while natural corner geometry and occluder-dependent intensity patterns supported passive inference and ordinary-camera periscopy (Bouman et al., 2017; Saunders et al., 2019). Faccio et al. (2020) place these developments within broader relationship between sensing and inverse reconstruction.Computational Mirrors uses deep matrix factorization of temporally varying indirect illumination to jointly recover hidden video and light transport (Aittala et al., 2019).Our EBLTR learns the wall-to-screen relation from paired observations and targets,enabling recovery of static display content with fixed model parameters at inference.Other learned approaches infer hidden-person activity from wall-illumination changes (Sharma et al., 2021) or recover trajectories through motion propagation and calibration (Wang et al., 2023).More recently,two-edge-resolved imaging combines orthogonal occluder edges with a tailored scene representation to recover three-dimensional colour information using an ordinary camera (Czajkowski & Murray-Bruce, 2024).These studies establish several ways to extract hidden-scene information from indirect light.EBLTR uses evidence remaining in the wall observation to determine where, and by how much, the recovered screen is refined.

Learned restoration provides complementary tools for representing image structure and adapting reconstruction to degraded observations.Encoder–decoder organization and residual learning provide foundational designs (Ronneberger et al., 2015; Zhang et al., 2017), while attention-based visual representations organize longer-range context (Dosovitskiy et al., 2021).NAFNet simplifies restoration blocks (Chen et al., 2022); Restormer and Uformer develop efficient long-range and window-based interactions (Zamir et al., 2022; Wang et al., 2022);and MIRNetv2 combines features across resolutions (Zamir et al., 2023). Observation-dependent adaptation is further developed through illumination-guided restoration in Retinexformer (Cai et al., 2023), task-specific visual prompts in DocRes (Zhang et al., 2024),and colour–intensity decoupling in HVI-CIDNet (Yan et al., 2025).AdaIR adaptively mines and modulates low- and high-frequency features for image restoration (Cui et al., 2025).EBLTR carries frequency-separated wall responses into explicit observability-dependent correction bounds,achieving a Game LPIPS of 0.4865 compared with 0.5754 for AdaIR in our evaluation (Table 1).Diffusion posterior sampling uses a known forward operator to guide iterative reconstruction toward the measurements (Chung et al., 2023).We combine evidence-bounded refinement with a frozen screen-to-wall surrogate for training-time observation consistency;at inference,EBLTR recovers screen content directly from the captured wall image.

3 METHOD

3.1 OVERALL ARCHITECTURE

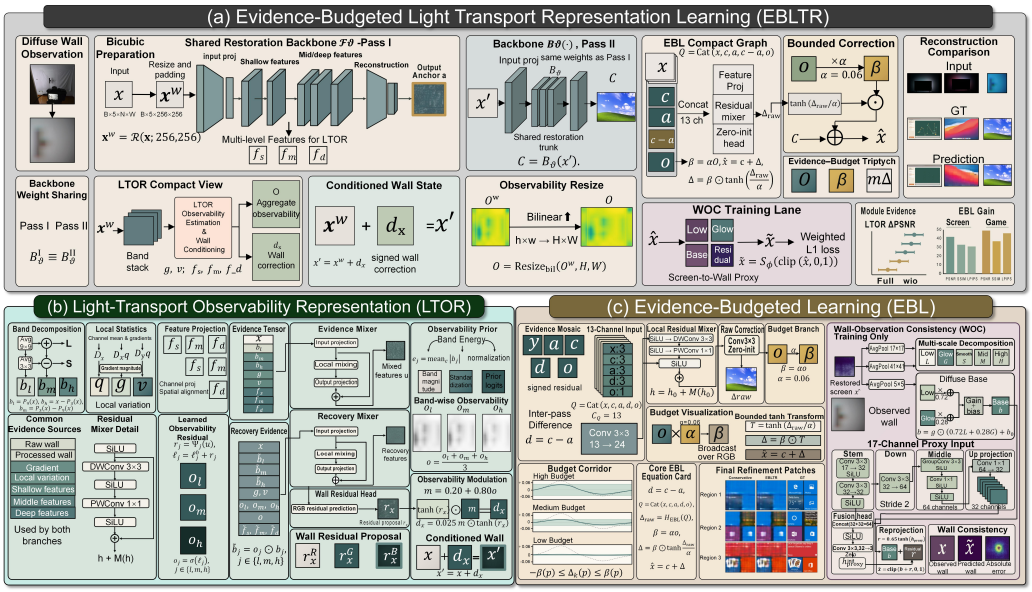

Figure 2: EBLTR architecture with solid inference paths and a dashed training-only observation check.

Given a wall observation $\mathbf{x} \in [0,1]^{B\times C\times H\times W}$, our goal is to recover the hidden screen $\widehat{\mathbf{Y}}$ from the wall light spot. EBLTR consists of a Light Transport Observability Representation (LTOR) and Evidence-Budgeted Learning (EBL). LTOR organizes wall-derived evidence across frequency bands and receptive ranges, and EBL converts the resulting observability into a pointwise bound on the reconstruction correction. A frozen screen-to-wall surrogate supplies a training-time observation check. Figure 2 summarizes the architecture, and Appendix A.2 provides the complete derivations.

For each colour response, we use the phenomenological sensing model

$$x_c(r) = \sum_{c'=1}^{C} \int_{\Omega_y} k_{cc'}(r,s) y_{c'}(s)\,\mathrm{d}s + a_c(r) + n_c(r), \qquad k_{cc'}(r,s) \geq 0.$$

Here $r \in \Omega_x$ and $s \in \Omega_y$ are wall and screen coordinates. The unknown kernel $k_{cc'}$ absorbs spatial propagation, diffuse reflection, colour mixing, and camera integration; a_c and n_c denote latent ambient response and residual perturbation. The effective-operator interpretation and its ill-posedness analysis are given in Appendix A.2.

EBLTR first forms an initial screen estimate with a shared reconstruction backbone. LTOR organizes the wall evidence and conditions the observation; the same backbone then produces the conservative reconstruction $\mathbf{C}$. EBL refines this estimate through a correction whose local magnitude is set by observability:

$$\begin{aligned} \mathbf{A} &= \mathcal{F}^0_\vartheta(\mathbf{x}), & \mathbf{f}_j &= \mathcal{F}^j_\vartheta(\mathbf{x}), \qquad j \in \{s,m,d\}, \\ (\mathbf{z}, \mathbf{O}, \mathbf{x}') &= \mathcal{M}_\theta(\mathbf{x}; \mathbf{f}_s, \mathbf{f}_m, \mathbf{f}_d), & \mathbf{C} &= \mathcal{F}^0_\vartheta(\mathbf{x}'), \\ \widehat{\mathbf{Y}} &= \mathbf{C} + \Delta_\psi(\mathbf{x}, \mathbf{A}, \mathbf{C}, \mathbf{O}). \end{aligned}$$

The shared backbone $\mathcal{F}_\vartheta$ produces the anchor $\mathbf{A}$ and shallow, middle, and deep features $\mathbf{f}_j \in \mathbb{R}^{B\times D_j\times H\times W}$. LTOR, denoted by $\mathcal{M}_\theta$, returns the evidence field $\mathbf{z} : \Omega_x \to \mathbb{E}$, aggregate observability $\mathbf{O} \in (0,1)^{B\times 1\times H\times W}$, and conditioned wall field $\mathbf{x}' \in \mathbb{R}^{B\times C\times H\times W}$. The initial and conditioned reconstruction passes share the same backbone parameters ϑ, while $\Delta_\psi(\mathbf{x}, \mathbf{A}, \mathbf{C}, \mathbf{O})$ denotes the EBL correction $\Delta\mathbf{Y}$. The learned parameters ϑ, θ, and ψ are fixed at inference, and every inference-time quantity derives from $\mathbf{x}$. Appendix A.2 gives the complete domain definitions and shared-parameter construction.

3.2 LIGHT TRANSPORT OBSERVABILITY REPRESENTATION

Diffuse transport generally preserves broad colour and layout more reliably than contours and fine texture. LTOR constructs broad, intermediate, and residual responses using fixed normalized smoothing operators:

$$\mathbf{b}_l = \Omega_{\tau_1}\mathbf{x}, \qquad \mathbf{b}_m = (\Omega_{\tau_2} - \Omega_{\tau_1})\mathbf{x}, \qquad \mathbf{b}_h = (I - \Omega_{\tau_2})\mathbf{x}, \qquad \mathbf{b}_l + \mathbf{b}_m + \mathbf{b}_h = \mathbf{x}.$$

The broader smoothing operator extracts large-scale colour structure, while differences between smoothing scales expose intermediate and residual responses. Their sum preserves the wall observation, allowing LTOR to assess broad layout cues and finer transitions together. Spatial gradients and local variation complement the three bands:

$$\begin{aligned} q(r) &= \frac{1}{C}\sum_{c=1}^{C} x_c(r), \\ g(r) &= \left(\sum_{\iota=1}^{2} |\partial_\iota q(r)|^2 + \epsilon\right)^{1/2}, \\ v(r) &= \left[\Omega_{\tau_v}(q^2)(r) - (\Omega_{\tau_v} q(r))^2\right]_+ + \eta\, \Omega_{\tau_v}\left(\frac{1}{C}\sum_{c=1}^{C} |b_{h,c}|\right)(r). \end{aligned}$$

Here q, g, and v are fields in $\mathbb{R}^{B\times 1\times H\times W}$ representing the colour mean, gradient magnitude, and nonnegative local variation. The derivative ∂_ι acts along spatial direction ι, Ω_{τ_v} is the fixed normalized local operator at scale τ_v, and $[\cdot]_+$ is the positive part. The stabilizer $\epsilon > 0$ and residual-energy weight $\eta \geq 0$ are fixed. Every term is computed from $\mathbf{x}$ and $\mathbf{b}_h$, so these wall-derived responses and the multilevel features form

$$\begin{aligned} \mathbf{z}(r) &= \big(\mathbf{x}(r), \mathbf{b}_l(r), \mathbf{b}_m(r), \mathbf{b}_h(r), g(r), v(r), \mathbf{f}_s(r), \mathbf{f}_m(r), \mathbf{f}_d(r)\big) \in \mathbb{E}, \\ \mathbb{E} &= \mathbb{R}^C \times (\mathbb{R}^C)^3 \times \mathbb{R}^2 \times \mathbb{R}^{D_s} \times \mathbb{R}^{D_m} \times \mathbb{R}^{D_d}. \end{aligned}$$

For each $j \in \{l, m, h\}$, the colour-averaged band magnitude is spatially standardized and mapped to an ordered prior:

$$\begin{aligned} e_j(r) &= \frac{1}{C}\sum_{c=1}^{C} |b_{j,c}(r)|, & \mu_j &= \frac{1}{|\Omega_x|}\int_{\Omega_x} e_j(r)\,\mathrm{d}r, \\ s_j &= \left(\frac{1}{|\Omega_x|}\int_{\Omega_x} (e_j(r) - \mu_j)^2\,\mathrm{d}r\right)^{1/2}, & n_j(r) &= \frac{e_j(r) - \mu_j}{\max(s_j, \epsilon)}, \\ \ell_j^0(r) &= \kappa n_j(r) + \rho_j, & p_j(r) &= \sigma(\ell_j^0(r)), \\ \kappa &> 0, & \rho_l &> \rho_m > \rho_h. \end{aligned}$$

Here e_j is the colour-averaged band magnitude, μ_j and s_j are its per-sample spatial mean and population standard deviation, and n_j is the stabilized standardized field. The positive slope κ preserves evidence ordering, while $\rho_l > \rho_m > \rho_h$ gives broad responses priority over high-frequency responses at equal standardized energy. The logistic output $p_j \in (0,1)^{B\times 1\times H\times W}$ is a deterministic wall-derived ordering cue that seeds contextual observability. Local context then corrects the prior log odds:

$$\begin{aligned} \mathbf{u}(r) &= \Phi_\theta\left(\mathbf{z}(r), \int_{\Omega_x} \gamma_\theta(r,s)[\mathbf{z}(s)]\,\mathrm{d}s\right), \quad \gamma_\theta(r,s) = 0 \qquad \text{for } s \notin \Omega_r, \\ \mathbf{O}_j(r) &= \sigma\big(\mathrm{logit}(p_j(r)) + \Psi_j^\theta[\mathbf{u}](r)\big), & \mathbf{O}(r) &= \frac{1}{3}\sum_{j\in\{l,m,h\}} \mathbf{O}_j(r). \end{aligned}$$

The learned coupling γ_θ has local support Ω_r and is distinct from $k_{cc'}$. The functional Φ_θ produces the contextual field $\mathbf{u}$, while Ψ_j^θ returns the scalar correction for band j. The bounded scores $\mathbf{O}_j$ and $\mathbf{O}$ order wall evidence, and $\{l, m, h\}$ indexes the complete three-band decomposition used by LTOR. They weight the bands before LTOR constructs the conditioned observation:

$$\widetilde{\mathbf{b}}_j = \mathbf{O}_j \otimes \mathbf{b}_j, \qquad j \in \{l, m, h\},$$
$$\mathbf{r}_x = \Gamma_\theta(\mathbf{x}, \widetilde{\mathbf{b}}_l, \widetilde{\mathbf{b}}_m, \widetilde{\mathbf{b}}_h, g, v, \mathbf{O}_l, \mathbf{O}_m, \mathbf{O}_h, \mathbf{O}, \mathbf{f}_s, \mathbf{f}_m, \mathbf{f}_d),$$
$$\mathbf{d}_x = \delta(\beta + \chi\mathbf{O}) \otimes \tanh(\mathbf{r}_x), \qquad \mathbf{x}' = \mathbf{x} + \mathbf{d}_x,$$
$$\beta \geq 0, \qquad \chi > 0, \qquad \beta + \chi = 1, \qquad \|\mathbf{d}_x(r)\|_\infty \leq \delta.$$

Thus, LTOR represents the effects of wall transport through spectral, spatial, and contextual evidence in the measured response. The bounded $\mathbf{d}_x$ changes the internal wall representation, and $\mathbf{x}'$ remains tied to the measured observation.

3.3 Evidence-Budgeted Learning and Observation Check

EBL turns aggregate observability into an explicit local correction limit around the conservative screen estimate. The tensors entering EBL share the computational grid used for refinement. The backbone produces the screen estimates, while the wall-derived tensors provide the local conditioning variables for the correction. EBL forms a local evidence field from the observation, conservative estimate, anchor, their difference, and observability, and then produces an unconstrained proposal:

$$\mathbf{z}_y(r) = \big(\mathbf{x}(r), \mathbf{C}(r), \mathbf{A}(r), \mathbf{C}(r) - \mathbf{A}(r), \mathbf{O}(r)\big) \in \mathbb{E}_y,$$
$$\mathbb{E}_y = \mathbb{R}^C \times (\mathbb{R}^C)^3 \times \mathbb{R}, \qquad \mathbf{D}(r) = \Phi_\psi\left(\mathbf{z}_y(r), \int_{\Omega_r} \gamma_\psi(r, s)[\mathbf{z}_y(s)]\,\mathrm{d}s\right).$$

All arguments of $\mathbf{D}$ derive from $\mathbf{x}$. The proposal enters through the observability-dependent feasible set

$$\mathbb{K}(\mathbf{O}; \mathbf{C}) = \left\{\mathbf{Y} \in \mathbb{R}^{B\times C\times H\times W} : |Y_{b,c,i,j} - C_{b,c,i,j}| \leq \alpha O_{b,1,i,j} \text{ for all } (b, c, i, j)\right\},$$

where b, c, i, and j index batch, colour, row, and column, respectively. The fixed $\alpha > 0$ is the maximum correction scale, and the scalar observability is broadcast over colour. Thus, $\mathbb{K}(\mathbf{O}; \mathbf{C})$ is an exact evidence-dependent feasible set centered at $\mathbf{C}$; $\mathbf{O}$ orders available evidence for budgeting. EBL maps the proposal smoothly into this set:

$$\Delta\mathbf{Y} = \alpha\mathbf{O} \otimes \tanh(\mathbf{D}/\alpha), \qquad \widehat{\mathbf{Y}} = \mathbf{C} + \Delta\mathbf{Y} \in \mathbb{K}(\mathbf{O}; \mathbf{C}),$$
$$|\Delta Y_{b,c,i,j}| \leq \alpha O_{b,1,i,j} \leq \alpha.$$

Consequently, low observability limits the magnitude of added detail while $\mathbf{C}$ remains the center of the admissible reconstruction. Appendix A.2 gives the exact norm and proposal-sensitivity bounds. Czajkowski & Murray-Bruce (2024) recover three-dimensional colour from ordinary-camera penumbrae through geometry-based inversion with sparsity and total-variation regularization. Our EBL makes refinement control local: for fixed observability $\mathbf{O}$, each correction has magnitude at most $\alpha\mathbf{O}$ and sensitivity to its corresponding residual proposal at most $\mathbf{O}$, permitting larger, more responsive updates where wall evidence is stronger and keeping weaker-evidence regions close to $\mathbf{C}$. During training only, the frozen surrogate supplies the auxiliary observation check

$$\widetilde{\mathbf{X}} = \mathcal{S}_\phi(\widehat{\mathbf{Y}}), \qquad L_1 = d(\widetilde{\mathbf{X}}, \mathbf{x}).$$

The observation-consistency term encourages the recovered screen to reproduce the captured wall response through the frozen surrogate. Screen reconstruction and wall consistency thus guide learning from the target and observation domains, respectively. The diffuse-plus-bounded-discrepancy parameterization of $\mathcal{S}_\phi$ is given in Appendix A.2. Its parameters ϕ are frozen during EBLTR training; inference follows the direct wall-to-screen path through the fixed parameters ϑ, θ, and ψ.

4 Experiment

4.1 Experimental Setup

We evaluate four screen-content categories—Screen, Game, Lock, and Home—using paired wall observations and screen targets. Screen retains fixed training, validation, and test lists containing 2121, 124, and 251 samples. Game uses 900, 100, and 250 samples, Lock uses 1008, 112, and 279 samples, and Home uses 2615, 291, and 726 samples for training, validation, and testing, respectively. For Game, Lock, and Home, each validation set is a fixed tenth of the original training portion. All compared models and ablations use the same lists. Each input is a colour wall image, and each target is the synchronously recorded screen image. Original aspect ratios are retained on a common landscape canvas, and evaluation is restricted to the valid image region.

We evaluate pixel fidelity with PSNR in dB (Huynh-Thu & Ghanbari, 2008), structural agreement with SSIM (Wang et al., 2004), and perceptual similarity with LPIPS (Zhang et al., 2018). Higher PSNR and SSIM and lower LPIPS indicate better reconstruction. The complete model, every listed baseline, and every ablation receive the same wall input at test time. Validation data are used for model selection under the shared evaluation, while the fixed test sets are reserved for the results.

4.2 Comparison with State-of-the-Art



Figure 3: Visual comparisons on Screen, Home, Lock, and Game. Each category contains two examples, with the wall input, ground-truth screen, EBLTR reconstruction, and 11 baseline reconstructions. Table 1 presents the quantitative comparison with all 19 baselines.

Figure 3 examines what screen content becomes recoverable from diffuse wall observations. EBLTR recovers screen-wide colour regions and principal content divisions from diffuse wall observations,

with interface and text-bearing structure becoming visible in the Screen examples. Several baseline outputs retain more of the wall-like appearance or merge distinct target regions. The comparison therefore connects the wall observation to recognizable organization in the hidden display. EBLTR leads all twelve content–metric columns in Table 1. Restormer (Zamir et al., 2022) models long-range interactions with an efficient Transformer for image restoration. Against this strongest Screen baseline in our evaluation, EBLTR raises PSNR from 22.60 to 24.27 dB—a gain of 1.67 dB—improves SSIM from 0.7774 to 0.8189, and reduces LPIPS from 0.3449 to 0.2913. We next test what enables this recovery: evidence representation, budgeted refinement, and their combination.

Table 1: Quantitative comparison on the fixed Screen, Game, Lock, and Home test sets (PSNR↑/SSIM↑/LPIPS↓). Best in red, second best in blue.

Method	Source	Screen			Game			Lock			Home		
		PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Classical CNN Baselines													
DnCNN (Zhang et al., 2017)	TIP'17	8.88	0.3466	0.7245	8.74	0.4275	0.7258	11.82	0.4787	0.7221	9.02	0.3979	0.7900
U-Net (Ronneberger et al., 2015)	MICCAI'15	9.18	0.2401	0.8061	10.29	0.4715	0.7581	12.18	0.4874	0.7618	8.73	0.3726	0.8223
General-Purpose and Multi-Task Image Restoration Frameworks													
AdaIR (Cui et al., 2025)	ICLR'25	21.71	0.7621	0.3873	17.44	0.6394	0.5754	15.17	0.5729	0.6455	13.96	0.5338	0.7324
NAFNet (Chen et al., 2022)	ECCV'22	21.62	0.7585	0.3794	18.16	0.6587	0.5759	15.24	0.5619	0.6637	14.58	0.5129	0.7203
Restormer (Zamir et al., 2022)	CVPR'22	22.60	0.7774	0.3449	17.37	0.6438	0.5676	16.75	0.6071	0.6329	12.80	0.5036	0.7451
Uformer (Wang et al., 2022)	CVPR'22	19.20	0.7068	0.4639	17.55	0.6470	0.6142	15.15	0.5652	0.6743	14.34	0.5356	0.7392
MIRNetv2 (Zamir et al., 2023)	TPAMI'23	18.24	0.6917	0.4923	17.68	0.6559	0.6120	15.21	0.6042	0.6575	12.79	0.5342	0.7598
AST (Zhou et al., 2024)	CVPR'24	19.39	0.7050	0.4668	17.46	0.6453	0.6158	15.07	0.5366	0.6920	12.90	0.4960	0.7672
ConvIR (Cui et al., 2024)	TPAMI'24	20.27	0.7283	0.4495	17.64	0.6515	0.6013	14.91	0.5925	0.6867	12.67	0.5168	0.7872
DeepSN-Net (Deng et al., 2025)	TPAMI'25	10.49	0.2954	0.7787	10.87	0.4440	0.7390	12.84	0.4228	0.7711	10.67	0.3751	0.8343
SFHformer (Jiang et al., 2024)	ECCV'24	8.82	0.3940	0.8066	10.01	0.4620	0.7598	12.46	0.5095	0.7404	10.03	0.4480	0.8530
Models for Related Degradations													
DocRes (Zhang et al., 2024)	CVPR'24	21.15	0.7490	0.3999	17.53	0.6472	0.5795	16.36	0.6094	0.6467	12.69	0.5153	0.7378
Retinexformer (Cai et al., 2023)	ICCV'23	19.57	0.6732	0.4554	17.93	0.6556	0.5924	14.60	0.5586	0.6506	12.66	0.5078	0.7343
HVI-CIDNet (Yan et al., 2025)	CVPR'25	19.69	0.7178	0.4724	17.77	0.6555	0.5979	15.61	0.5758	0.6714	13.85	0.5331	0.7420
Ada-Deblur (Wu et al., 2023)	TIP'23	19.43	0.7067	0.4637	17.75	0.6577	0.5878	16.27	0.6086	0.6608	14.06	0.5358	0.7522
DEA-Net (Chen et al., 2024)	TIP'24	20.65	0.7500	0.4240	17.73	0.6579	0.5879	15.96	0.6069	0.6742	13.69	0.5334	0.7610
DehazeFormer (Song et al., 2023)	TIP'23	19.37	0.6931	0.4794	17.77	0.6494	0.6074	14.87	0.5375	0.6687	12.79	0.4704	0.7511
Pixel2Pixel (Ma et al., 2025)	TPAMI'25	10.32	0.4816	0.6637	11.14	0.4502	0.6917	12.41	0.5178	0.7303	10.42	0.3787	0.7669
DarkIR (Feijoo et al., 2025)	CVPR'25	9.22	0.1749	0.8673	4.46	0.0409	0.9361	8.79	0.0949	0.8040	6.58	0.0665	0.8897
EBLTR	**Ours**	**24.27**	**0.8189**	**0.2913**	**18.71**	**0.6665**	**0.4865**	**17.16**	**0.6295**	**0.6299**	**14.69**	**0.5462**	**0.7077**

4.3 Ablation Study

Table 2: Overall ablation and replacement studies of LTOR and EBL on Screen, Game, Lock, and Home. Best in red, second best in blue.

Method	Screen			Game			Lock			Home		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
LTOR Overall Ablation and Replacement												
w/o LTOR	24.06	0.8139	0.3175	18.40	0.6633	0.5463	17.05	0.6288	0.6361	14.47	0.5425	0.7362
LTOR→Constant Map	23.42	0.8043	0.3321	18.26	0.6653	0.5596	16.85	0.6293	0.6323	14.56	0.5455	0.7362
LTOR →Band-Shuffled Evidence	23.80	0.8108	0.3250	18.28	0.6638	0.5588	16.85	0.6212	0.6347	14.62	0.5435	0.7390
LTOR→Flat Evidence	23.61	0.8096	0.3259	18.30	0.6646	0.5552	16.92	0.6272	0.6334	14.60	0.5445	0.7387
EBL Overall Ablation and Replacement												
w/o EBL	23.72	0.8098	0.3288	18.26	0.6657	0.5600	16.92	0.6289	0.6351	14.54	0.5411	0.7346
EBL→Constant Low Budget	23.63	0.8107	0.3257	18.30	0.6623	0.5581	16.76	0.6263	0.6405	14.59	0.5454	0.7375
EBL→Shuffled Low Budget	23.58	0.8069	0.3244	18.32	0.6651	0.5559	16.93	0.6292	0.6404	14.67	0.5447	0.7412
EBL→Plain Residual Refiner	23.72	0.8083	0.3264	17.74	0.6557	0.5689	17.01	0.6243	0.6436	14.42	0.5407	0.7354
EBLTR (full)	**24.27**	**0.8189**	**0.2913**	**18.71**	**0.6665**	**0.4865**	**17.16**	**0.6295**	**0.6299**	**14.69**	**0.5462**	**0.7077**

The complete LTOR–EBL chain leads all twelve columns in Table 2. Replacing evidence organization with constant, shuffled, or flattened alternatives reduces reconstruction quality, as does replacing evidence-budgeted refinement with the corresponding simpler routes. These chain-level comparisons test the two linked decisions introduced in the method: locating useful wall evidence and controlling the correction applied to the screen estimate. Their combination produces the strongest recovery across all four content categories.

4.4 Robustness to Distance and Effective Sampling

Does hidden-screen structure remain recoverable under changes in the effective spatial sampling of the wall response? The Lock study compares 20 methods on 279 test samples at five factors from 1.10 to 1.30. EBLTR leads all three metrics at every factor in Tables 3 and B16, as visualized in

Figure 4. The corresponding reconstructions in Figure B20 retain the screen's broad colour bands and principal divisions, showing that the recovery advantage persists as the sampled observation changes.

Table 3: Complete Lock distance and effective-sampling-resolution data for all methods and metrics.

Method	Sampling factor 1.10			Sampling factor 1.15			Sampling factor 1.20			Sampling factor 1.25			Sampling factor 1.30		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4792	0.7240	11.82	0.4791	0.7242	11.82	0.4792	0.7244	11.82	0.4792	0.7244	11.82	0.4792	0.7249
U-Net	12.15	0.4868	0.7638	12.14	0.4868	0.7637	12.16	0.4877	0.7639	12.15	0.4866	0.7647	12.15	0.4868	0.7649
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.17	0.5724	0.6462	15.17	0.5726	0.6458	15.17	0.5730	0.6463	15.17	0.5725	0.6462	15.17	0.5728	0.6458
NAFNet	15.24	0.5626	0.6641	15.24	0.5626	0.6643	15.24	0.5625	0.6642	15.24	0.5627	0.6646	15.24	0.5628	0.6647
Restormer	16.72	0.6051	0.6361	16.72	0.6054	0.6360	16.68	0.6048	0.6377	16.58	0.6033	0.6390	16.57	0.6035	0.6392
Uformer	15.18	0.5661	0.6755	15.18	0.5657	0.6761	15.18	0.5661	0.6754	15.17	0.5661	0.6756	15.17	0.5663	0.6757
MIRNetv2	15.21	0.6044	0.6580	15.22	0.6045	0.6580	15.22	0.6045	0.6583	15.22	0.6045	0.6582	15.22	0.6045	0.6583
AST	15.07	0.5370	0.6930	15.07	0.5369	0.6930	15.07	0.5369	0.6930	15.07	0.5369	0.6935	15.08	0.5370	0.6936
ConvIR	9.72	0.2907	0.7803	9.72	0.2907	0.7804	9.72	0.2907	0.7807	9.72	0.2907	0.7806	9.72	0.2907	0.7810
DeepSN-Net	12.84	0.4232	0.7739	12.84	0.4232	0.7740	12.84	0.4233	0.7741	12.84	0.4233	0.7748	12.84	0.4234	0.7751
SFHformer	12.46	0.5099	0.7404	12.46	0.5098	0.7406	12.46	0.5097	0.7404	12.46	0.5100	0.7407	12.46	0.5100	0.7408
Models for Related Degradations															
DocRes	16.35	0.6111	0.6474	16.35	0.6110	0.6468	16.41	0.6125	0.6461	16.39	0.6109	0.6477	16.39	0.6116	0.6466
Retinexformer	14.61	0.5599	0.6511	14.61	0.5599	0.6510	14.61	0.5601	0.6513	14.61	0.5602	0.6514	14.61	0.5603	0.6515
HVI-CIDNet	15.60	0.5749	0.6778	15.62	0.5755	0.6769	15.60	0.5750	0.6771	15.59	0.5752	0.6744	15.60	0.5751	0.6769
Ada-Deblur	16.21	0.6060	0.6609	16.25	0.6073	0.6608	16.19	0.6055	0.6611	16.18	0.6053	0.6611	16.15	0.6041	0.6611
DEA-Net	15.54	0.5952	0.6863	15.68	0.5976	0.6850	15.70	0.5966	0.6817	15.36	0.5898	0.6860	15.65	0.5950	0.6825
DehazeFormer	14.87	0.5381	0.6687	14.87	0.5381	0.6688	14.87	0.5381	0.6688	14.87	0.5382	0.6689	14.87	0.5382	0.6689
Pixel2Pixel	12.41	0.5179	0.7302	12.41	0.5178	0.7301	12.41	0.5179	0.7303	12.41	0.5179	0.7303	12.41	0.5179	0.7303
DarkIR	8.79	0.0950	0.8053	8.79	0.0950	0.8053	8.79	0.0950	0.8055	8.79	0.0950	0.8057	8.79	0.0950	0.8059
EBLTR	**17.01**	**0.6268**	**0.6319**	**17.02**	**0.6273**	**0.6314**	**16.99**	**0.6265**	**0.6319**	**16.76**	**0.6228**	**0.6352**	**16.63**	**0.6206**	**0.6361**

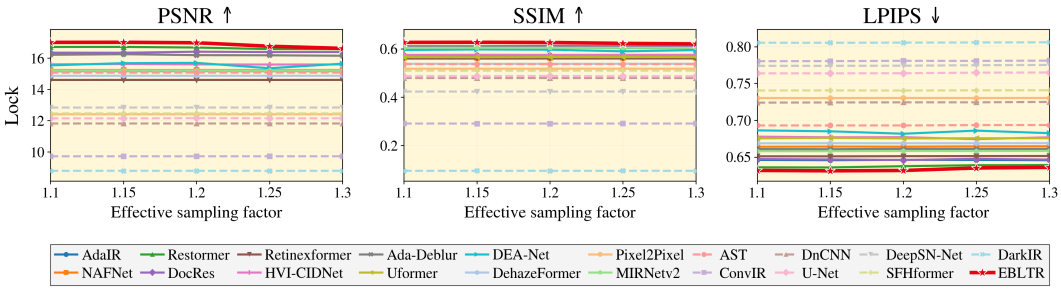

Figure 4: Lock distance and effective-sampling-resolution robustness for 20 methods at five factors, with PSNR and SSIM higher being better and LPIPS lower being better.

4.5 Same-Backbone Representation Comparison

Table 4: Complete same-backbone representation comparison on Screen, Game, Lock, and Home. Red and blue indicate the best and second-best distinct displayed values; ties share the same colour.

Representation	Screen			Game			Lock			Home		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Uniform	23.57	0.8039	0.3146	18.75	0.6658	0.4912	16.96	0.6215	0.6307	14.57	0.5373	0.7054
CNN-Pyramid	23.67	0.8045	0.3137	18.75	0.6656	0.4913	16.96	0.6213	0.6309	14.54	0.5381	0.7046
Window-Attention	23.65	0.8039	0.3146	18.75	0.6657	0.4913	16.96	0.6215	0.6307	14.58	0.5375	0.7053
Frequency-Prompt	23.62	0.8041	0.3143	18.75	0.6657	0.4913	16.96	0.6215	0.6308	14.55	0.5381	0.7047
LTOR	**23.90**	**0.8096**	**0.3045**	**18.76**	**0.6671**	**0.4864**	**16.97**	**0.6238**	**0.6304**	**14.59**	**0.5387**	**0.7041**

LTOR achieves the best reconstruction scores across all four content categories with the backbone and learning rule fixed in Table 4. The boundary diagnostics reveal what this advantage changes: relative to CNN-Pyramid, the false-edge rate in low-observability regions defined by the fixed evaluation proxy decreases from 0.4233 to 0.4118 in Table B7. Improved image quality is therefore accompanied by fewer erroneous boundaries in these low-observability regions. This connects LTOR's spatial–scale–frequency organization directly to the goal of recovering screen structure with less spurious local detail.

Figure 5 compares five representations under the same restoration backbone, decoder, learning rule, wall observation, and target; the representation is the controlled variable. The first column presents the wall input, ground truth, and enlarged target detail. Each remaining column shows the complete

reconstruction, the shared red-box region, its absolute error relative to the ground truth, and the three image-quality measures. The enlarged desktop region makes the representation advantage visible: LTOR retains clearer text-bearing boundaries and window structure. Its PSNR increases from 24.33 to 25.47 dB relative to Uniform, connecting the local improvement to the full-image reconstruction. The aligned gains in global quality, local detail, and perceptual similarity show that observability-aware representation strengthens hidden-screen recovery.

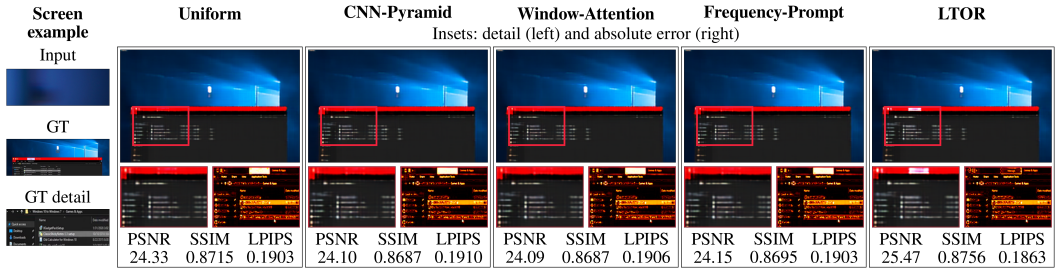

Figure 5: Same-backbone representation comparison on Screen. Red boxes mark the reconstruction regions shown in enlarged panels; the ground truth provides corresponding screen reference. Enlarged panels show recovered detail and its error. Absolute error is averaged over RGB channels and displayed with a shared scale set by the 98th percentile of pooled errors in the enlarged regions.

4.6 3-by-3 Grid Occlusion and Spatial Influence Heatmap Analysis

For one Screen test wall image, the nine equal-area occlusions in Figure 6 reveal spatially uneven reconstruction influence. The influence scores range from 1.00 for the middle-left region and 0.92 for the upper-left region to 0.00 and 0.01 for the middle-right and lower-right regions. Equal-area occlusions produce markedly different reconstruction changes. In this example, the left-side regions exert greater influence than the right-side regions, showing that the contribution of a wall observation is spatially uneven. This variation motivates a local observability representation and location-dependent refinement budgets.

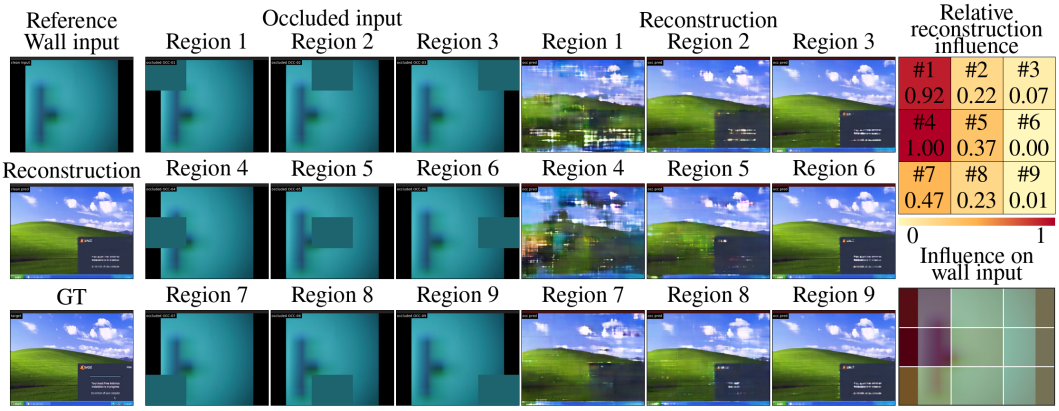

Figure 6: Nine-region occlusion responses for one Screen test example, including reconstructions and a relative reconstruction influence map. For each metric, positive degradation magnitudes are min–max scaled across the nine regions; the score averages the three scaled magnitudes.

5 Conclusion

We demonstrate the preliminary feasibility of recovering hidden screen content from diffuse wall illumination recorded by an ordinary camera. EBLTR organizes the evidence that survives wall transport, converts it into local refinement budgets, and connects the recovered screen back to the wall observation during training. Across four screen-content categories and varied observation conditions, this design outperforms 19 reconstruction baselines while preserving screen colour, layout, and recoverable local structure. The diagnostic and visual results connect these gains to fewer spurious edges, textures, interface contours, and text strokes. Previous work has recovered hidden motion by jointly estimating scene content and light transport (Aittala et al., 2019), and has obtained three-dimensional colour reconstructions through two-edge-resolved imaging with an ordinary camera (Czajkowski & Murray-Bruce, 2024). EBLTR makes the wall observation govern the detail added to the recovered screen, providing an evidence-budgeted learning route from diffuse wall illumination to hidden display content.

REFERENCES

Miika Aittala, Prafull Sharma, Lukas Murmann, Adam Yedidia, Gregory Wornell, Bill Freeman, and Fredo 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/hash/5a2afca61e35f45a7dd44ca46e0225f4-Abstract.html.

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 *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2289–2297, 2017. doi: 10.1109/ICCV.2017.249.

Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12470–12479, 2023. doi: 10.1109/ICCV51070.2023.01149.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *Computer Vision – ECCV 2022*, volume 13667 of *Lecture Notes in Computer Science*, pp. 17–33, 2022. doi: 10.1007/978-3-031-20071-7_2.

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.

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

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, 2024. doi: 10.1109/TPAMI.2024.3419007.

Yuning Cui, Syed Waqas Zamir, Salman H. Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz Khan. AdaIR: Adaptive all-in-one image restoration via frequency mining and modulation. In *The Thirteenth International Conference on Learning Representations (ICLR)*. OpenReview.net, 2025. URL https://openreview.net/forum?id=M5t0WvjfCg.

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.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *The Ninth International Conference on Learning Representations (ICLR)*. OpenReview.net, 2021. URL https://openreview.net/forum?id=YicbFdNTTy.

Daniele Faccio, Andreas Velten, and Gordon Wetzstein. Non-line-of-sight imaging. *Nature Reviews Physics*, 2(6):318–327, 2020. doi: 10.1038/s42254-020-0174-8.

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*, pp. 10879–10889, 2025. doi: 10.1109/CVPR52734.2025.01016.

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

Xingyu Jiang, Xiuhui Zhang, Ning Gao, and Yue Deng. When fast fourier transform meets transformer for image restoration. In *Computer Vision – ECCV 2024*, pp. 381–402, 2024. doi: 10.1007/978-3-031-72995-9_22.

Qing Ma, Junjun Jiang, Xiong Zhou, Pengwei Liang, Xianming Liu, and Jiayi Ma. Pixel2Pixel: A pixelwise approach for zero-shot single image denoising. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(6):4614–4629, 2025. doi: 10.1109/TPAMI.2025.3546870.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351 of *Lecture Notes in Computer Science*, 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.

Prafull Sharma, Miika Aittala, Yoav Y. Schechner, Antonio Torralba, Gregory W. Wornell, William T. Freeman, and Frédo Durand. What you can learn by staring at a blank wall. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2330–2339, 2021.

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.

Yihao Wang, Zhigang Wang, Bin Zhao, Dong Wang, Mulin Chen, and Xuelong Li. Propagate and calibrate: Real-time passive non-line-of-sight tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 972–981, 2023.

Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17662–17672, 2022. doi: 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.

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*, pp. 5678–5687, 2025. doi: 10.1109/CVPR52734.2025.00533.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5718–5729, 2022. doi: 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.

Jiaxin Zhang, Dezhi Peng, Chongyu Liu, Peirong Zhang, and Lianwen Jin. DocRes: A generalist model toward unifying document image restoration tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15654–15664, 2024. doi: 10.1109/CVPR52733.2024.01482.

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*, pp. 586–595, 2018. doi: 10.1109/CVPR.2018.00068.

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*, pp. 2952–2963, 2024. doi: 10.1109/CVPR52733.2024.00285.

APPENDIX FOR EVIDENCE-BUDGETED LIGHT TRANSPORT REPRESENTATION LEARNING: WALL-REFLECTED SCREEN CONTENT RECOVERY

A DATA-ACQUISITION PROTOCOL, MECHANISTIC EXPLANATION, AND COMPLETE EXPERIMENTAL SUPPLEMENT FOR LTOR AND EBL

A.1 DATA ACQUISITION PROTOCOL AND COMPLETE MAIN COMPARISON

The experiments use four subsets of the computer-screen dataset: Screen, Home, Lock, and Game. They share the same paired recovery objective while covering different screen-content organizations. Screen is recorded in an indoor occlusion scene: a 60-inch self-emissive LED display has a T-shaped baffle placed approximately 0.5 m in front of it. The display illuminates a white diffuse wall, while the T-shaped baffle shapes the observed shadow pattern. A Sony A7SII camera records this wall response from approximately 1.8 m away. Each sample pairs this wall-shadow image with the synchronized screen image containing webpages, desktops, charts, user interfaces, or password characters. The task is to recover the hidden display content from the captured wall-shadow observation.

Home, Lock, and Game are recorded in an indoor reflection scene. A 60-inch self-emissive LED display faces a white diffuse wall at a separation of approximately 0.5 m and forms a bright wall spot. The Sony A7SII camera observes the wall from behind the display, recording the wall spot together with the display-back silhouette. Each sample pairs this composite wall observation with the synchronized screen image. The three subsets provide varied desktop-window arrangements, broad lock-screen bands, and regular game layouts, so the same wall-to-screen mapping is tested across distinct content structures.

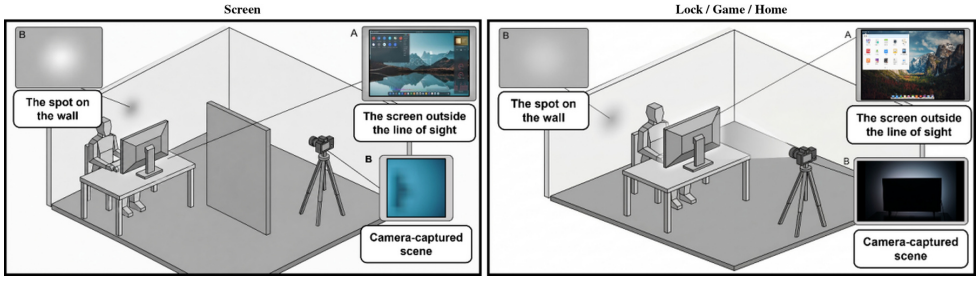

Figure A1: The two acquisition geometries used by the computer-screen dataset. Screen uses the indoor occlusion arrangement, whereas Home, Lock, and Game use the direct wall-reflection arrangement.



Figure A2: Real acquisition setup for the Screen subset in the indoor occlusion scene.

Screen examines desktop windows and text-bearing webpages in the occlusion scene. Each column in Figure A3 pairs a wall-shadow input in the upper row with its screen target below. The desktop and webpage examples show how distinct window arrangements and text-bearing regions produce diffuse wall responses. Screen recovery connects these responses to the target's page divisions, window positions, and local boundaries.

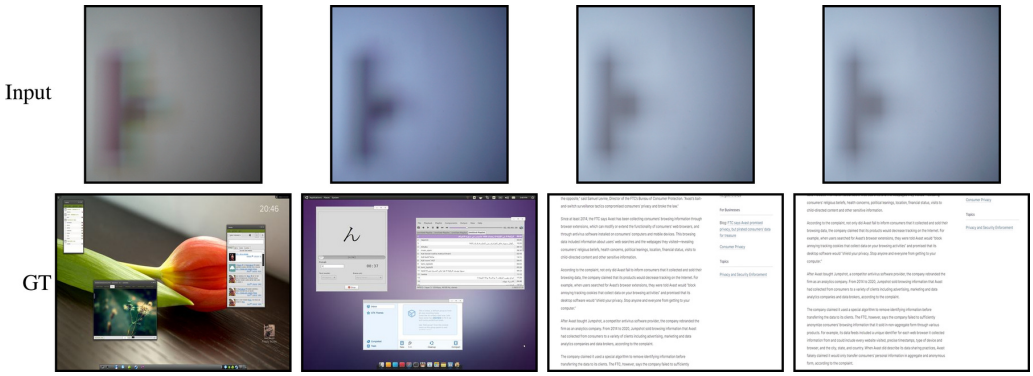

Figure A3: Four Screen examples from the occlusion scene. Wall-shadow inputs are shown above their paired desktop and webpage targets.

Home examines the arrangement of application windows within an operating-system desktop. Figure A4 shows file-manager, settings, terminal, and welcome-window examples against different desktop backgrounds. Each input contains the display-back silhouette surrounded by diffuse wall illumination, while its paired target specifies the window positions and background composition. The blue, green, and violet wall responses vary with the displayed content. Recovering the dominant colour distribution and window layout from these responses connects the Home examples directly to the paper's screen-recovery objective.

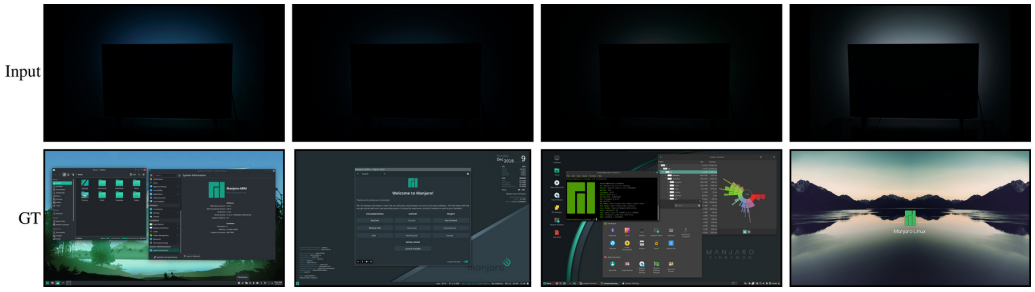

Figure A4: Four Home examples from the reflection scene. The upper row shows the captured display-back silhouette and surrounding wall illumination; the lower row shows the paired desktop targets with different application-window arrangements.

Lock examines large wallpaper regions together with clock or login elements. Figure A5 pairs mountain, night-sky, and green-landscape backgrounds, together with a plain login interface, with their wall observations. Their paired observations show different brightness levels and colour spreads around the display-back silhouette. These examples connect broad background recovery with the placement of smaller interface elements: the wallpaper determines the broad scene layout and major boundaries, while the overlaid interface occupies smaller regions. The pairs place the recovery of global screen organization and principal contours in a familiar lock-screen setting.

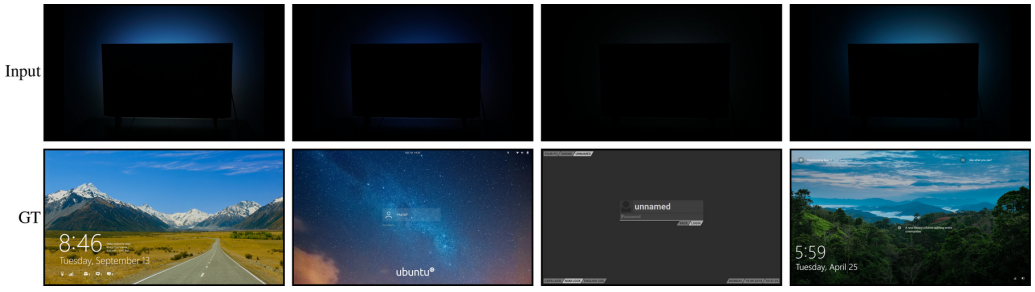

Figure A5: Four Lock examples from the reflection scene. Wall observations are shown in the upper row, with their paired wallpaper, clock, and login-screen targets below.

Game examines structured playing areas and their surrounding interface panels. Figure A6 includes light puzzle grids and blue strategy-game views, with status or control panels beside the main playing area. The paired wall inputs range from bright, nearly neutral illumination to darker blue responses. This content variation makes both the overall brightness and the division between the playing area and side panel relevant reconstruction targets, extending the same wall-to-screen recovery question to game interfaces.

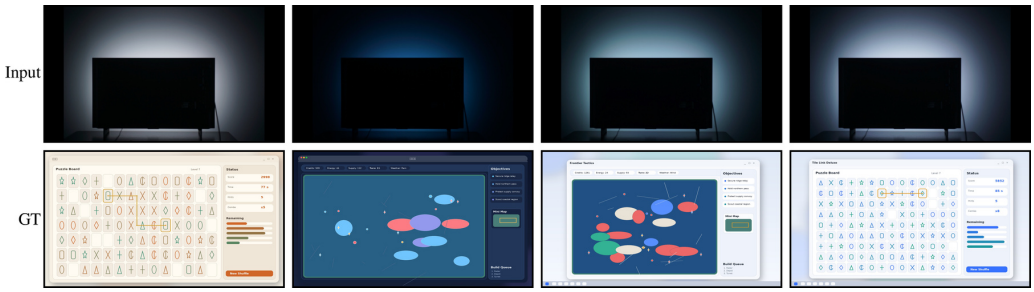

Figure A6: Four Game examples from the reflection scene. The upper row shows wall observations, and the lower row shows their paired game targets, including the playing area and adjacent interface panels.

The acquisition settings share a common observation path: light from a self-emissive display reaches a diffuse wall, and the camera records the wall response while the displayed content remains hidden from direct view. Figure A1 shows the occlusion and direct-reflection arrangements used for the four content categories. Synchronous wall and screen captures form the paired training and evaluation samples, while inference takes the wall image as its observation.

The four content categories test the same recovery principle through distinct screen structures. The fixed splits are 2121 training, 124 validation, and 251 test samples for Screen; 900, 100, and 250 for Game; 1008, 112, and 279 for Lock; and 2615, 291, and 726 for Home. Every method uses the same lists and valid evaluation regions, so Appendix Table A1 compares how effectively each model turns an identical wall observation into the hidden screen's brightness distribution, regional layout, and principal contours.

The main comparison places EBLTR against the complete set of 19 baselines: AdaIR (Cui et al., 2025), NAFNet (Chen et al., 2022), Restormer (Zamir et al., 2022), DocRes (Zhang et al., 2024), Retinexformer (Cai et al., 2023), HVI-CIDNet (Yan et al., 2025), Ada-Deblur (Wu et al., 2023), Uformer (Wang et al., 2022), DEA-Net (Chen et al., 2024), DehazeFormer (Song et al., 2023), Pixel2Pixel (Ma et al., 2025), MIRNetv2 (Zamir et al., 2023), AST (Zhou et al., 2024), ConvIR (Cui et al., 2024), DnCNN (Zhang et al., 2017), U-Net (Ronneberger et al., 2015), DeepSN-Net (Deng et al., 2025), SFHformer (Jiang et al., 2024), and DarkIR (Feijoo et al., 2025). The controlled studies then sharpen the comparison in two stages: the same-backbone study isolates the value of LTOR's evidence organization, and the fixed-LTOR study isolates the value of EBL's evidence-aligned refinement. Chain-level and internal ablations complete this progression under the same partitions and model-selection procedure.

EBLTR leads every PSNR, SSIM, and LPIPS column across Screen, Game, Lock, and Home in Appendix Table A1, providing the first quantitative evidence for screen-recovery feasibility. The agreement across pixel fidelity, structural similarity, and perceptual distance shows that the advantage extends from brightness recovery to regional organization and principal contours.

The breadth of the comparison strengthens the central claim within the capture arrangement. Across classical CNNs, general restoration architectures, and models designed for related degradations, the evidence-budgeted design turns the blurred wall response into a more faithful approximation of the hidden display, establishing ordinary-camera wall observation as a workable route to preliminary screen-content recovery.

Table A1: Complete results for EBLTR and the 19 baselines across the four content categories. Best in red, second best in blue.

Method	Source	Screen			Game			Lock			Home		
		PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Classical CNN Baselines													
DnCNN	TIP'17	8.88	0.3466	0.7245	8.74	0.4275	0.7258	11.82	0.4787	0.7221	9.02	0.3979	0.7900
U-Net	MICCAI'15	9.18	0.2401	0.8061	10.29	0.4715	0.7581	12.18	0.4874	0.7618	8.73	0.3726	0.8223
General-Purpose and Multi-Task Image Restoration Frameworks													
AdaIR	ICLR'25	21.71	0.7621	0.3873	17.44	0.6394	0.5754	15.17	0.5729	0.6455	13.96	0.5338	0.7324
NAFNet	ECCV'22	21.62	0.7585	0.3794	18.16	0.6587	0.5759	15.24	0.5619	0.6637	14.58	0.5129	0.7203
Restormer	CVPR'22	22.60	0.7774	0.3449	17.37	0.6438	0.5676	16.75	0.6071	0.6329	12.80	0.5036	0.7451
Uformer	CVPR'22	19.20	0.7068	0.4639	17.55	0.6470	0.6142	15.15	0.5652	0.6743	14.34	0.5356	0.7392
MIRNetv2	TPAMI'23	18.24	0.6917	0.4923	17.68	0.6559	0.6120	15.21	0.6042	0.6575	12.79	0.5342	0.7598
AST	CVPR'24	19.39	0.7050	0.4668	17.46	0.6453	0.6158	15.07	0.5366	0.6920	12.90	0.4960	0.7672
ConvIR	TPAMI'24	20.27	0.7283	0.4495	17.64	0.6515	0.6013	14.91	0.5925	0.6867	12.67	0.5168	0.7872
DeepSN-Net	TPAMI'25	10.49	0.2954	0.7787	10.87	0.4440	0.7390	12.84	0.4228	0.7711	10.67	0.3751	0.8343
SFHformer	ECCV'24	8.82	0.3940	0.8066	10.01	0.4620	0.7598	12.46	0.5095	0.7404	10.03	0.4480	0.8530
Models for Related Degradations													
DocRes	CVPR'24	21.15	0.7490	0.3999	17.53	0.6472	0.5795	16.36	0.6094	0.6467	12.69	0.5153	0.7378
Retinexformer	ICCV'23	19.57	0.6732	0.4554	17.93	0.6556	0.5924	14.60	0.5586	0.6506	12.66	0.5078	0.7343
HVI-CIDNet	CVPR'25	19.69	0.7178	0.4724	17.77	0.6555	0.5979	15.61	0.5758	0.6714	13.85	0.5331	0.7420
Ada-Deblur	TIP'23	19.43	0.7067	0.4637	17.75	0.6577	0.5878	16.27	0.6086	0.6608	14.06	0.5358	0.7522
DEA-Net	TIP'24	20.65	0.7500	0.4240	17.73	0.6579	0.5879	15.96	0.6069	0.6742	13.69	0.5334	0.7610
DehazeFormer	TIP'23	19.37	0.6931	0.4794	17.77	0.6494	0.6074	14.87	0.5375	0.6687	12.79	0.4704	0.7511
Pixel2Pixel	TPAMI'25	10.32	0.4816	0.6637	11.14	0.4502	0.6917	12.41	0.5178	0.7303	10.42	0.3787	0.7669
DarkIR	CVPR'25	9.22	0.1749	0.8673	4.46	0.0409	0.9361	8.79	0.0949	0.8040	6.58	0.0665	0.8897
EBLTR	**Ours**	**24.27**	**0.8189**	**0.2913**	**18.71**	**0.6665**	**0.4865**	**17.16**	**0.6295**	**0.6299**	**14.69**	**0.5462**	**0.7077**

A.2 Detailed Method Formulation and Same-Backbone Representation Evidence

The detailed formulation exposes how EBLTR converts a diffuse wall response into a controlled reconstruction. It follows the same variables and operators as the main method: LTOR organizes the surviving wall evidence, EBL translates observability into an exact local correction budget, and the frozen screen-to-wall surrogate evaluates the reconstruction through wall-domain consistency during training.

Effective transport model and observability structure

For each colour response, the captured wall image is represented phenomenologically as

$$x_c(r) = \sum_{c'=1}^{C} \int_{\Omega_y} k_{cc'}(r,s) y_{c'}(s)\,\mathrm{d}s + a_c(r) + n_c(r), \qquad k_{cc'}(r,s) \geq 0.$$

Here $r \in \Omega_x$, $s \in \Omega_y$, and $c, c' \in \{1,\dots,C\}$; $y_{c'}(s)$ is the hidden screen signal and $x_c(r)$ is the measured wall response. The effective kernel $k_{cc'}$ combines spatial propagation, diffuse reflection, cross-colour response, and camera integration, while a_c and n_c denote latent ambient response and observation perturbation. After sampling, $\mathbf{x} \in [0,1]^{B\times C\times H\times W}$, where B, C, H, and W are the batch, colour, height, and width dimensions. Inference derives every quantity from $\mathbf{x}$, making this effective transport abstraction the starting point for evidence-aware recovery of the hidden $\mathbf{y}$.

For a square-integrable effective kernel on finite acquisition domains, the induced operator $\mathcal{S}_\star$ is Hilbert–Schmidt and therefore compact. Its nonzero singular modes satisfy

$$\mathcal{S}_\star \mathbf{v}_\nu = \varsigma_\nu \mathbf{u}_\nu, \qquad \mathcal{S}_\star^* \mathbf{u}_\nu = \varsigma_\nu \mathbf{v}_\nu,$$

$$\langle \mathbf{x} - \mathbf{a}, \mathbf{u}_\nu \rangle_x = \varsigma_\nu \langle \mathbf{y}, \mathbf{v}_\nu \rangle_y + \langle \mathbf{n}, \mathbf{u}_\nu \rangle_x, \quad |\mathrm{d}\langle \mathbf{y}, \mathbf{v}_\nu \rangle_y| = \frac{|\langle \mathrm{d}\mathbf{n}, \mathbf{u}_\nu \rangle_x|}{\varsigma_\nu}.$$

The index ν denotes a transport mode; $\mathbf{u}_\nu \in L^2(\Omega_x;\mathbb{R}^C)$ and $\mathbf{v}_\nu \in L^2(\Omega_y;\mathbb{R}^C)$ are observation and screen basis functions, $\varsigma_\nu > 0$ is the corresponding singular value, and $\mathcal{S}_\star^*$ is the Hilbert-space

adjoint. The subscripts on $\langle \cdot, \cdot \rangle$ identify the observation and screen inner products. Small values of ς_ν amplify observation perturbations, so the strength of recoverable evidence varies across transport modes. This variation supplies the mathematical motivation for LTOR's wall-derived observability ordering and EBL's evidence-dependent refinement budget.

The inference map used by EBLTR is

$$\begin{aligned} \mathbf{A} &= \mathcal{F}_\vartheta^0(\mathbf{x}), & \mathbf{f}_j &= \mathcal{F}_\vartheta^j(\mathbf{x}), \qquad j \in \{s, m, d\}, \\ (\mathbf{z}, \mathbf{O}, \mathbf{x}') &= \mathcal{M}_\theta(\mathbf{x}; \mathbf{f}_s, \mathbf{f}_m, \mathbf{f}_d), & \mathbf{C} &= \mathcal{F}_\vartheta^0(\mathbf{x}'), \\ \widehat{\mathbf{Y}} &= \mathbf{C} + \Delta\mathbf{Y}. \end{aligned}$$

The two appearances of $\mathcal{F}_\vartheta^0$ share the same parameters ϑ. The feature fields have domains $\mathbf{f}_j \in \mathbb{R}^{B\times D_j\times H\times W}$; D_s, D_m, and D_d denote the channel dimensions of the shallow, middle, and deep features. LTOR returns $\mathbf{z} : \Omega_x \to \mathbb{E}$, $\mathbf{O} \in (0,1)^{B\times 1\times H\times W}$, and $\mathbf{x}' \in \mathbb{R}^{B\times C\times H\times W}$. The anchor $\mathbf{A}$, conservative estimate $\mathbf{C}$, correction $\Delta\mathbf{Y}$, and output $\widehat{\mathbf{Y}}$ lie in $\mathbb{R}^{B\times C\times H\times W}$, with $\Delta_\psi(\mathbf{x}, \mathbf{A}, \mathbf{C}, \mathbf{O}) = \Delta\mathbf{Y}$. LTOR has parameters θ, EBL has parameters ψ, and all learned parameters are fixed during inference.

In the implementation, EBL aligns the wall tensor, anchor, and observability map to the spatial size of $\mathbf{C}$ by bilinear interpolation with `align_corners=False`; tensors already at that size pass through unchanged. The aligned tensors supply the common computational grid for local refinement.

Normalized LTOR scale space

Let $\mathbf{x}^0$ be the zero extension of $\mathbf{x}$ from Ω_x to $\mathbb{R}^2$, and let ω_τ be a nonnegative normalized smoothing kernel. The continuous analysis form is

$$\begin{aligned} (\Omega_\tau \mathbf{x}^0)_c(r) &= \int_{\mathbb{R}^2} \omega_\tau(r-s) x_c^0(s)\,\mathrm{d}s, \\ \omega_\tau(s) \geq 0, &\qquad \int_{\mathbb{R}^2} \omega_\tau(s)\,\mathrm{d}s = 1, \\ \mathbf{b}_l^0 &= \Omega_{\tau_1}\mathbf{x}^0, \\ \mathbf{b}_m^0 &= (\Omega_{\tau_2} - \Omega_{\tau_1})\mathbf{x}^0, \\ \mathbf{b}_h^0 &= (I - \Omega_{\tau_2})\mathbf{x}^0. \end{aligned}$$

The corresponding transfer functions are

$$\begin{aligned} \widehat{\mathbf{b}_l^0}(\xi) &= \widehat{\omega}_{\tau_1}(\xi)\widehat{\mathbf{x}^0}(\xi), \\ \widehat{\mathbf{b}_m^0}(\xi) &= (\widehat{\omega}_{\tau_2}(\xi) - \widehat{\omega}_{\tau_1}(\xi))\,\widehat{\mathbf{x}^0}(\xi), \\ \widehat{\mathbf{b}_h^0}(\xi) &= (1 - \widehat{\omega}_{\tau_2}(\xi))\,\widehat{\mathbf{x}^0}(\xi). \end{aligned}$$

Restricting the extended fields to Ω_x gives

$$\mathbf{b}_j = \mathbf{b}_j^0\big|_{\Omega_x}, \qquad j \in \{l, m, h\}, \qquad \mathbf{b}_l + \mathbf{b}_m + \mathbf{b}_h = \mathbf{x}.$$

The fixed scales obey the same broad-to-narrow support ordering used in the main paper,

$$\mathrm{supp}(\omega_{\tau_2}) \subsetneq \mathrm{supp}(\omega_{\tau_1}).$$

The three transfer functions sum algebraically to one, preserving the full wall observation while exposing broad, intermediate, and residual evidence to separate analysis. Their overlapping signed responses form an exact decomposition, and the sampled realization follows the same zero-boundary convention and finite-support normalized sums.

The deterministic local statistics are

$$
\begin{aligned}
q(r) &= \frac{1}{C}\sum_{c=1}^{C} x_c(r), \\
g(r) &= \left(\sum_{\iota=1}^{2} |\partial_\iota q(r)|^2 + \epsilon\right)^{1/2}, \\
v(r) &= \left[\Omega_{\tau_v}(q^2)(r) - (\Omega_{\tau_v} q(r))^2\right]_+ + \eta\, \Omega_{\tau_v}\left(\frac{1}{C}\sum_{c=1}^{C} |b_{h,c}|\right)(r).
\end{aligned}
$$

Here q, g, and v have shape $B \times 1 \times H \times W$ and represent the colour mean, gradient magnitude, and nonnegative local variation. The derivative ∂_ι acts in spatial direction ι, $[\cdot]_+$ is the positive part, and $\epsilon > 0$, $\eta \geq 0$, and τ_v are fixed. Whenever Ω_{τ_v} acts on a finite-domain field, the same deterministic zero extension is used.

The full LTOR product representation is

$$
\begin{aligned}
\mathbf{z}(r) &= \big(\mathbf{x}(r), \mathbf{b}_l(r), \mathbf{b}_m(r), \mathbf{b}_h(r), g(r), v(r), \mathbf{f}_s(r), \mathbf{f}_m(r), \mathbf{f}_d(r)\big) \in \mathbb{E}, \\
\mathbb{E} &= \mathbb{R}^C \times (\mathbb{R}^C)^3 \times \mathbb{R}^2 \times \mathbb{R}^{D_s} \times \mathbb{R}^{D_m} \times \mathbb{R}^{D_d}.
\end{aligned}
$$

Every component of $\mathbf{z}$ is computed from $\mathbf{x}$ and $\mathbf{b}_h$, so LTOR obtains its full evidence field directly from the captured wall response at inference. The deterministic statistics expose brightness, variation, and frequency structure, while the learned feature coordinates add shallow, middle, and deep context; together they organize the wall traces that support hidden-screen recovery.

Band priors and contextual observability

For each $j \in \{l, m, h\}$, the colour-averaged magnitude, spatial mean, and population standard deviation are

$$
\begin{aligned}
e_j(r) &= \frac{1}{C}\sum_{c=1}^{C} |b_{j,c}(r)|, & \mu_j &= \frac{1}{|\Omega_x|}\int_{\Omega_x} e_j(r)\,\mathrm{d}r, \\
s_j &= \left(\frac{1}{|\Omega_x|}\int_{\Omega_x} (e_j(r) - \mu_j)^2\,\mathrm{d}r\right)^{1/2}, & n_j(r) &= \frac{e_j(r) - \mu_j}{\max(s_j, \epsilon)}.
\end{aligned}
$$

The deterministic prior is

$$
\ell_j^0(r) = \kappa n_j(r) + \rho_j, \qquad p_j(r) = \sigma(\ell_j^0(r)), \qquad \kappa > 0, \qquad \rho_l > \rho_m > \rho_h.
$$

Here e_j is the colour-averaged band magnitude, μ_j and s_j are its per-sample spatial mean and population standard deviation, and n_j is the stabilized standardized field. The positive slope κ preserves evidence ordering, while the ordered offsets give broad colour and layout evidence priority over finer responses at equal standardized energy. The logistic output p_j is therefore a deterministic wall-derived ordering cue that initiates LTOR's spatial–scale–frequency observability representation.

Context corrects the deterministic prior through

$$
\begin{aligned}
\mathbf{u}(r) &= \Phi_\theta\left(\mathbf{z}(r), \int_{\Omega_x} \gamma_\theta(r, s)[\mathbf{z}(s)]\,\mathrm{d}s\right), \\
\gamma_\theta(r, s) &= 0 \qquad \text{for } s \notin \Omega_r, \\
\mathbf{O}_j(r) &= \sigma\big(\mathrm{logit}(p_j(r)) + \Psi_j^\theta[\mathbf{u}](r)\big), \\
\mathbf{O}(r) &= \frac{1}{3}\sum_{j \in \{l,m,h\}} \mathbf{O}_j(r).
\end{aligned}
$$

Here Ω_r is the local support of the learned interaction, Φ_θ produces the contextual field $\mathbf{u}$, and Ψ_j^θ produces a scalar correction for band j. The learned coupling γ_θ integrates neighbouring representation evidence, while the physical kernel $k_{cc'}$ describes the acquisition process. The bounded scores

$\mathbf{O}_j$ and $\mathbf{O}$ order the evidence available for recovery, and $\{l, m, h\}$ indexes the complete three-band decomposition used by LTOR.

The conditioned observation is obtained from

$$\begin{aligned}
\widetilde{\mathbf{b}}_j &= \mathbf{O}_j \otimes \mathbf{b}_j, \qquad j \in \{l, m, h\}, \\
\mathbf{r}_x &= \Gamma_\theta(\mathbf{x}, \widetilde{\mathbf{b}}_l, \widetilde{\mathbf{b}}_m, \widetilde{\mathbf{b}}_h, g, v, \mathbf{O}_l, \mathbf{O}_m, \mathbf{O}_h, \mathbf{O}, \mathbf{f}_s, \mathbf{f}_m, \mathbf{f}_d), \\
\mathbf{d}_x &= \delta(\beta + \chi\mathbf{O}) \otimes \tanh(\mathbf{r}_x), \\
\mathbf{x}' &= \mathbf{x} + \mathbf{d}_x, \\
\beta \geq 0, &\qquad \chi > 0, \qquad \beta + \chi = 1.
\end{aligned}$$

Since $|\tanh(t)| \leq 1$ and $0 < \beta + \chi O_{b,1,i,j} \leq 1$, the wall-domain adjustment obeys

$$\|\mathbf{d}_x(r)\|_\infty \leq \delta.$$

Thus, LTOR first weights each band by observability and then applies a bounded adjustment before computing $\mathbf{C} = \mathcal{F}^0_\vartheta(\mathbf{x}')$. The conditioned internal wall field $\mathbf{x}'$ carries the selected brightness, layout, and boundary evidence into the conservative screen estimate while remaining anchored to the measured $\mathbf{x}$.

EBL FEASIBLE SET AND EXACT BOUNDS

EBL constructs its local evidence field and unconstrained proposal as

$$\begin{aligned}
\mathbf{z}_y(r) &= \big(\mathbf{x}(r), \mathbf{C}(r), \mathbf{A}(r), \mathbf{C}(r) - \mathbf{A}(r), \mathbf{O}(r)\big) \in \mathbb{E}_y, \\
\mathbb{E}_y &= \mathbb{R}^C \times (\mathbb{R}^C)^3 \times \mathbb{R}, \\
\mathbf{D}(r) &= \Phi_\psi\left(\mathbf{z}_y(r), \int_{\Omega_r} \gamma_\psi(r, s)[\mathbf{z}_y(s)]\, \mathrm{d}s\right).
\end{aligned}$$

All arguments of $\mathbf{D}$ are inference-visible functions of $\mathbf{x}$. The proposal gathers the measured wall response, the anchor, the conservative estimate, their difference, and observability into a local candidate correction whose admissible magnitude is then set by evidence.

The evidence-dependent feasible set is

$$\mathbb{K}(\mathbf{O}; \mathbf{C}) = \left\{\mathbf{Y} \in \mathbb{R}^{B\times C\times H\times W} : |Y_{b,c,i,j} - C_{b,c,i,j}| \leq \alpha O_{b,1,i,j} \text{ for all } (b, c, i, j)\right\}.$$

Here b, c, i, and j index batch, colour, row, and column, respectively. The fixed $\alpha > 0$ is the maximum correction scale, and the scalar $O_{b,1,i,j}$ is broadcast over colour. Centering the feasible set at $\mathbf{C}$ gives EBL its central mechanism: wall-derived observability sets the local distance that refinement may travel from the conservative reconstruction.

The smooth bounded parameterization is

$$\Delta\mathbf{Y} = \alpha\mathbf{O} \otimes \tanh(\mathbf{D}/\alpha), \qquad \widehat{\mathbf{Y}} = \mathbf{C} + \Delta\mathbf{Y}.$$

It gives the entrywise and norm bounds below, together with the proposal-sensitivity bound for fixed observability $\mathbf{O}$

$$\begin{aligned}
|\Delta Y_{b,c,i,j}| &\leq \alpha O_{b,1,i,j} \leq \alpha, \\
\|\Delta\mathbf{Y}\|_p &\leq \alpha\|\mathbf{O} \otimes \mathbf{1}_C\|_p, \qquad 1 \leq p \leq \infty, \\
0 \leq \left|\frac{\partial \Delta Y_{b,c,i,j}}{\partial D_{b,c,i,j}}\right| &= O_{b,1,i,j}\, \mathrm{sech}^2(D_{b,c,i,j}/\alpha) \leq O_{b,1,i,j}.
\end{aligned}$$

The first inequality follows directly from $|\tanh(t)| \leq 1$; applying the same entrywise bound yields the stated p-norm inequality, and differentiating the smooth parameterization gives the sensitivity bound. Consequently, $\widehat{\mathbf{Y}} \in \mathbb{K}(\mathbf{O}; \mathbf{C})$. These exact properties make the evidence budget operational: stronger wall evidence permits a larger, more responsive correction, while weaker evidence keeps the output close to the conservative estimate. The training-time observation check then evaluates the resulting reconstruction through the wall domain.

TRAINING-ONLY OBSERVATION CHECK

The frozen surrogate maps a recovered screen back to the wall domain through a diffuse base response and bounded discrepancy:

$$\begin{aligned}\mathbf{g}_\phi &= \upsilon_0 \sigma(\boldsymbol{\zeta}_\phi),\\ \mathbf{s}_0 &= \mathbf{g}_\phi \otimes \Omega_{\tau_s}(\widehat{\mathbf{Y}}) + \mathbf{b}_\phi,\\ \mathcal{S}_\phi(\widehat{\mathbf{Y}}) &= \Pi_{[0,1]}\Big(\mathbf{s}_0 + \delta_2 \tanh\Big(\Gamma_\phi(\widehat{\mathbf{Y}}, \mathbf{s}_0, \nabla\widehat{\mathbf{Y}})\Big)\Big).\end{aligned}$$

Here $\boldsymbol{\zeta}_\phi, \mathbf{b}_\phi \in \mathbb{R}^{1\times C\times 1\times 1}$, $\upsilon_0 > 0$ is fixed, and $\mathbf{g}_\phi \in \mathbb{R}_+^{1\times C\times 1\times 1}$ is a positive colour gain. The operator Ω_{τ_s} supplies the fixed diffuse smoothing response, $\mathbf{s}_0$ is the resulting base field, $\delta_2 > 0$ bounds the discrepancy produced by Γ_ϕ, and $\Pi_{[0,1]}$ projects onto the intensity domain. The auxiliary wall response and discrepancy are

$$\widetilde{\mathbf{X}} = \mathcal{S}_\phi(\widehat{\mathbf{Y}}), \qquad L_1 = d(\widetilde{\mathbf{X}}, \mathbf{x}).$$

During training, the frozen screen-to-wall surrogate predicts the wall response of the recovered screen and compares it with the captured wall image through L_1. This connects the reconstructed content back to the observation that supports it. The surrogate parameters ϕ stay frozen during EBLTR training. At inference, EBLTR uses $\mathbf{x}$, its derived quantities, and the fixed parameters ϑ, θ, and ψ, so ordinary-camera wall capture remains the complete input to recovery.

The implemented discrepancy is an observability-weighted absolute error,

$$d(\widetilde{\mathbf{X}}, \mathbf{x}) = \frac{\sum_{b,c,i,j} w_{b,i,j} |\widetilde{X}_{b,c,i,j} - x_{b,c,i,j}|}{C \sum_{b,i,j} w_{b,i,j} + 10^{-8}}, \qquad w = m\,(0.5 + 0.5\ \mathrm{stopgrad}(\mathbf{O}))\,,$$

where m is the valid-image mask. The reconstruction is clipped to $[0, 1]$ before entering the surrogate, and its predicted wall response is bilinearly aligned to the observed wall size when needed. The surrogate is loaded, placed in evaluation mode, and frozen before the reconstruction optimizer is created; gradients through its input update the reconstruction network.

IMPLEMENTATION DETAILS

For the complete EBLTR model, the reconstruction backbone is a six-stage encoder–decoder with encoder channels $(64, 128, 256, 512, 512, 512)$ and a mirrored decoder channel configuration. The backbone incorporates learned path compensation and learnable nearest-neighbour upsampling in the final two decoder stages. Two additional decoder blocks map 128 to 64 channels with a branch width of 16, and 256 to 128 channels with a maximum RMS residual ratio of 0.02, respectively. LTOR and EBL augment this reconstruction pathway.

The 384×512 input canvas is resampled to a 256×256 working grid for the two backbone passes and LTOR. Reconstruction tensors return to the input canvas by bicubic interpolation, and observability returns by bilinear interpolation before EBL refinement. The reconstruction model has 32,745,071 trainable parameters. Its final training stage initializes from the corresponding content-specific predecessor with the same backbone, LTOR, EBL, and wall-observation consistency (WOC), and uses AdamW with batch size 8, learning rate 4×10^{-6}, weight decay 10^{-4}, gradient-norm clipping at 1, and random seed 42. Training runs for at most 60 epochs, with early stopping after 10 epochs without a training-loss improvement of 10^{-4}.

A frozen copy of the initialization model provides the reference predictions for the structural losses. A separate frozen backbone predecessor without LTOR or EBL calibrates the band observability scores. These two references and the screen-to-wall surrogate remain frozen; all 32,745,071 parameters of the reconstruction model are trainable. The complete model's final-stage objective is

$$\begin{aligned}\mathcal{L}_{\mathrm{base}} &= \mathcal{L}_{\mathrm{pix}} + \lambda_{\mathrm{lap}}(t)\mathcal{L}_{\mathrm{lap}} + 0.02\mathcal{L}_{\mathrm{LPIPS}},\\ \mathcal{L}_{\mathrm{train}} &= \mathcal{L}_{\mathrm{base}} + \lambda_s^{\mathrm{eff}}(t, g)\mathcal{L}_{\mathrm{comp}} + \lambda_n(t)\mathcal{L}_{\mathrm{nr}} + \lambda_l(t)\mathcal{L}_{\mathrm{L1g}}\\ &\quad + \lambda_c(t)\mathcal{L}_{\mathrm{cg}} + 0.05\mathcal{L}_{\mathrm{LTOR}} + 0.005\mathcal{L}_{\mathrm{WOC}}.\end{aligned}$$

Here $\mathcal{L}_{\mathrm{pix}}$ is masked RGB absolute error, $\mathcal{L}_{\mathrm{lap}}$ is the corresponding absolute error after the four-neighbour Laplacian, and $\mathcal{L}_{\mathrm{LPIPS}}$ uses AlexNet features on valid-region crops bilinearly resized

to 224×224. The term $\mathcal{L}_{\mathrm{WOC}}$ is the observability-weighted wall discrepancy d defined above. Training progress is $t = ((e-1) + s/S)/60$, clipped to $[0, 1]$, for epoch e, step s, and S steps per epoch. The Laplacian weight is $\lambda_{\mathrm{lap}}(t) = 0.05(1 - 0.75t)$. The tuple $(\lambda_s, \lambda_n, \lambda_l, \lambda_c)$ is $(0.018, 0.090, 0.035, 0.045)$ through $t = 0.10$, interpolates linearly to $(0.038, 0.120, 0.045, 0.060)$ at $t = 0.55$, and remains there. Only the component weight is adjusted by gradient agreement: $\lambda_s^{\mathrm{eff}} = \lambda_s \,\mathrm{clip}((g + 0.010)/0.050, 0, 1)$, where g is the cosine between the output gradients of $\mathcal{L}_{\mathrm{base}}$ and $\mathcal{L}_{\mathrm{comp}}$.

The component loss $\mathcal{L}_{\mathrm{comp}}$ combines local RGB SSIM luminance, contrast, and structure deficits with weights $(0.25, 0.30, 0.45)$, using windows of widths $(7, 11, 21)$ with weights $(0.30, 0.45, 0.25)$. Its spatial weight combines the valid mask, target structure, detached observability, and the frozen reference's local SSIM error. The no-regression loss $\mathcal{L}_{\mathrm{nr}}$ uses width-11 SSIM, luminance, contrast, and structure similarities: for each, it penalizes $[a_{\mathrm{ref}} - a_{\mathrm{pred}} + \mu]_+$, with respective weights $(0.40, 0.15, 0.20, 0.25)$ and margins $(0.0015, 0.0010, 0.0010, 0.0015)$. The losses $\mathcal{L}_{\mathrm{L1g}}$ and $\mathcal{L}_{\mathrm{cg}}$ penalize increases beyond 0.0002 over the frozen reference's width-11 local RGB absolute error and width-9 local luminance-standard-deviation error, respectively. The no-regression and L1 guards are averaged over the valid mask; the contrast guard uses the component loss's spatial weights.

LTOR calibration averages masked Brier losses over the three bands. Its detached target for band j is

$$T_j = \exp\left[-\frac{1}{0.05C}\sum_{c=1}^{C}\sqrt{(b_{j,c}^{\mathrm{teacher}} - b_{j,c}^{\mathrm{target}})^2 + 10^{-6}}\right],$$

where the teacher reconstruction and target screen are aligned to the observability grid before band decomposition. The loss compares $\mathbf{O}_j$ with T_j under the valid mask. This calibration reference determines band reliability, while the separate initialization-model reference supplies the structural comparisons.

The LTOR panels in Figure A7 emphasize transitions in the wall observation, while the later EBL panels respond around reconstructed screen regions. The visualization connects the organization of wall evidence to the locations where screen refinement occurs.

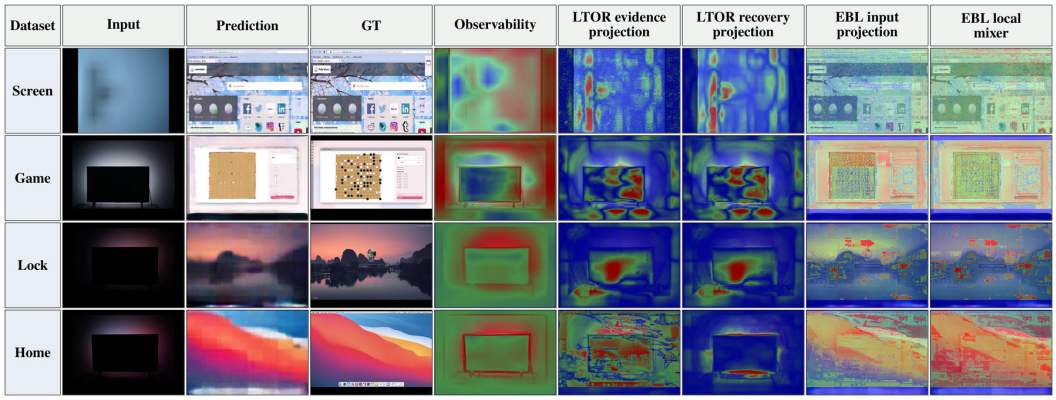

Figure A7: Observability and feature attribution for representative Screen, Game, Lock, and Home examples.

The observability overlay shows the channel-mean LTOR output. The four attribution panels show the LTOR evidence projection, LTOR recovery projection, EBL input projection, and EBL local mixer. Attribution is the channel sum of absolute activation–gradient products for negative reconstruction L1 error on the valid target region. Each displayed map is shifted by its minimum, divided by its 99th percentile, and clipped to $[0, 1]$. LTOR attribution is overlaid on the wall input and EBL attribution on the reconstruction.

The four content categories reveal the same mechanism through different structures: page regions in Screen, board geometry in Game, broad bands in Lock, and scene organization in Home. This repeated progression connects the internal representation directly to the target claim that a diffuse wall spot retains enough structured evidence for preliminary recovery of brightness, layout, and principal outlines.

How should the surviving wall evidence be represented before detail budgeting? Appendix Table A2 isolates this first step in EBLTR. The shared backbone, decoder, learning rule, and partitions turn the table into a controlled comparison of evidence organization.

Table A2: Complete same-backbone representation comparison on all four content categories. Red and blue indicate the best and second-best distinct displayed values; ties share the same colour.

Representation	Learning	Screen			Game			Lock			Home		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Uniform	Plain learning	23.57	0.8039	0.3146	18.75	0.6658	0.4912	16.96	0.6215	0.6307	14.57	0.5373	0.7054
CNN-Pyramid	Plain learning	23.67	0.8045	0.3137	18.75	0.6656	0.4913	16.96	0.6213	0.6309	14.54	0.5381	0.7046
Window-Attention	Plain learning	23.65	0.8039	0.3146	18.75	0.6657	0.4913	16.96	0.6215	0.6307	14.58	0.5375	0.7053
Frequency-Prompt	Plain learning	23.62	0.8041	0.3143	18.75	0.6657	0.4913	16.96	0.6215	0.6308	14.55	0.5381	0.7047
LTOR	**Plain learning**	**23.90**	**0.8096**	**0.3045**	**18.76**	**0.6671**	**0.4864**	**16.97**	**0.6238**	**0.6304**	**14.59**	**0.5387**	**0.7041**

LTOR leads all twelve content–metric columns in Appendix Table A2. Its advantage across Screen, Game, Lock, and Home shows that explicitly arranging spatial, multilevel, and frequency evidence provides a stronger basis for recovering hidden-screen structure than uniform, pyramid, window-attention, or frequency-prompt representations under the shared reconstruction system. The controlled comparison identifies how evidence organization contributes to hidden-screen recovery. The wall spot carries heterogeneous evidence; LTOR turns that heterogeneity into a useful observability ordering, and the consistent PSNR, SSIM, and LPIPS direction shows that this ordering benefits brightness fidelity, regional structure, and perceptual agreement together.

A.3 Overall Chain-Level Ablation

Appendix Table A3 tests the full progression from observing evidence to spending detail budget. Its compound variants replace the corresponding LTOR or EBL chain together with the associated training components, preserving each row as a complete alternative reconstruction route.

The complete EBLTR chain leads all twelve content–metric columns in Appendix Table A3. Constant, shuffled, flattened, and plain-residual alternatives each weaken the link between observed wall evidence and reconstructed detail, while the complete design preserves that link from LTOR through EBL.

Table A3: Complete overall ablation of the LTOR and EBL chains. Best in red, second best in blue.

Method	Screen			Game			Lock			Home		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
LTOR Overall Ablation and Replacement												
w/o LTOR	24.06	0.8139	0.3175	18.40	0.6633	0.5463	17.05	0.6288	0.6361	14.47	0.5425	0.7362
LTOR→Constant Map	23.42	0.8043	0.3321	18.26	0.6653	0.5596	16.85	0.6293	0.6323	14.56	0.5455	0.7362
LTOR→Band-Shuffled Evidence	23.80	0.8108	0.3250	18.28	0.6638	0.5588	16.85	0.6212	0.6347	14.62	0.5435	0.7390
LTOR →Flat Evidence	23.61	0.8096	0.3259	18.30	0.6646	0.5552	16.92	0.6272	0.6334	14.60	0.5445	0.7387
EBL Overall Ablation and Replacement												
w/o EBL	23.72	0.8098	0.3288	18.26	0.6657	0.5600	16.92	0.6289	0.6351	14.54	0.5411	0.7346
EBL →Constant Low Budget	23.63	0.8107	0.3257	18.30	0.6623	0.5581	16.76	0.6263	0.6405	14.59	0.5454	0.7375
EBL →Shuffled Low Budget	23.58	0.8069	0.3244	18.32	0.6651	0.5559	16.93	0.6292	0.6404	14.67	0.5447	0.7412
EBL →Plain Residual Refiner	23.72	0.8083	0.3264	17.74	0.6557	0.5689	17.01	0.6243	0.6436	14.42	0.5407	0.7354
EBLTR	**24.27**	**0.8189**	**0.2913**	**18.71**	**0.6665**	**0.4865**	**17.16**	**0.6295**	**0.6299**	**14.69**	**0.5462**	**0.7077**

This agreement across Screen, Game, Lock, and Home supports the complete chain from wall-evidence representation to budgeted screen refinement. Useful recovery begins by locating the evidence that survives wall transport and reaches its strongest form when the available refinement is allocated to those locations, explaining why the complete sequence yields the most faithful hidden-screen reconstruction.

A.4 LTOR Internal Ablation

Appendix Table A4 asks which ingredients let LTOR recognize useful wall evidence. The variants jointly examine multilevel backbone cues and the ordered band prior, including their associated calibration and training components under the shared evaluation procedure.

Table A4: Complete LTOR internal ablation across the four fixed test sets. Best in red, second best in blue.

Method	Screen			Game			Lock			Home		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
LTOR Backbone Evidence Variants												
LTOR w/o Backbone Evidence	23.39	0.8079	0.3268	18.39	0.6655	0.5622	16.92	0.6240	0.6421	14.45	0.5392	0.7364
LTOR Backbone Evidence →Shallow Only	23.95	0.8126	0.3207	18.23	0.6587	0.5602	16.66	0.6148	0.6400	14.63	0.5412	0.7398
LTOR Backbone Evidence →Middle Only	23.68	0.8072	0.3288	18.09	0.6567	0.5711	16.70	0.6178	0.6421	14.48	0.5406	0.7333
LTOR Backbone Evidence →Deep Only	23.88	0.8124	0.3206	18.45	0.6653	0.5521	17.09	0.6270	0.6408	14.30	0.5382	0.7331
LTOR Band Prior Variants												
LTOR w/o Band Prior	23.79	0.8099	0.3268	18.31	0.6648	0.5557	16.78	0.6208	0.6396	14.45	0.5380	0.7325
LTOR Band Prior →Luma Gradient Prior	23.77	0.8102	0.3277	18.38	0.6625	0.5547	17.03	0.6286	0.6396	14.52	0.5421	0.7347
LTOR Band Prior →Texture Energy Prior	23.56	0.8053	0.3293	18.37	0.6644	0.5510	16.93	0.6248	0.6360	14.52	0.5411	0.7364
LTOR Band Prior →Reversed Energy Prior	23.84	0.8091	0.3255	18.41	0.6643	0.5522	17.06	0.6283	0.6387	14.50	0.5418	0.7390
EBLTR	**24.27**	**0.8189**	**0.2913**	**18.71**	**0.6665**	**0.4865**	**17.16**	**0.6295**	**0.6299**	**14.69**	**0.5462**	**0.7077**

The complete LTOR design leads all twelve content–metric columns in Appendix Table A4. Multilevel backbone evidence supplies complementary receptive ranges for broad regions and local boundaries, while the ordered band prior reflects the transport pattern in which brightness and layout usually survive more strongly than fine contours. The single-level variants reveal that useful receptive ranges depend on screen structure: shallow evidence is strongest among the single-level alternatives on Screen in PSNR and SSIM, while deeper evidence benefits Game. LTOR brings these complementary ranges together with the ordered band prior. Broad regions and local transitions can therefore contribute to the same observability estimate, so refinement can respond to the different screen structures across all four content categories.

A.5 EBL Internal Ablation and Additional Visual Diagnostics

Table A5: Complete EBL internal ablation across the four fixed test sets. Best in red, second best in blue.

Method	Screen			Game			Lock			Home		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
EBL Budget and Structure Guard Variants												
EBL w/o Structure Guard	23.55	0.8054	0.3269	18.26	0.6613	0.5558	16.91	0.6256	0.6438	14.49	0.5374	0.7334
EBL Budget →Block-Shuffled Low Budget	23.59	0.8067	0.3311	18.33	0.6644	0.5546	17.01	0.6275	0.6420	14.42	0.5403	0.7334
EBL Budget →Reversed Tile Low Budget	23.68	0.8122	0.3223	18.12	0.6571	0.5746	17.03	0.6242	0.6429	14.58	0.5404	0.7403
EBL Budget →Anti-Evidence Low Budget	23.69	0.8112	0.3251	18.39	0.6646	0.5515	17.05	0.6277	0.6454	14.42	0.5408	0.7349
EBL Local Mixer Variants												
EBL w/o Local Mixer	23.63	0.8084	0.3255	18.04	0.6561	0.5638	16.96	0.6237	0.6449	14.51	0.5421	0.7361
EBL Local Mixer →Plain Conv Mixer	23.61	0.8104	0.3261	18.44	0.6639	0.5565	17.07	0.6293	0.6413	14.56	0.5396	0.7395
EBL Local Mixer →Channel MLP Mixer	23.82	0.8118	0.3225	18.07	0.6552	0.5625	17.02	0.6284	0.6412	14.46	0.5399	0.7348
EBL Local Mixer →Dilated Conv Mixer	23.56	0.8055	0.3333	18.42	0.6636	0.5535	16.97	0.6282	0.6379	14.46	0.5415	0.7361
EBLTR	**24.27**	**0.8189**	**0.2913**	**18.71**	**0.6665**	**0.4865**	**17.16**	**0.6295**	**0.6299**	**14.69**	**0.5462**	**0.7077**

Once LTOR has organized the wall evidence, Appendix Table A5 tests how EBL turns it into reconstructed detail. The structural-protection, budget-permutation, and local-mixer variants probe the content, placement, and neighbourhood use of the local refinement budget.

The complete EBL design leads all twelve columns in Appendix Table A5. The low-budget controls redistribute and halve the observability budget, while the complete model places refinement where LTOR finds usable screen evidence. The local mixer adds neighbourhood context, and structural protection keeps the correction organized around the conservative reconstruction.

The block-shuffled control averages observability over 8×8 tiles, circularly shifts the tile grid, and restores the original size by nearest-neighbour interpolation before multiplying by 0.5. The reversed-tile control reverses the flattened tile order and applies the same interpolation and scaling. The anti-evidence control uses $0.5(1 - \mathbf{O})$. Local-mixer replacements use a plain 3×3 convolution, a 1×1 channel MLP, or a dilation-2 depthwise convolution followed by a pointwise convolution; removing the mixer leaves the input projection and residual head in place.

Together, these variants establish that detail capacity and detail placement are coupled. EBL converts LTOR's evidence map into a local feasible set, so supported regions receive stronger contour and layout refinement while weakly observed regions remain close to their stable brightness and regional estimate. This is the operational advantage of evidence-budgeted learning over uniform sharpening.

Appendix Table A6 isolates evidence use after LTOR has already organized the wall response. With LTOR fixed, EBL achieves the best or joint-best reconstruction results across all four content categories. The comparison connects the benefit of observability representation to the way that evidence controls subsequent refinement. Anchor-only, plain-residual, constant-budget, and shuffled-budget rules provide the corresponding alternative uses of the same representation.

Table A6: Complete fixed-LTOR comparison of five detail-learning rules on all four test sets. Red and blue indicate the best and second-best distinct displayed values; ties share the same colour.

Representation	Learning Rule	Screen			Game			Lock			Home		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
LTOR	Anchor-only	23.78	0.8094	0.3046	18.75	0.6667	0.4863	17.05	0.6242	0.6315	14.75	0.5493	0.7092
LTOR	Plain residual	23.76	0.8095	0.3046	18.73	0.6660	0.4888	17.10	0.6261	0.6314	14.75	0.5474	0.7100
LTOR	Constant budget	23.80	0.8092	0.3055	18.73	0.6660	0.4888	17.05	0.6240	0.6315	14.75	0.5474	0.7101
LTOR	Shuffled budget	23.71	0.8091	0.3057	18.73	0.6660	0.4888	17.05	0.6240	0.6315	14.75	0.5473	0.7101
LTOR	**EBL**	**23.82**	**0.8097**	**0.3045**	**18.76**	**0.6668**	**0.4861**	**17.18**	**0.6293**	**0.6314**	**14.75**	**0.5494**	**0.7091**

With LTOR fixed, Table A6 isolates how the learning rule uses the same observability estimate to control refinement. LTOR answers where the wall spot still carries brightness, layout, and contour information; EBL answers how much refinement those locations receive. The consistent direction across Screen, Game, Lock, and Home shows that spending detail capacity according to observability turns the same representation into a more faithful hidden-screen reconstruction.

B Expanded Evidence for Wall-Reflected Screen Recovery

The additional experiments examine why hidden-screen recovery improves when refinement follows the evidence in a diffuse wall observation. We connect spatial influence, reconstruction risk, local boundary errors, and refinement allocation to the evidence used by EBLTR.

B.1 Mechanistic Evidence for Evidence Representation and Budgeting

We examine whether improved reconstruction is accompanied by fewer erroneous boundaries, with additional diagnostics in regions of low observability. Table B7 connects representation choice directly to reconstruction behaviour on Screen. LTOR combines the strongest PSNR, SSIM, and LPIPS with the lowest false-edge quantities, linking its global fidelity advantage to a concrete local effect: fewer reconstructed screen boundaries disagree with the target.

The boundary diagnostics use luminance $q_{\mathrm{lum}} = 0.299R + 0.587G + 0.114B$ and the magnitude of horizontal and vertical first differences. Predicted and target edges are pixels at or above the respective image's 85th gradient percentile. The false-edge rate is the fraction of predicted edge pixels absent from the target edge set. Low- and high-observability regions are defined by the bottom and top 40% of a fixed wall proxy: $N(0.55|\nabla q_{\mathrm{lum},x}| + 0.45N(|\nabla^2 q_{\mathrm{lum},x}|))$, where N linearly rescales the 1st–99th percentile interval to $[0, 1]$ and clips the result. The Laplacian uses a

four-neighbour stencil. Regional gradient error averages the absolute difference between predicted and target gradient magnitudes; the regional false-edge denominator counts predicted edges in that region. Each diagnostic is computed in the valid image region and averaged over test samples.

Table B7: Screen diagnostics for evidence representations. Image quality and regional boundary fidelity are evaluated together. Red and blue indicate the best and second-best distinct displayed values; ties share the same colour.

Representation	Image quality			All pixels	Low observability		High observability
	PSNR↑	SSIM↑	LPIPS↓	False-edge rate↓	Gradient error↓	False-edge rate↓	Gradient error↓
Uniform	23.57	0.8039	0.3146	0.4143	0.0186	0.4241	0.0244
CNN-Pyramid	23.67	0.8045	0.3137	0.4137	0.0186	0.4233	0.0244
Window-Attention	23.65	0.8039	0.3146	0.4144	0.0186	0.4238	0.0244
Frequency-Prompt	23.62	0.8041	0.3143	0.4139	0.0186	0.4234	0.0244
LTOR	**23.90**	**0.8096**	**0.3045**	**0.4036**	**0.0183**	**0.4118**	**0.0240**

This result sharpens the meaning of observability representation. LTOR organizes wall responses by their value for screen structure, so page layout and contour refinement follow transport evidence while spurious local edges receive a smaller role. The improvement appears jointly in image quality and boundary fidelity: LTOR produces a closer screen reconstruction with fewer erroneous local edges.

Table B8: Recovery-risk ordering of the observability representations on Screen. An em dash denotes an undefined correlation for a constant score.

Representation	Top-50% risk reduction↑	Top-20% risk reduction↑	Low–high risk gap↑	Spearman correlation with negative reconstruction error↑
Uniform	0.0000	0.0000	0.0000	—
CNN-Pyramid	0.1820	0.2785	0.0091	0.0380
Window-Attention	0.1856	0.2900	0.0094	0.0341
Frequency-Prompt	-0.2009	-0.3282	-0.0092	-0.0730
LTOR	**0.2159**	**0.3585**	**0.0112**	**0.1169**

Table B8 tests whether higher observability identifies reconstruction regions with lower recovery risk. LTOR achieves the strongest risk-coverage gains and low–high risk separation among the compared representations, with top-50% and top-20% risk reductions of 0.2159 and 0.3585. This ordering supplies the decision variable needed by EBL: where additional refinement has stronger observational support.

Here risk is the channel-mean absolute reconstruction error per pixel. For the highest-scoring 50% or 20% of pixels, risk reduction is one minus their mean risk divided by the full-image mean risk; the low–high gap is the lower-half mean risk minus the upper-half mean risk. The correlation compares the observability score with negative reconstruction error. These quantities are evaluated per image and then averaged. Uniform assigns a constant score and therefore has zero risk separation and an undefined rank correlation.

The representation therefore connects wall evidence to a concrete reconstruction decision. Once the wall image has been arranged by spatially varying recovery evidence, EBL can allocate more correction to informative regions and preserve the conservative layout where evidence is weak. The table measures the representation property that makes evidence budgeting possible.

The local perturbations in Table B9 test whether changes in wall brightness, texture, and high-frequency variation introduce erroneous screen boundaries. LTOR maintains the lowest false-edge rate together with the best reconstruction metrics under brightness, texture, and high-frequency perturbations. The paired improvement is important: the recovered image becomes more faithful while fewer competing wall patterns appear as screen boundaries.

By combining local statistics with multiscale and frequency context, LTOR evaluates each response through the surrounding evidence before it contributes to reconstruction. This evidence-centred route preserves more of the hidden screen layout while reducing the transfer of irrelevant wall structure into the final reconstruction.

Table B9: Screen evidence-selectivity tests under the local perturbations. Best in red, second best in blue, evaluated separately within each perturbation group.

Perturbation	Representation	Amplitude	PSNR↑	SSIM↑	LPIPS↓	False-edge rate↓
Local Brightness	Uniform	0.060000	20.58	0.7394	0.3920	0.4802
Local Brightness	CNN-Pyramid	0.060000	20.67	0.7408	0.3903	0.4786
Local Brightness	Window-Attention	0.060000	20.56	0.7393	0.3923	0.4800
Local Brightness	Frequency-Prompt	0.060000	20.60	0.7400	0.3912	0.4794
Local Brightness	**LTOR**	0.060000	**20.70**	**0.7459**	**0.3818**	**0.4681**
Wall Texture	Uniform	0.040000	23.44	0.8027	0.3159	0.4151
Wall Texture	CNN-Pyramid	0.040000	23.62	0.8034	0.3149	0.4144
Wall Texture	Window-Attention	0.040000	23.58	0.8028	0.3158	0.4151
Wall Texture	Frequency-Prompt	0.040000	23.49	0.8030	0.3156	0.4147
Wall Texture	**LTOR**	0.040000	**23.88**	**0.8088**	**0.3051**	**0.4041**
High-frequency Variation	Uniform	0.025000	23.41	0.8020	0.3163	0.4148
High-frequency Variation	CNN-Pyramid	0.025000	23.57	0.8027	0.3153	0.4142
High-frequency Variation	Window-Attention	0.025000	23.49	0.8021	0.3162	0.4149
High-frequency Variation	Frequency-Prompt	0.025000	23.45	0.8023	0.3160	0.4145
High-frequency Variation	**LTOR**	0.025000	**23.83**	**0.8079**	**0.3059**	**0.4041**

Together with Supplementary Figure A7, these perturbation results connect quantitative evidence selectivity with the same internal feature-attribution pattern observed across the model. Across the four content categories, LTOR responses concentrate on wall regions carrying informative brightness transitions, layout cues, and principal boundaries, while EBL aligns subsequent refinement with the recovered page, game board, lock screen, and home scene. The correspondence between the perturbation metrics and these internal responses provides a direct link between measured robustness and the model's intended mechanism.

Figure B8 further exposes the spatially uneven value of a wall observation. Equal-area occlusions produce different reconstruction changes across the nine regions, showing spatially uneven reconstruction influence under occlusions of identical area and comparable spatial extent. Regions with stronger influence correspond to portions of the wall that contribute more substantially to the recovered brightness distribution, layout, or principal boundaries, while other regions have a noticeably smaller effect.

This uneven influence motivates the local representation used by EBLTR. LTOR estimates the wall-derived support for recovered brightness, layout, and structural boundaries, and EBL directs refinement capacity according to that estimate. The framework adapts reconstruction effort to the evidence carried by the observation. The representative examples from all four content categories show that spatially uneven influence recurs across distinct screen structures and content organizations. The corresponding influence score in Figure 6 uses $\max(0, -\Delta\mathrm{PSNR})$, $\max(0, -\Delta\mathrm{SSIM})$, and $\max(0, \Delta\mathrm{LPIPS})$. Each quantity is min–max scaled over the nine occlusions, with a constant quantity assigned zero, and the three scaled quantities are averaged for each region.

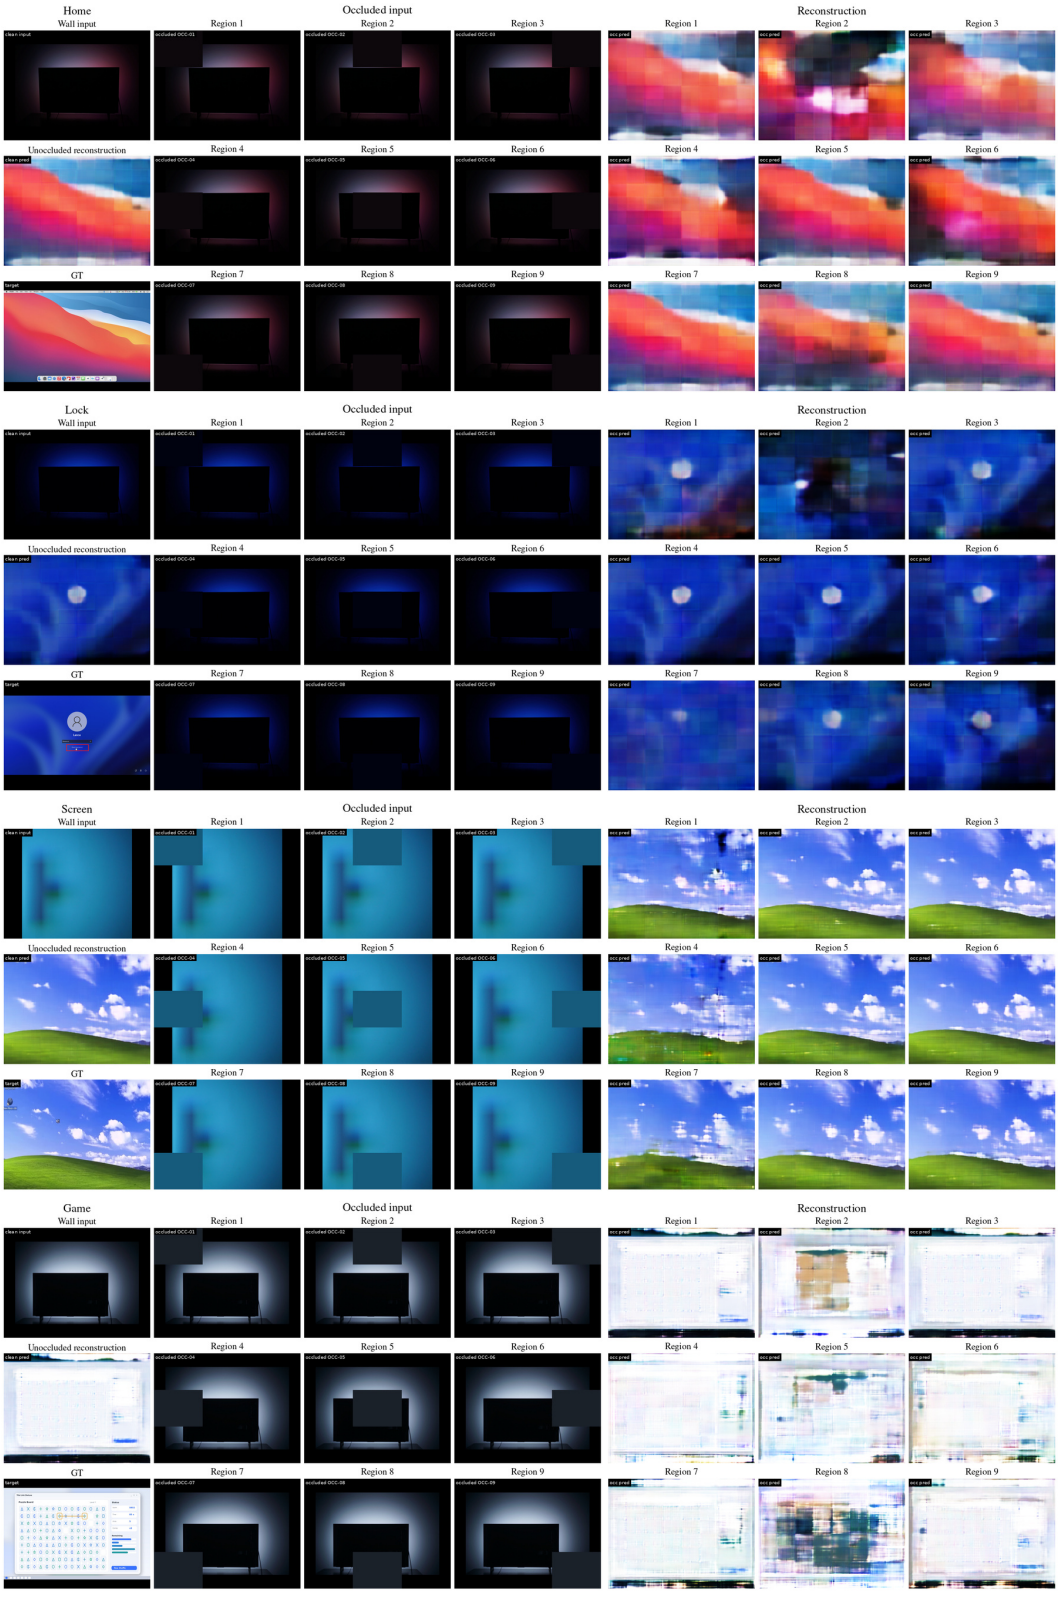

Figure B8: Nine-region occlusion diagnostic across representative Screen, Game, Lock, and Home examples.

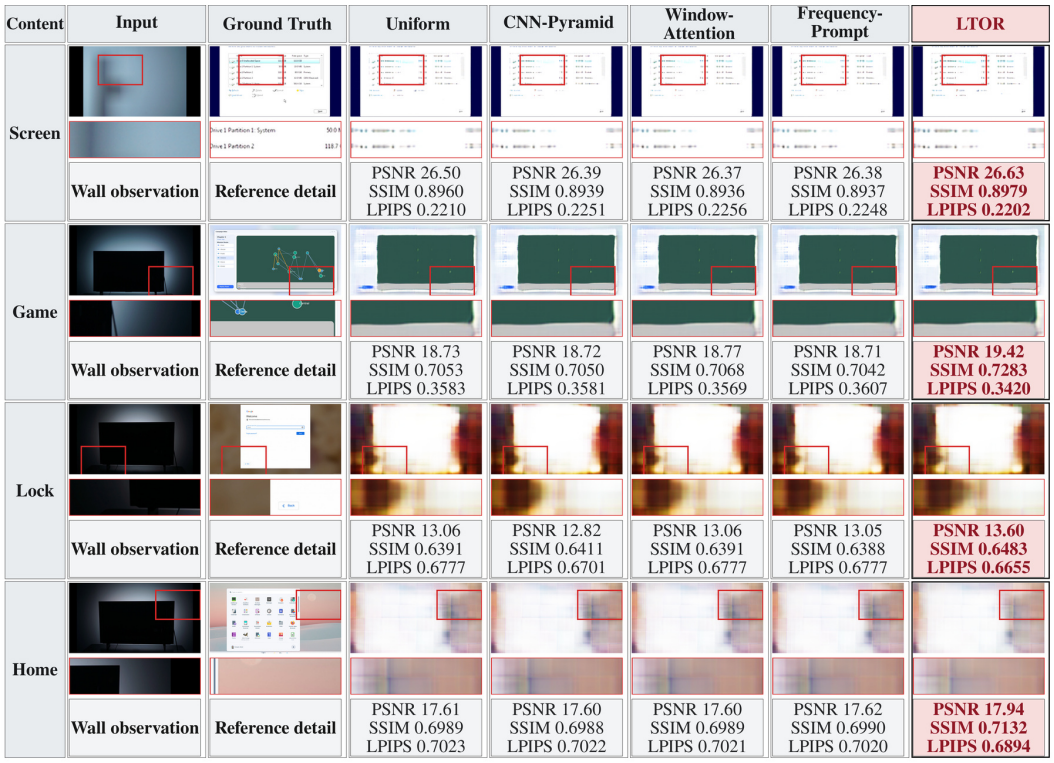

Figure B9: Same-backbone visual comparison of evidence representations across Screen, Game, Lock, and Home.

Figure B9 isolates the value of evidence organization through a visual comparison conducted under a shared reconstruction backbone. All variants use the same reconstruction system, allowing the differences in their reconstructed outputs to be considered in relation to how each representation organizes the information carried by the wall observation. The comparison therefore focuses on the effectiveness of evidence organization within the shared system, rather than on differences between reconstruction backbones. Across Screen, Game, Lock, and Home, LTOR more consistently transfers wall-derived colour partitions, regional layout, and structural boundaries into the reconstructed output. These visual aspects provide a basis for examining the correspondence between the diffuse wall response and the hidden display structure. In particular, the comparison concerns how the colour partitions, spatial arrangement of regions, and structural boundaries conveyed by the wall observation are reflected in the recovered screen. Viewed across the four content categories, the outputs illustrate the clearer correspondence that LTOR establishes through its organization of wall-derived evidence.

The visual comparison gives concrete form to Appendix Table A2, presenting the reconstructed appearance associated with the evidence representations compared in the table. Uniform and alternative feature organizations expose the same input to the same reconstruction backbone; the distinction of interest is therefore how that input is organized as evidence for reconstruction. LTOR explicitly organizes broad layout evidence together with finer boundary cues, incorporating both aspects into its organization of the information available from the wall observation. Its emphasis is on the structures that remain useful for recovering the hidden screen, including the broad layout and finer boundaries described above. This organization provides a stronger starting point for reconstructing the screen's brightness distribution, major regional divisions, and principal content boundaries. These are the aspects of reconstruction through which the visual comparison illustrates the value of the organized evidence. Across Screen, Game, Lock, and Home, the figure thus complements the tabulated comparison by showing how LTOR's evidence organization is reflected in the recovery of the hidden screen's regional and boundary structure.

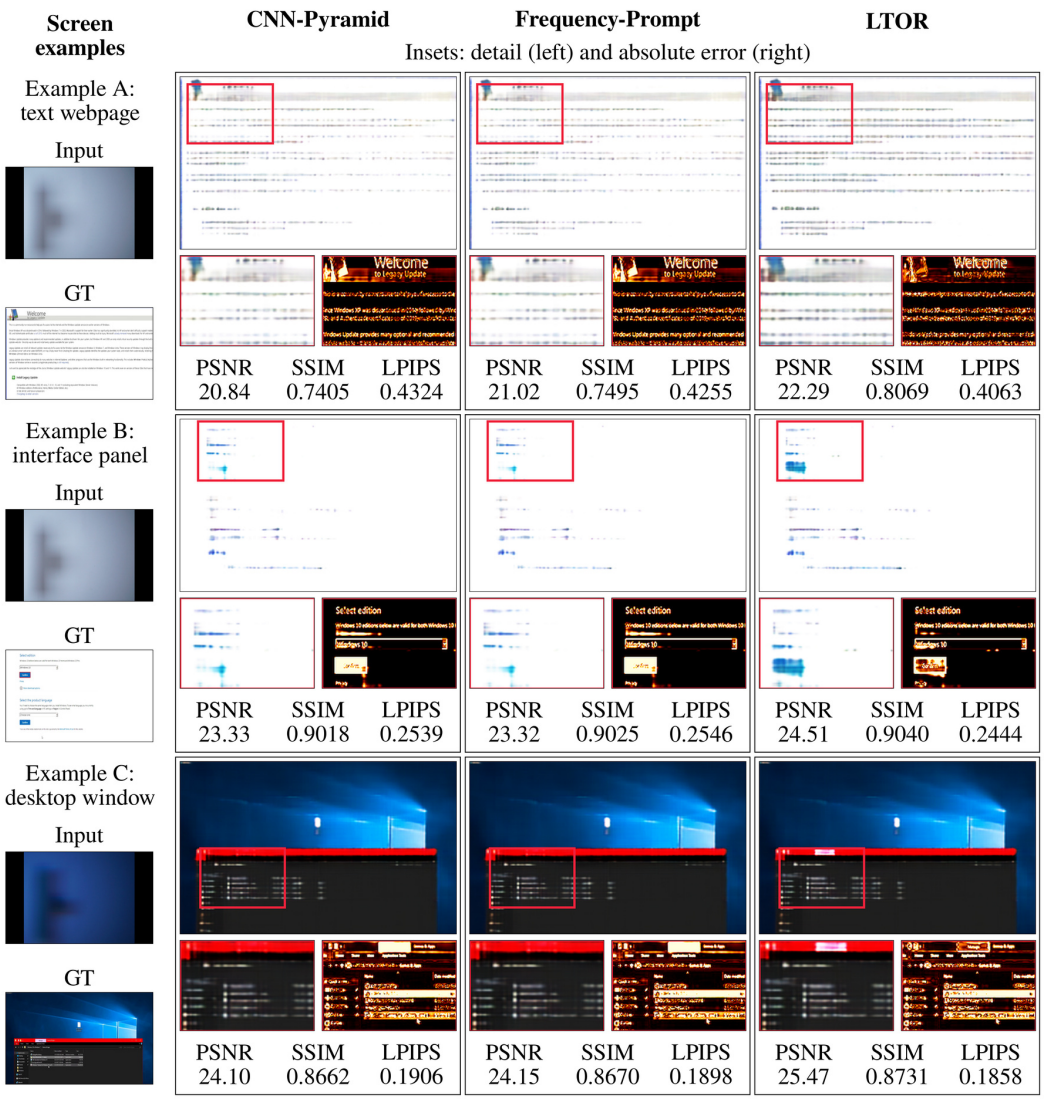

Figure B10: Screen examples for the same-backbone representation study. Red boxes mark the reconstruction regions shown in the enlarged panels; the ground truth provides the corresponding screen reference. Enlarged panels show the recovered detail and its error. Error is the RGB-channel-mean absolute difference from the target, with each prediction's full-image 98th-percentile error used as its display scale.

Figure B10 focuses on representative Screen cases where differences in evidence representation produce visible changes in local reconstruction. LTOR preserves page regions, brightness transitions, and boundary cues that connect the blurred wall observation to the hidden display. Compared with less structured representations, these responses remain more spatially coherent and better aligned with the dominant organization of the target screen.

Across the selected examples, the representation advantage appears at the level required by the recovery task. Broad page brightness is separated into meaningful regions, while the strongest surviving transitions are retained as principal contours distinct from unrelated wall responses. These organized cues provide the later evidence-budgeted stage with a more informative observability map, allowing EBL to set larger correction limits where the captured wall signal provides stronger support and smaller limits around the conservative estimate where support is weaker.

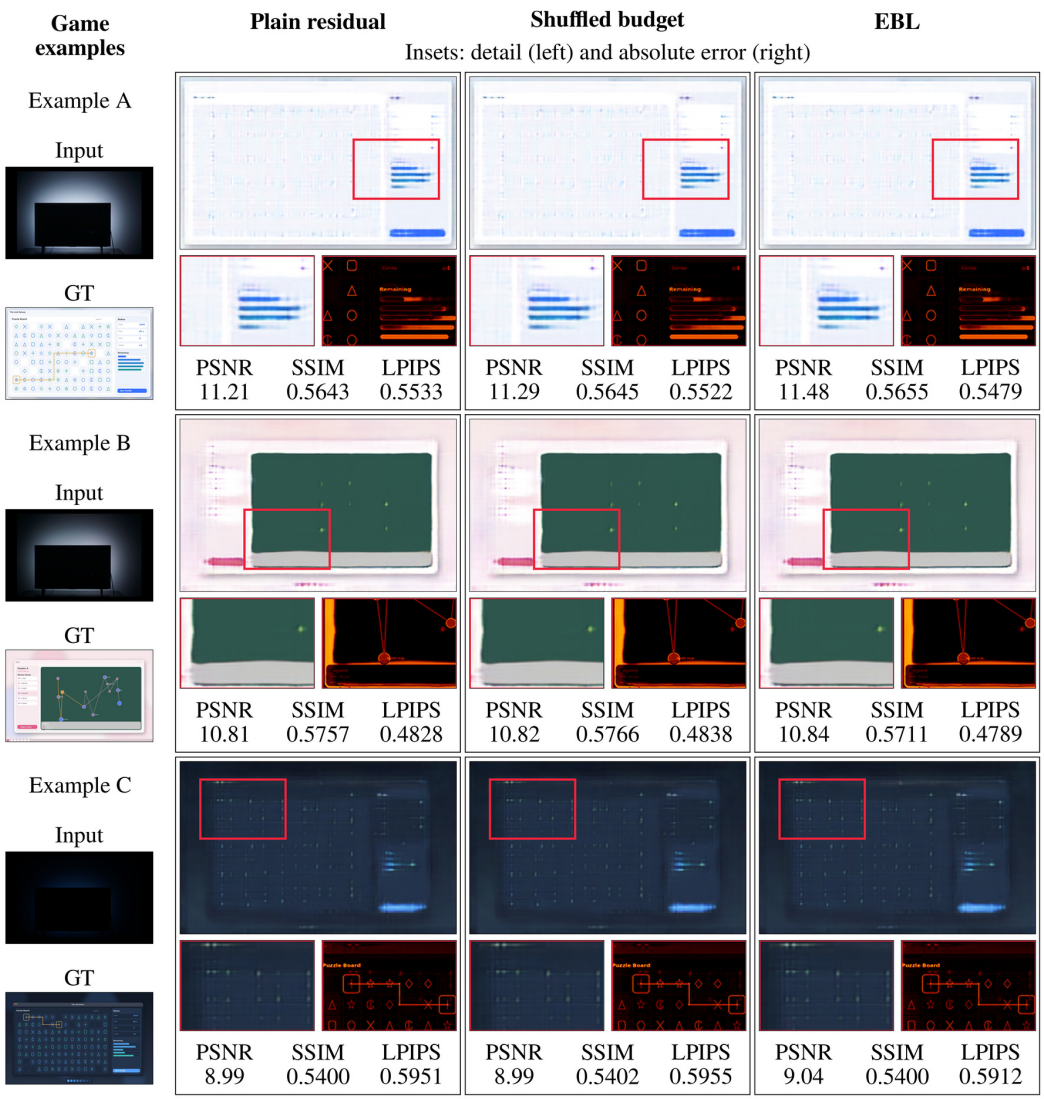

Figure B11: Game examples for fixed LTOR with alternative detail-learning rules. Red boxes mark the reconstruction regions shown in the enlarged panels; the ground truth provides the corresponding screen reference. Enlarged panels show the recovered detail and its error. Error is the RGB-channel-mean absolute difference from the target, with each prediction's full-image 98th-percentile error used as its display scale.

With LTOR fixed, Figure B11 compares alternative detail-learning rules within the same evidence representation. In this setting, EBL links local correction to the evidence available in each example, making the relationship between evidence and refinement the focus of the comparison. The enlarged regions show how refinement acts on bright interface elements, playing-area boundaries, and dark screen regions. These local views complement the full-screen reconstructions, allowing the recovered details to be examined in the context of the overall screen arrangement that the refinement preserves.

The sequence illustrates why budget placement matters after the representation has organized the wall response. EBL uses the spatially varying evidence estimate to define a corresponding range of permitted corrections, so the correction range follows the evidence available at different screen locations. Budget placement thus connects the organized wall-derived evidence to the local refinement of the reconstruction. Through this connection, the permitted changes to individual regions remain linked to the recovered screen's overall organization, relating the treatment of local detail to the arrangement of the screen as a whole.

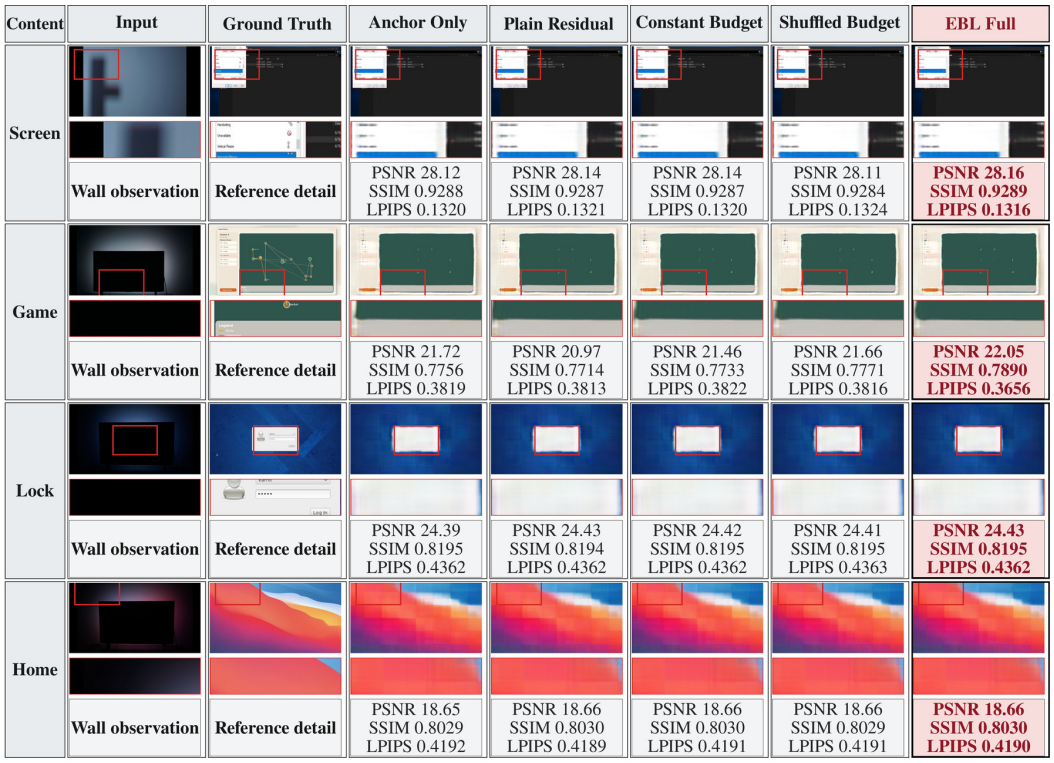

Figure B12: Visual comparison of detail-learning rules with LTOR fixed across all four content categories.

Figure B12 shows how the learning rule turns the same LTOR representation into an output. Evidence-aligned budgeting preserves the broad scene organization in all four content categories, from page and board divisions to lock-screen bands and home-scene regions. The fixed representation makes this a direct visual comparison of refinement rules. EBL carries LTOR's observability into the output more coherently than uniform or misaligned alternatives, connecting the method's second stage directly to recoverable brightness distributions, regional layouts, and principal contours. Appendix Table A6 provides the complete quantitative counterpart to Figure B12. Together they show that the advantage follows how the learning rule uses the same wall-derived evidence: EBL keeps refinement spatially matched to observability across all four content categories.

Table B10: Spatial allocation of the detail budget with LTOR fixed. Content groups report PSNR↑/SSIM↑/LPIPS↓. Red and blue indicate the best and second-best distinct displayed values; ties share the same colour.

Learning	Screen			Game			Lock			Home		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Constant budget	23.80	0.8092	0.3055	18.73	0.6660	0.4888	17.05	0.6240	0.6315	14.75	0.5474	0.7101
Shuffled budget	23.71	0.8091	0.3057	18.73	0.6660	0.4888	17.05	0.6240	0.6315	14.75	0.5473	0.7101
EBL	**23.82**	**0.8097**	**0.3045**	**18.76**	**0.6668**	**0.4861**	**17.18**	**0.6293**	**0.6314**	**14.75**	**0.5494**	**0.7091**

EBL achieves the highest SSIM and lowest LPIPS across the four content categories in Table B10. This advantage links the allocation of local correction to the structure and appearance of the recovered screen. This controlled comparison captures EBL's conceptual advantage in its simplest form. A constant budget assigns the same correction limit everywhere, and a shuffled budget relocates the limits; EBL preserves their alignment with wall-derived evidence, directing refinement capacity to screen regions and contours supported by the observation. The constant control substitutes $\mathbf{O} = 0.5$ at every location. The shuffled control circularly shifts the observability map by $\max(1, \lfloor H/4 \rfloor)$ rows and $\max(1, \lfloor W/6 \rfloor)$ columns. This shift preserves the map's sum while changing its spatial alignment; EBL uses the original observability map in the local budget $\alpha\mathbf{O}$.

Table B11 compares representation and learning choices across the four content categories. The LTOR–EBL configuration achieves the best or joint-best scores across all four content categories, supporting their combined use for screen recovery.

Table B11: Factorial study of the representations and learning rules. Content groups report PSNR↑/SSIM↑/LPIPS↓. Red and blue indicate the best and second-best distinct displayed values; ties share the same colour.

Representation	Learning	Screen			Game			Lock			Home		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
CNN-Pyramid	Budgeted learning	23.81	0.8077	0.3073	18.73	0.6666	0.4870	17.15	0.6276	0.6319	14.75	0.5481	0.7097
CNN-Pyramid	Plain learning	23.81	0.8086	0.3061	18.72	0.6665	0.4869	17.15	0.6279	0.6318	14.75	0.5480	0.7098
Frequency-Prompt	Budgeted learning	23.76	0.8093	0.3046	18.73	0.6664	0.4867	17.15	0.6276	0.6319	14.75	0.5480	0.7097
Frequency-Prompt	Plain learning	23.77	0.8087	0.3060	18.73	0.6663	0.4869	17.15	0.6279	0.6317	14.75	0.5480	0.7098
Uniform	Budgeted learning	23.71	0.8094	0.3046	18.72	0.6662	0.4869	17.15	0.6276	0.6319	14.75	0.5479	0.7092
Uniform	Plain learning	23.76	0.8093	0.3050	18.73	0.6667	0.4870	17.15	0.6279	0.6318	14.75	0.5480	0.7098
Window-Attention	Budgeted learning	23.78	0.8087	0.3059	18.72	0.6661	0.4867	17.15	0.6276	0.6319	14.75	0.5480	0.7097
Window-Attention	Plain learning	23.77	0.8087	0.3060	18.72	0.6662	0.4870	17.15	0.6279	0.6318	14.75	0.5480	0.7098
LTOR	Plain learning	23.74	0.8091	0.3050	18.73	0.6660	0.4889	17.05	0.6243	0.6314	14.75	0.5493	0.7092
LTOR	**EBL**	**23.82**	**0.8097**	**0.3045**	**18.76**	**0.6668**	**0.4861**	**17.18**	**0.6293**	**0.6314**	**14.75**	**0.5494**	**0.7091**

The factorial structure clarifies the division of labour. Representation first establishes where brightness, regions, and contours remain observable through the wall; learning then converts that ordering into spatially matched correction. Their combination forms the complete EBLTR chain and gives the wall observation its most effective route to the hidden screen.

Table B12: Reconstruction performance under global and radial illumination shifts. Content columns report the mean under each control and shift level over 251 Screen, 250 Game, 279 Lock, and a fixed 300-sample Home test subset. The zero-shift reference is recomputed with the same checkpoint and sample set as each shifted sequence.

Shift (EV)	Metric	Matched global				Radial			
		Screen	Game	Lock	Home	Screen	Game	Lock	Home
0.00	PSNR↑	24.27	18.71	17.16	14.74	24.27	18.71	17.16	14.74
	SSIM↑	0.8189	0.6665	0.6295	0.5486	0.8189	0.6665	0.6295	0.5486
	LPIPS↓	0.2913	0.4864	0.6299	0.7098	0.2913	0.4864	0.6299	0.7098
0.25	PSNR↑	23.62	18.71	17.16	14.73	23.03	18.67	17.14	14.72
	SSIM↑	0.8046	0.6663	0.6296	0.5482	0.7919	0.6651	0.6294	0.5481
	LPIPS↓	0.3142	0.4867	0.6302	0.7106	0.3336	0.4883	0.6300	0.7100
0.50	PSNR↑	22.42	18.70	17.15	14.72	21.45	18.50	17.11	14.67
	SSIM↑	0.7711	0.6657	0.6296	0.5477	0.7422	0.6596	0.6287	0.5471
	LPIPS↓	0.3603	0.4874	0.6307	0.7116	0.3912	0.4940	0.6308	0.7106
0.75	PSNR↑	21.30	18.68	17.14	14.70	19.67	18.29	17.04	14.60
	SSIM↑	0.7322	0.6648	0.6295	0.5470	0.6836	0.6530	0.6275	0.5455
	LPIPS↓	0.4032	0.4886	0.6314	0.7127	0.4495	0.5024	0.6320	0.7118
1.00	PSNR↑	20.11	18.65	17.12	14.67	17.79	18.06	16.95	14.49
	SSIM↑	0.6902	0.6635	0.6293	0.5462	0.6269	0.6464	0.6256	0.5434
	LPIPS↓	0.4451	0.4902	0.6321	0.7139	0.5040	0.5123	0.6336	0.7133

Global and radial illumination changes produce distinct reconstruction responses in Table B12. This dependence on the spatial form of illumination motivates organizing recovery evidence locally. EBLTR makes that organization actionable by converting observability into a correction budget, linking the illumination study to the spatial-allocation comparisons in Table B10.

B.2 Additional Robustness Studies

Does hidden-screen content remain recoverable as its wall observation changes? Illumination, sampling, and noise alter the strength and clarity of screen cues; camera geometry changes their observed spatial arrangement; wall texture and reflectance change the surrounding response. The following

comparisons connect these changes to the recovered Game and Lock content. The quantitative sequences apply controlled test-time transformations to the existing paired wall inputs, retaining each screen target and the trained reconstruction weights. The accompanying photographs document the illumination, camera-placement, and wall-material settings motivating these transformations. Brightness uses a multiplier in encoded RGB, reflectance uses a gain in linear RGB, and wall texture uses a smooth multiplicative field specified by its RMS amplitude. Gaussian noise is parameterized by input SNR: $\mathrm{SNR} = 10\log_{10}(\mathrm{Var}(\mathbf{x})/\sigma^2)$, with variance computed over the encoded RGB input and independent zero-mean noise of standard deviation σ, followed by clipping and 8-bit quantization. Lower SNR corresponds to stronger added noise.

B.2.1 Wall Brightness



Figure B13: Indoor illumination settings associated with the wall-brightness study. The photographs illustrate the acquisition conditions motivating the controlled brightness transformations.

Brightness changes the strength of the wall spot while the hidden-screen layout remains fixed. LTOR re-evaluates the evidence carried by the altered response, and EBL matches local refinement to that evidence, supporting stable recovery of the broad brightness distribution and regional layout.

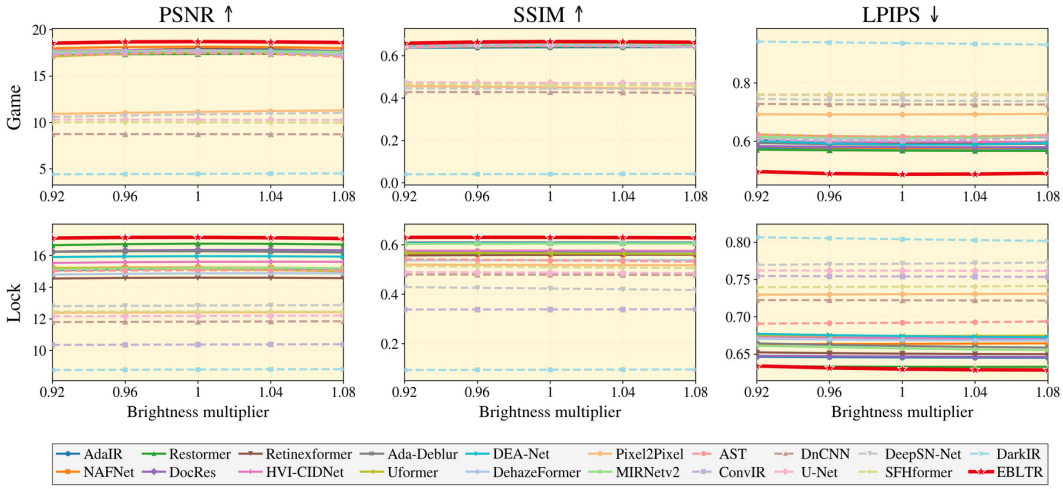

Figure B14: Game and Lock robustness curves under the wall-brightness factors.

EBLTR achieves the highest PSNR and SSIM at every brightness factor on Game and Lock in Figure B14. The result connects changing wall intensity to stable recovery of screen-wide organization. The brightness sequence tests how changes in the strength of the same wall response affect screen recovery. LTOR processes the revised brightness response, EBL uses the resulting local correction budget, and the PSNR and SSIM trajectories preserve the closest connection to the hidden screen across both content categories.

Table B13 confirms the Game-side curve across all baselines and factors. EBLTR leads the PSNR, SSIM, and LPIPS values throughout the sequence, showing that evidence-budgeted recovery continues to preserve the game's brightness distribution and regional structure as wall intensity varies.

Retinexformer uses illumination-guided restoration, while HVI-CIDNet decouples colour and intensity for low-light enhancement (Cai et al., 2023; Yan et al., 2025). Our EBLTR controls refinement through wall-derived observability and achieves higher PSNR than both methods at every Game brightness factor in our comparison, with minimum gains of 0.79 and 0.78 dB, respectively (Table B13).

Table B13: Complete Game comparison under wall brightness. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Brightness multiplier 0.92			Brightness multiplier 0.96			Brightness multiplier 1.0			Brightness multiplier 1.04			Brightness multiplier 1.08		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.75	0.4278	0.7276	8.74	0.4279	0.7264	8.74	0.4273	0.7259	8.73	0.4260	0.7255	8.72	0.4239	0.7255
U-Net	10.30	0.4744	0.7585	10.30	0.4730	0.7584	10.29	0.4715	0.7581	10.27	0.4701	0.7577	10.26	0.4687	0.7575
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.39	0.6370	0.5784	17.42	0.6383	0.5765	17.44	0.6394	0.5754	17.44	0.6401	0.5748	17.45	0.6403	0.5747
NAFNet	17.99	0.6523	0.5823	18.11	0.6564	0.5783	18.16	0.6587	0.5759	18.12	0.6585	0.5759	18.00	0.6558	0.5787
Restormer	17.30	0.6414	0.5707	17.34	0.6428	0.5688	17.37	0.6438	0.5676	17.38	0.6446	0.5667	17.38	0.6449	0.5665
Uformer	17.11	0.6393	0.6232	17.40	0.6442	0.6169	17.55	0.6470	0.6142	17.50	0.6461	0.6149	17.21	0.6413	0.6203
MIRNetv2	17.63	0.6529	0.6144	17.68	0.6551	0.6128	17.68	0.6559	0.6119	17.65	0.6556	0.6123	17.58	0.6542	0.6138
AST	17.22	0.6402	0.6221	17.39	0.6438	0.6176	17.46	0.6452	0.6158	17.35	0.6440	0.6170	17.07	0.6401	0.6211
ConvIR	17.46	0.6486	0.6060	17.61	0.6514	0.6020	17.64	0.6514	0.6013	17.52	0.6488	0.6047	17.25	0.6429	0.6125
DeepSN-Net	10.59	0.4447	0.7446	10.74	0.4447	0.7410	10.87	0.4440	0.7390	10.96	0.4426	0.7378	11.03	0.4406	0.7373
SFHformer	10.01	0.4661	0.7609	10.01	0.4640	0.7601	10.01	0.4620	0.7598	9.99	0.4600	0.7599	9.97	0.4579	0.7603
Models for Related Degradations															
DocRes	17.45	0.6442	0.5830	17.50	0.6461	0.5806	17.53	0.6472	0.5794	17.53	0.6474	0.5789	17.52	0.6468	0.5793
Retinexformer	17.45	0.6359	0.6043	17.79	0.6491	0.5958	17.92	0.6556	0.5924	17.87	0.6558	0.5935	17.71	0.6524	0.5979
HVI-CIDNet	17.76	0.6538	0.5994	17.77	0.6550	0.5984	17.77	0.6555	0.5978	17.76	0.6556	0.5978	17.74	0.6554	0.5982
Ada-Deblur	17.53	0.6491	0.5938	17.68	0.6548	0.5894	17.75	0.6577	0.5878	17.74	0.6582	0.5882	17.64	0.6567	0.5908
DEA-Net	17.39	0.6472	0.6007	17.60	0.6534	0.5922	17.73	0.6579	0.5879	17.73	0.6591	0.5884	17.58	0.6560	0.5938
DehazeFormer	17.40	0.6374	0.6129	17.68	0.6461	0.6087	17.77	0.6494	0.6074	17.67	0.6474	0.6104	17.34	0.6399	0.6177
Pixel2Pixel	10.92	0.4563	0.6923	11.03	0.4535	0.6915	11.14	0.4502	0.6917	11.23	0.4466	0.6924	11.29	0.4429	0.6933
DarkIR	4.42	0.0398	0.9414	4.44	0.0404	0.9387	4.46	0.0409	0.9361	4.49	0.0414	0.9336	4.51	0.0419	0.9312
EBLTR	**18.54**	**0.6587**	**0.4955**	**18.68**	**0.6645**	**0.4889**	**18.71**	**0.6665**	**0.4864**	**18.67**	**0.6657**	**0.4872**	**18.61**	**0.6635**	**0.4902**

The blue playing field and adjacent panel in Figure B15 provide fixed landmarks across the brightness sequence. EBLTR preserves their overall arrangement as the observed wall intensity changes, connecting brightness robustness to a stable recovered layout. Table B13 evaluates the same brightness variation across all 19 reconstruction baselines. Together, the two elements show that the changing wall spot still carries enough evidence for preliminary recovery and that EBLTR uses that evidence more effectively across the tested intensity range. Table B13 supplies the dataset-level PSNR, SSIM, and LPIPS evidence underlying Figure B15, linking the representative board reconstruction to the same fidelity, structure, and perceptual criteria used throughout the paper.

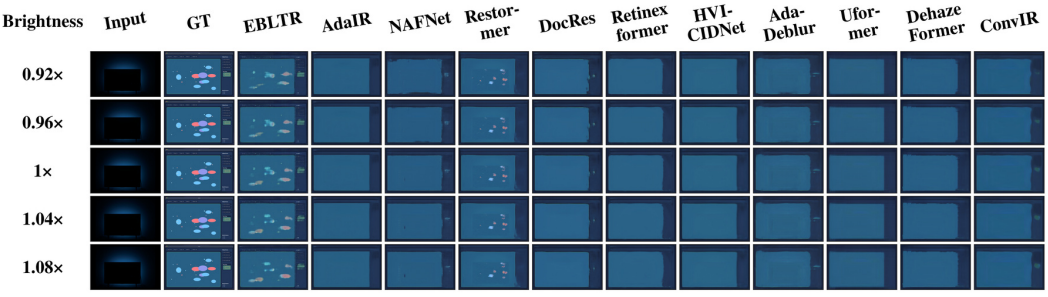

Figure B15: Representative Game reconstructions under the wall-brightness sequence.

Across the brightness sequence in Table B14, EBLTR maintains the closest intensity and structural agreement with the hidden Lock screen, achieving the highest PSNR and SSIM at every factor. The recovered colour bands and principal divisions remain coherent as the wall intensity changes. This second content family strengthens the mechanism claim. Brightness changes alter the wall evidence, while LTOR and EBL continue to recover the stable screen organization from that revised signal.

The agreement between Game and Lock shows a common evidence-budgeted response across two distinct layouts.

Table B14: Complete Lock comparison under wall brightness. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Brightness multiplier 0.92			Brightness multiplier 0.96			Brightness multiplier 1			Brightness multiplier 1.04			Brightness multiplier 1.08		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.79	0.4796	0.7223	11.81	0.4792	0.7222	11.82	0.4787	0.7221	11.83	0.4782	0.7219	11.84	0.4777	0.7217
U-Net	12.15	0.4890	0.7621	12.17	0.4882	0.7619	12.18	0.4874	0.7618	12.19	0.4865	0.7617	12.20	0.4855	0.7615
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.08	0.5720	0.6463	15.13	0.5726	0.6459	15.17	0.5729	0.6455	15.18	0.5730	0.6452	15.19	0.5730	0.6450
NAFNet	15.21	0.5648	0.6634	15.24	0.5634	0.6634	15.24	0.5619	0.6637	15.23	0.5605	0.6640	15.20	0.5592	0.6644
Restormer	16.66	0.6074	0.6339	16.73	0.6074	0.6333	16.75	0.6070	0.6329	16.74	0.6064	0.6329	16.70	0.6057	0.6330
Uformer	15.22	0.5666	0.6752	15.20	0.5660	0.6747	15.15	0.5651	0.6744	15.08	0.5641	0.6743	15.00	0.5627	0.6746
MIRNetv2	15.13	0.6029	0.6610	15.18	0.6037	0.6592	15.21	0.6042	0.6575	15.22	0.6044	0.6562	15.20	0.6043	0.6553
AST	15.02	0.5413	0.6907	15.06	0.5390	0.6913	15.07	0.5365	0.6920	15.07	0.5340	0.6928	15.06	0.5315	0.6936
ConvIR	10.35	0.3381	0.7547	10.36	0.3383	0.7543	10.37	0.3385	0.7540	10.38	0.3387	0.7537	10.39	0.3388	0.7535
DeepSN-Net	12.80	0.4286	0.7696	12.82	0.4257	0.7703	12.84	0.4228	0.7711	12.85	0.4199	0.7719	12.87	0.4169	0.7727
SFHformer	12.46	0.5134	0.7396	12.46	0.5115	0.7399	12.46	0.5095	0.7404	12.46	0.5074	0.7408	12.46	0.5054	0.7413
Models for Related Degradations															
DocRes	16.25	0.6092	0.6476	16.32	0.6094	0.6471	16.36	0.6094	0.6467	16.36	0.6092	0.6463	16.34	0.6089	0.6461
Retinexformer	14.55	0.5563	0.6523	14.59	0.5577	0.6512	14.60	0.5586	0.6506	14.60	0.5592	0.6501	14.57	0.5595	0.6499
HVI-CIDNet	15.54	0.5757	0.6735	15.58	0.5759	0.6724	15.61	0.5758	0.6714	15.61	0.5755	0.6708	15.61	0.5750	0.6705
Ada-Deblur	16.23	0.6078	0.6643	16.27	0.6084	0.6624	16.27	0.6086	0.6608	16.25	0.6084	0.6596	16.20	0.6079	0.6587
DEA-Net	15.91	0.6060	0.6771	15.95	0.6066	0.6754	15.96	0.6069	0.6741	15.96	0.6068	0.6732	15.93	0.6062	0.6726
DehazeFormer	14.80	0.5365	0.6703	14.85	0.5371	0.6694	14.87	0.5375	0.6687	14.88	0.5375	0.6683	14.87	0.5373	0.6681
Pixel2Pixel	12.39	0.5192	0.7297	12.40	0.5185	0.7299	12.41	0.5178	0.7302	12.41	0.5171	0.7305	12.42	0.5164	0.7309
DarkIR	8.76	0.0938	0.8065	8.77	0.0944	0.8052	8.79	0.0949	0.8040	8.80	0.0955	0.8028	8.81	0.0960	0.8017
EBLTR	**17.11**	**0.6295**	**0.6340**	**17.15**	**0.6297**	**0.6315**	**17.16**	**0.6295**	**0.6298**	**17.13**	**0.6287**	**0.6289**	**17.08**	**0.6277**	**0.6286**

Figure B16 supplies the complementary Lock sequence. The recovered colour bands, regional divisions, and main outline remain visible as the wall response changes, making the retained low-frequency and boundary evidence directly interpretable. Paired with Table B14, the figure shows both the visual content and the dataset-level direction of the result. EBLTR preserves the strongest correspondence to the hidden lock screen throughout the tested brightness sequence.

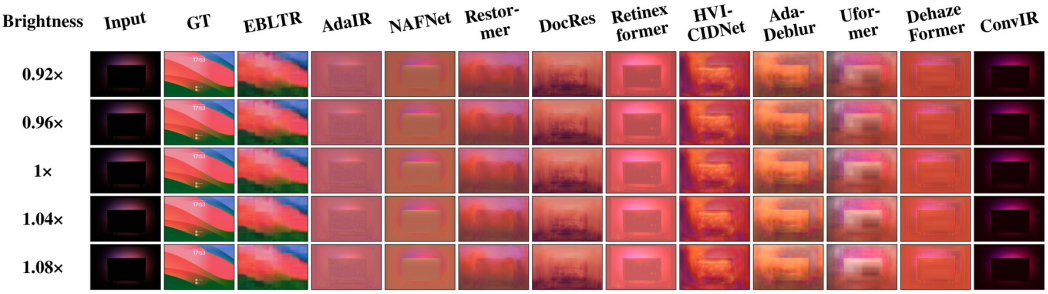

Figure B16: Representative Lock reconstructions under the wall-brightness sequence.

B.2.2 Distance and Effective Sampling Resolution



Figure B17: Camera-placement photographs associated with the distance and effective-sampling study.

The distance-motivated sampling study varies the effective spatial resolution of the wall observation. The experiment highlights the coarse screen information that survives this change: EBLTR keeps the recovery centered on color fields, layout, and main contours that remain represented in the observation.

Table B15: Complete Game comparison under distance and effective sampling resolution. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Sampling factor 1.1			Sampling factor 1.15			Sampling factor 1.2			Sampling factor 1.25			Sampling factor 1.3		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.74	0.4280	0.7275	8.74	0.4280	0.7275	8.74	0.4281	0.7276	8.74	0.4280	0.7282	8.74	0.4281	0.7281
U-Net	10.24	0.4695	0.7634	10.25	0.4699	0.7616	10.27	0.4706	0.7624	10.22	0.4685	0.7609	10.25	0.4698	0.7618
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.23	0.6307	0.5842	17.28	0.6333	0.5818	17.33	0.6351	0.5805	17.14	0.6263	0.5890	17.30	0.6347	0.5819
NAFNet	18.14	0.6580	0.5759	18.16	0.6581	0.5758	18.13	0.6570	0.5761	18.13	0.6571	0.5761	18.13	0.6572	0.5765
Restormer	17.32	0.6404	0.5685	17.34	0.6412	0.5683	17.33	0.6407	0.5661	17.30	0.6405	0.5705	17.35	0.6419	0.5672
Uformer	17.53	0.6464	0.6111	17.54	0.6467	0.6113	17.51	0.6457	0.6107	17.54	0.6463	0.6120	17.53	0.6459	0.6122
MIRNetv2	17.67	0.6559	0.6144	17.67	0.6561	0.6142	17.68	0.6561	0.6136	17.66	0.6557	0.6150	17.68	0.6561	0.6141
AST	17.41	0.6444	0.6137	17.28	0.6433	0.6155	17.43	0.6445	0.6129	17.21	0.6429	0.6166	17.33	0.6435	0.6155
ConvIR	17.61	0.6506	0.6031	17.63	0.6511	0.6032	17.57	0.6496	0.6027	17.62	0.6510	0.6035	17.63	0.6514	0.6030
DeepSN-Net	10.88	0.4453	0.7403	10.88	0.4451	0.7405	10.88	0.4450	0.7399	10.88	0.4455	0.7410	10.88	0.4453	0.7423
SFHformer	10.01	0.4634	0.7626	10.01	0.4632	0.7623	10.01	0.4630	0.7626	10.01	0.4637	0.7631	10.01	0.4634	0.7630
Models for Related Degradations															
DocRes	17.38	0.6389	0.5874	17.44	0.6424	0.5852	17.45	0.6431	0.5839	17.24	0.6338	0.5932	17.42	0.6410	0.5858
Retinexformer	17.85	0.6528	0.5926	17.86	0.6537	0.5920	17.87	0.6542	0.5909	17.77	0.6496	0.5953	17.85	0.6532	0.5928
HVI-CIDNet	16.59	0.6165	0.6357	16.92	0.6269	0.6267	16.69	0.6192	0.6270	17.32	0.6405	0.6117	17.56	0.6456	0.6043
Ada-Deblur	17.64	0.6544	0.5930	17.62	0.6543	0.5929	17.66	0.6557	0.5909	17.61	0.6528	0.5948	17.66	0.6555	0.5921
DEA-Net	15.54	0.6185	0.6332	16.52	0.6333	0.6237	15.28	0.6138	0.6294	15.76	0.6204	0.6389	16.17	0.6273	0.6244
DehazeFormer	17.73	0.6480	0.6083	17.74	0.6484	0.6083	17.69	0.6472	0.6092	17.73	0.6479	0.6087	17.70	0.6474	0.6089
Pixel2Pixel	11.15	0.4418	0.6992	11.16	0.4451	0.6966	11.16	0.4455	0.6968	11.15	0.4406	0.6982	11.16	0.4437	0.6974
DarkIR	4.46	0.0409	0.9398	4.46	0.0409	0.9402	4.46	0.0409	0.9419	4.46	0.0409	0.9412	4.46	0.0409	0.9408
EBLTR	**18.70**	**0.6661**	**0.4875**	**18.68**	**0.6656**	**0.4874**	**18.45**	**0.6566**	**0.4952**	**18.68**	**0.6653**	**0.4882**	**18.40**	**0.6550**	**0.4963**

Table B15 records the complete Game comparison and places every baseline under the same effective-sampling factors. EBLTR's highest PSNR and lowest LPIPS show that organizing the remaining wall evidence yields stronger recovery of the playing-area layout and major colour regions across the sampling sequence.

The multi-factor layout reveals more than a nominal-condition ranking. As spatial evidence changes, EBLTR consistently converts the available response into a closer hidden-screen estimate, supporting evidence-budgeted learning as a better match to this observation-limited task.

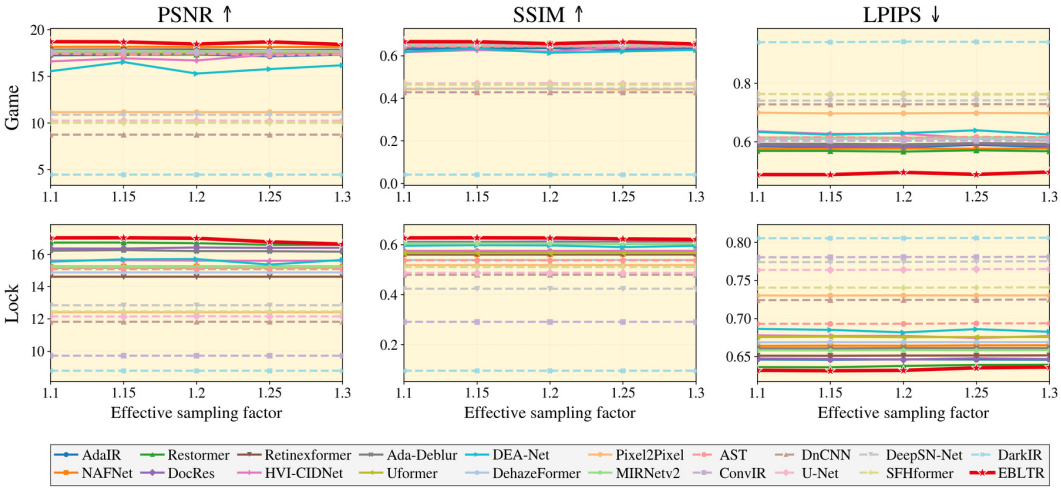

Figure B18: Game and Lock robustness curves under the distance and effective-sampling factors.

EBLTR maintains the highest PSNR and lowest LPIPS across the sampling factors on both content categories in Figure B18. The corresponding images retain the main screen divisions as the spatial sampling of the observation changes.

The retained colour fields and principal divisions match LTOR's use of broad responses together with finer transitions; EBL uses their aggregate observability to bound correction around the conservative screen estimate. Each screen target stays fixed as the spatial evidence supporting its recovery changes.

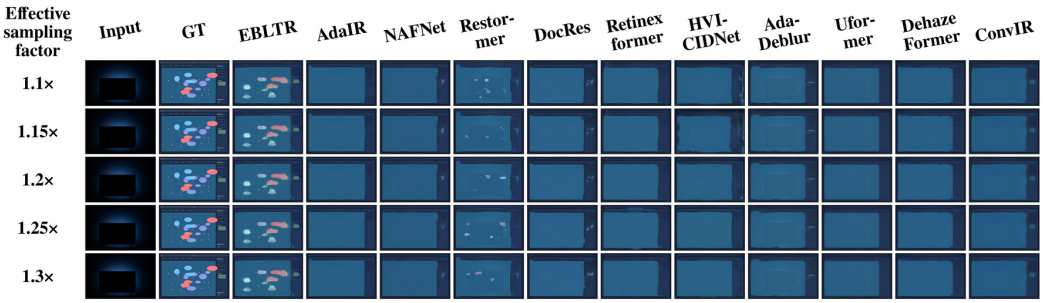

Figure B19: Representative Game reconstructions across the effective-sampling sequence.

The main playing area and side panel in Figure B19 remain separated as the wall response is sampled differently. This preservation connects the quantitative advantage to screen organization that remains recoverable across the sampling sequence.

The result illustrates the hierarchy encoded by LTOR: broad screen organization persists through coarser sampling, while EBL bounds finer corrections using the estimated observational support. The representative output thus visualizes the same mechanism measured across the complete Game table.

Table B15 anchors Figure B19 in full-test-set PSNR, SSIM, and LPIPS, confirming that the visible preservation of board structure follows the dataset-level comparison.

Table B16: Complete Lock comparison under distance and effective sampling resolution. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Sampling factor 1.1			Sampling factor 1.15			Sampling factor 1.2			Sampling factor 1.25			Sampling factor 1.3		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4792	0.7240	11.82	0.4791	0.7242	11.82	0.4792	0.7244	11.82	0.4792	0.7244	11.82	0.4792	0.7249
U-Net	12.15	0.4868	0.7638	12.14	0.4868	0.7637	12.16	0.4877	0.7639	12.15	0.4866	0.7647	12.15	0.4868	0.7649
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.17	0.5724	0.6462	15.17	0.5726	0.6458	15.17	0.5730	0.6463	15.17	0.5725	0.6462	15.17	0.5728	0.6458
NAFNet	15.24	0.5626	0.6641	15.24	0.5626	0.6643	15.24	0.5625	0.6642	15.24	0.5627	0.6646	15.24	0.5628	0.6647
Restormer	16.72	0.6051	0.6361	16.72	0.6054	0.6360	16.68	0.6048	0.6377	16.58	0.6033	0.6390	16.57	0.6035	0.6392
Uformer	15.18	0.5661	0.6755	15.18	0.5657	0.6761	15.18	0.5661	0.6754	15.17	0.5661	0.6756	15.17	0.5663	0.6757
MIRNetv2	15.21	0.6044	0.6580	15.22	0.6045	0.6580	15.22	0.6045	0.6583	15.22	0.6045	0.6582	15.22	0.6045	0.6583
AST	15.07	0.5370	0.6930	15.07	0.5369	0.6930	15.07	0.5369	0.6930	15.07	0.5369	0.6935	15.08	0.5370	0.6936
ConvIR	9.72	0.2907	0.7803	9.72	0.2907	0.7804	9.72	0.2907	0.7807	9.72	0.2907	0.7806	9.72	0.2907	0.7810
DeepSN-Net	12.84	0.4232	0.7739	12.84	0.4232	0.7740	12.84	0.4233	0.7741	12.84	0.4233	0.7748	12.84	0.4234	0.7751
SFHformer	12.46	0.5099	0.7404	12.46	0.5098	0.7406	12.46	0.5097	0.7404	12.46	0.5100	0.7407	12.46	0.5100	0.7408
Models for Related Degradations															
DocRes	16.35	0.6111	0.6474	16.35	0.6110	0.6468	16.41	0.6125	0.6461	16.39	0.6109	0.6477	16.39	0.6116	0.6466
Retinexformer	14.61	0.5599	0.6511	14.61	0.5599	0.6510	14.61	0.5601	0.6513	14.61	0.5602	0.6514	14.61	0.5603	0.6515
HVI-CIDNet	15.60	0.5749	0.6778	15.62	0.5755	0.6769	15.60	0.5750	0.6771	15.59	0.5752	0.6744	15.60	0.5751	0.6769
Ada-Deblur	16.21	0.6060	0.6609	16.25	0.6073	0.6608	16.19	0.6055	0.6611	16.18	0.6053	0.6611	16.15	0.6041	0.6611
DEA-Net	15.54	0.5952	0.6863	15.68	0.5976	0.6850	15.70	0.5966	0.6817	15.36	0.5898	0.6860	15.65	0.5950	0.6825
DehazeFormer	14.87	0.5381	0.6687	14.87	0.5381	0.6688	14.87	0.5381	0.6688	14.87	0.5382	0.6689	14.87	0.5382	0.6689
Pixel2Pixel	12.41	0.5179	0.7302	12.41	0.5178	0.7301	12.41	0.5179	0.7303	12.41	0.5179	0.7303	12.41	0.5179	0.7303
DarkIR	8.79	0.0950	0.8053	8.79	0.0950	0.8053	8.79	0.0950	0.8055	8.79	0.0950	0.8057	8.79	0.0950	0.8059
EBLTR	**17.01**	**0.6268**	**0.6319**	**17.02**	**0.6273**	**0.6314**	**16.99**	**0.6265**	**0.6319**	**16.76**	**0.6228**	**0.6352**	**16.63**	**0.6206**	**0.6361**

Table B16 provides the complementary complete Lock comparison. EBLTR leads the reconstruction comparison at all effective-sampling factors, extending the spatial-evidence advantage to the lock screen's broad bands, regions, and principal divisions.

The Lock result also clarifies what remains recoverable under reduced sampling. Large colour and layout components retain stronger transport support than fine details, and EBLTR's LTOR–EBL chain is designed around exactly this ordering.

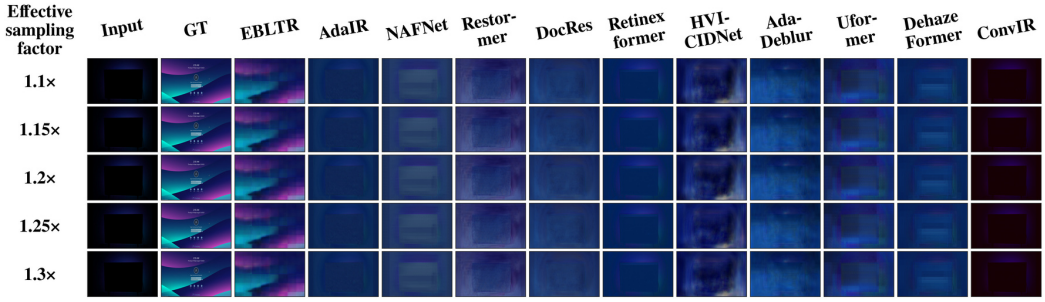

Figure B20: Representative Lock reconstructions across the effective-sampling sequence.

Figure B20 shows the corresponding Lock sequence. The broad colour bands, regional divisions, and main outline remain recoverable across the factors, while their changing clarity follows the changing spatial evidence.

Together with Table B16, the sequence demonstrates preliminary hidden-screen recovery across the tested effective resolutions and connects the result directly to EBLTR's evidence-dependent detail budget.

B.2.3 Gaussian Observation Noise

Observation noise adds local wall fluctuations that compete with the screen-derived signal. LTOR emphasizes the responses that remain coherent across space and frequency, while EBL keeps detail corrections tied to this retained evidence.

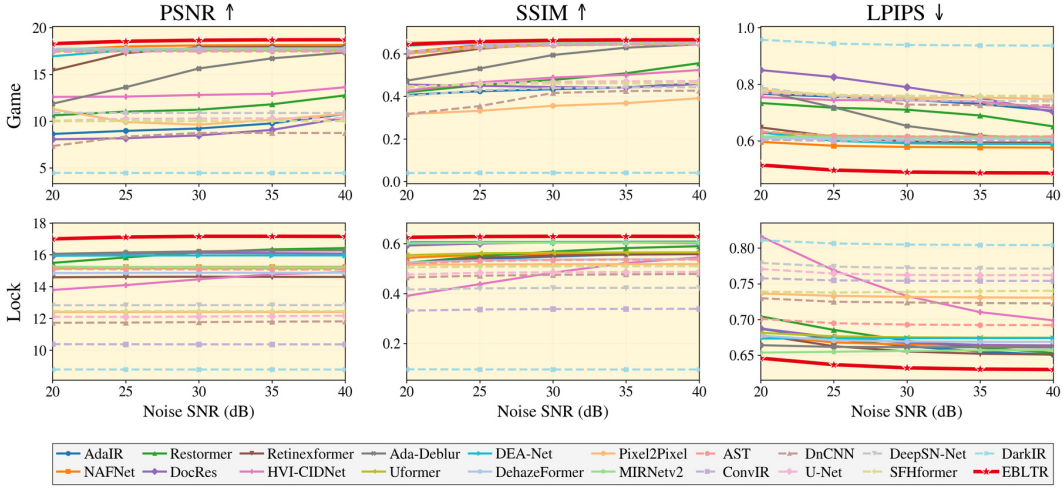

Figure B21: Game and Lock robustness curves under the Gaussian-noise levels.

EBLTR achieves the highest PSNR and lowest LPIPS across the noise conditions on Game and Lock in Figure B21. These gains accompany the preservation of screen structure when local wall fluctuations compete with the transported signal.

Noise makes the evidence-selection problem explicit. LTOR combines frequency, local statistics, and multilevel context to organize screen-related responses, and EBL restricts the detail budget

where the resulting evidence is weaker. The full-test-set comparison shows how this evidence-budgeted model recovers screen structure under noise.

Table B17: Complete Game comparison under gaussian observation noise. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Input SNR (dB) 20			Input SNR (dB) 25			Input SNR (dB) 30			Input SNR (dB) 35			Input SNR (dB) 40		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	7.34	0.3170	0.7689	8.35	0.3545	0.7598	8.75	0.4160	0.7280	8.74	0.4253	0.7259	8.74	0.4268	0.7256
U-Net	10.03	0.4415	0.7716	10.26	0.4616	0.7602	10.28	0.4689	0.7563	10.30	0.4712	0.7564	10.32	0.4722	0.7566
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	8.63	0.4068	0.7662	8.96	0.4243	0.7559	9.22	0.4336	0.7460	9.76	0.4424	0.7300	10.78	0.4615	0.7058
NAFNet	17.59	0.6087	0.5962	17.99	0.6411	0.5826	18.10	0.6526	0.5783	18.13	0.6562	0.5767	18.14	0.6577	0.5761
Restormer	10.62	0.4195	0.7337	11.02	0.4531	0.7167	11.22	0.4774	0.7099	11.80	0.5085	0.6895	12.76	0.5561	0.6509
Uformer	17.53	0.6018	0.6300	17.54	0.6303	0.6173	17.54	0.6411	0.6144	17.55	0.6449	0.6141	17.55	0.6463	0.6142
MIRNetv2	17.61	0.6533	0.6155	17.67	0.6555	0.6134	17.68	0.6559	0.6126	17.68	0.6559	0.6123	17.68	0.6559	0.6121
AST	17.42	0.5994	0.6327	17.45	0.6283	0.6186	17.46	0.6393	0.6157	17.46	0.6431	0.6156	17.46	0.6445	0.6157
ConvIR	17.58	0.6084	0.6037	17.62	0.6357	0.5999	17.64	0.6459	0.6003	17.64	0.6495	0.6009	17.64	0.6508	0.6012
DeepSN-Net	10.90	0.4014	0.7877	10.88	0.4285	0.7562	10.88	0.4386	0.7444	10.87	0.4421	0.7408	10.87	0.4434	0.7396
SFHformer	9.99	0.4472	0.7793	10.00	0.4570	0.7630	10.01	0.4604	0.7589	10.01	0.4614	0.7590	10.01	0.4618	0.7594
Models for Related Degradations															
DocRes	8.06	0.4554	0.8498	8.18	0.4498	0.8253	8.48	0.4437	0.7895	9.07	0.4411	0.7510	10.37	0.4581	0.7018
Retinexformer	15.40	0.5790	0.6472	17.27	0.6240	0.6148	17.90	0.6492	0.5981	17.92	0.6537	0.5943	17.92	0.6550	0.5932
HVI-CIDNet	12.60	0.4267	0.7541	12.63	0.4660	0.7439	12.81	0.4885	0.7432	12.93	0.5012	0.7386	13.63	0.5238	0.7160
Ada-Deblur	11.87	0.4737	0.7765	13.63	0.5312	0.7184	15.63	0.5945	0.6525	16.71	0.6288	0.6190	17.34	0.6448	0.6024
DEA-Net	16.94	0.6380	0.6294	17.58	0.6538	0.6000	17.69	0.6568	0.5913	17.72	0.6576	0.5889	17.73	0.6579	0.5883
DehazeFormer	17.75	0.6442	0.6079	17.77	0.6476	0.6073	17.77	0.6488	0.6072	17.77	0.6491	0.6073	17.77	0.6493	0.6073
Pixel2Pixel	11.29	0.3167	0.7791	9.91	0.3322	0.7599	9.63	0.3557	0.7513	10.15	0.3682	0.7531	10.78	0.3913	0.7465
DarkIR	4.47	0.0401	0.9571	4.47	0.0407	0.9434	4.47	0.0408	0.9386	4.47	0.0409	0.9370	4.46	0.0409	0.9364
EBLTR	**18.28**	**0.6449**	**0.5145**	**18.54**	**0.6578**	**0.4966**	**18.65**	**0.6637**	**0.4898**	**18.69**	**0.6657**	**0.4872**	**18.70**	**0.6663**	**0.4864**

Table B17 confirms the Game comparison across every method and noise level. EBLTR leads PSNR and LPIPS throughout, linking evidence selection to stronger preservation of the game's playing area and regional organization.

Each noisy wall input is shared across methods, so Table B17 compares how effectively they recover the same hidden screen. EBLTR gains its advantage by distinguishing usable transported structure from competing perturbation and by matching refinement magnitude to that distinction.

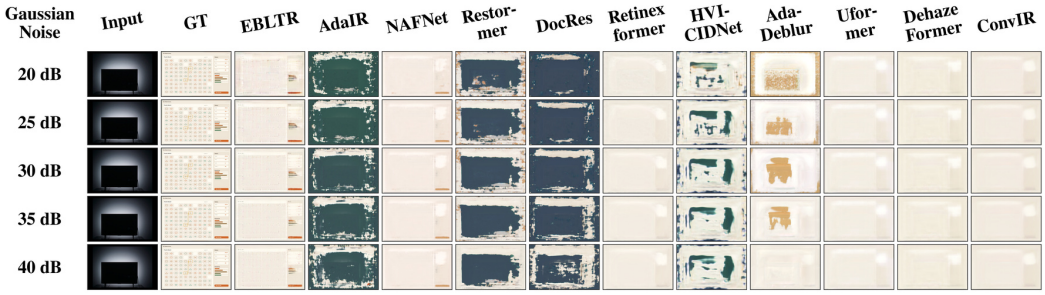

Figure B22: Representative Game reconstructions under the Gaussian-noise sequence.

The light-background game layout and its side panel remain identifiable in the EBLTR reconstructions across the noise sequence in Figure B22. The example shows how the principal target regions can persist despite competing fluctuations in the wall observation. The example also illustrates the role of budgeting: broad evidence continues to guide layout recovery, while detail refinement follows the strength of the local signal. Across the evaluated noise conditions, EBLTR preserves the screen's main spatial organization while maintaining its reconstruction advantage.

Table B18: Complete Lock comparison under gaussian observation noise. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Input SNR (dB) 20			Input SNR (dB) 25			Input SNR (dB) 30			Input SNR (dB) 35			Input SNR (dB) 40		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.73	0.4646	0.7298	11.74	0.4714	0.7249	11.77	0.4748	0.7238	11.80	0.4769	0.7231	11.81	0.4781	0.7225
U-Net	12.10	0.4759	0.7707	12.10	0.4833	0.7644	12.11	0.4850	0.7623	12.15	0.4861	0.7621	12.17	0.4869	0.7621
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.18	0.5213	0.6870	15.17	0.5375	0.6726	15.16	0.5483	0.6626	15.14	0.5572	0.6552	15.15	0.5644	0.6506
NAFNet	15.22	0.5435	0.6774	15.23	0.5550	0.6679	15.24	0.5595	0.6650	15.24	0.5611	0.6641	15.24	0.5617	0.6638
Restormer	15.49	0.5229	0.7044	15.84	0.5495	0.6855	16.16	0.5684	0.6703	16.34	0.5825	0.6606	16.42	0.5897	0.6533
Uformer	15.15	0.5539	0.6815	15.15	0.5612	0.6765	15.15	0.5638	0.6750	15.15	0.5648	0.6746	15.15	0.5651	0.6744
MIRNetv2	15.25	0.6008	0.6537	15.23	0.6035	0.6549	15.22	0.6040	0.6561	15.21	0.6041	0.6569	15.21	0.6042	0.6573
AST	15.11	0.5200	0.7009	15.09	0.5305	0.6948	15.08	0.5344	0.6929	15.08	0.5358	0.6923	15.07	0.5363	0.6921
ConvIR	10.38	0.3312	0.7577	10.37	0.3360	0.7548	10.37	0.3377	0.7542	10.37	0.3382	0.7541	10.37	0.3384	0.7541
DeepSN-Net	12.83	0.4160	0.7792	12.83	0.4205	0.7740	12.84	0.4221	0.7722	12.84	0.4226	0.7715	12.84	0.4228	0.7712
SFHformer	12.46	0.5047	0.7390	12.46	0.5079	0.7377	12.46	0.5089	0.7388	12.46	0.5093	0.7397	12.46	0.5094	0.7402
Models for Related Degradations															
DocRes	15.92	0.5924	0.6875	16.09	0.6004	0.6749	16.12	0.6057	0.6677	16.11	0.6047	0.6646	16.08	0.6016	0.6637
Retinexformer	14.59	0.5206	0.6783	14.61	0.5430	0.6627	14.62	0.5524	0.6553	14.62	0.5560	0.6521	14.61	0.5578	0.6510
HVI-CIDNet	13.79	0.3904	0.8165	14.09	0.4375	0.7688	14.46	0.4835	0.7328	14.78	0.5224	0.7103	14.92	0.5454	0.6988
Ada-Deblur	16.06	0.6019	0.6641	16.15	0.6043	0.6617	16.21	0.6063	0.6613	16.27	0.6077	0.6612	16.28	0.6084	0.6610
DEA-Net	15.94	0.6061	0.6733	15.96	0.6067	0.6737	15.96	0.6069	0.6739	15.96	0.6069	0.6740	15.96	0.6069	0.6741
DehazeFormer	14.85	0.5239	0.6775	14.87	0.5327	0.6711	14.87	0.5359	0.6693	14.87	0.5371	0.6689	14.87	0.5374	0.6688
Pixel2Pixel	12.40	0.5152	0.7364	12.40	0.5167	0.7331	12.40	0.5174	0.7316	12.40	0.5177	0.7309	12.41	0.5178	0.7306
DarkIR	8.79	0.0957	0.8109	8.79	0.0953	0.8064	8.79	0.0951	0.8048	8.79	0.0950	0.8043	8.79	0.0949	0.8041
EBLTR	**17.00**	**0.6254**	**0.6459**	**17.12**	**0.6287**	**0.6368**	**17.16**	**0.6296**	**0.6324**	**17.17**	**0.6296**	**0.6306**	**17.16**	**0.6295**	**0.6300**

Table B18 extends the noise result to the complete Lock comparison. EBLTR leads the reconstruction results at all levels, showing that the evidence-aware advantage applies to broad lock-screen bands and layout divisions as well as the denser Game structure. This cross-content agreement reinforces the central claim. A noisy wall spot still retains recoverable screen organization, and EBLTR makes stronger use of that organization by coupling observability with local correction capacity.

Figure B23 supplies the complementary Lock sequence. The recovered colour bands, regional divisions, and main outline remain visible through the noise levels, revealing the broad screen evidence that survives the perturbation. The figure and Table B18 together show that this visible structure follows a complete test-set advantage. EBLTR therefore continues to demonstrate preliminary screen-content recovery when observation noise weakens the wall trace.

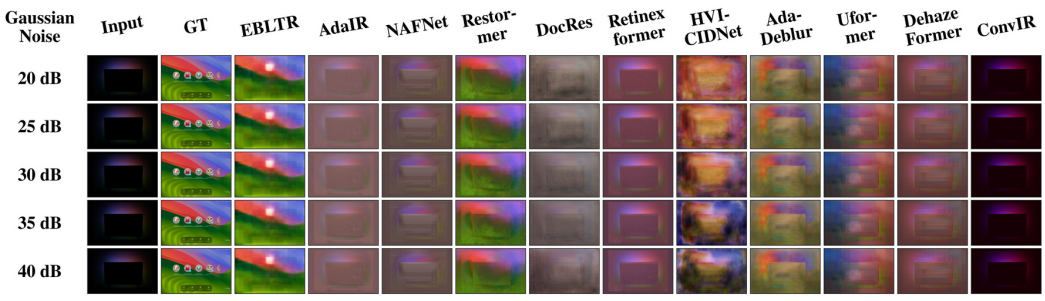

Figure B23: Representative Lock reconstructions under the Gaussian-noise sequence.

B.2.4 Camera Jitter

Camera jitter shifts local wall responses while the hidden display remains the same. The robustness sequence tests whether screen structure remains recoverable through this spatial change; LTOR and EBL preserve the evidence-to-budget route that anchors recovery to the screen layout.

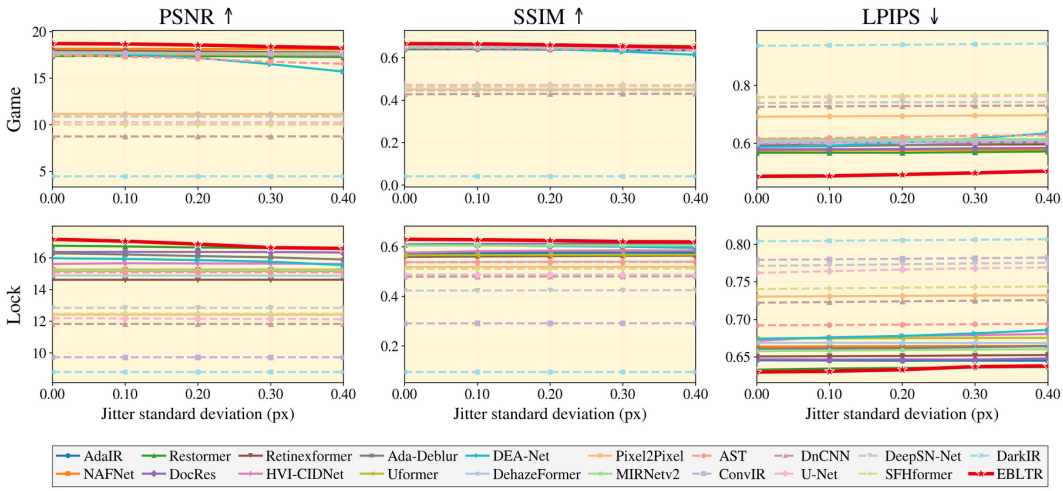

Figure B24: Game and Lock robustness curves under the camera-jitter sequence.

EBLTR achieves the highest PSNR at every jitter level on Game and Lock in Figure B24, together with the lowest LPIPS on Game and highest SSIM on Lock. The recovered screen remains organized around the same principal regions as the wall image shifts. This perturbation tests spatially adaptive screen recovery. LTOR re-evaluates the shifted wall pattern through local and multilevel context, and EBL uses the resulting map to preserve the screen's major regions and boundaries. The curves expose that adaptation across the sequence.



Figure B25: Representative Game reconstructions under the camera-jitter sequence.

The light-background playing area in Figure B25 retains its position relative to the side panel as the wall image shifts. EBLTR thus connects displaced observations to the same screen arrangement.

Paired with Table B19, the figure shows that LTOR recomputes observability from the shifted wall input and EBL converts it into a stable screen-level reconstruction across the test set.

Table B19: Complete Game comparison under camera jitter. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Jitter sigma (pixels) 0.0			Jitter sigma (pixels) 0.1			Jitter sigma (pixels) 0.2			Jitter sigma (pixels) 0.3			Jitter sigma (pixels) 0.4		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.74	0.4273	0.7259	8.74	0.4285	0.7272	8.74	0.4293	0.7284	8.74	0.4297	0.7292	8.75	0.4300	0.7299
U-Net	10.29	0.4715	0.7581	10.27	0.4714	0.7597	10.26	0.4710	0.7613	10.25	0.4706	0.7626	10.23	0.4703	0.7642
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.44	0.6394	0.5754	17.44	0.6403	0.5763	17.39	0.6381	0.5788	17.34	0.6361	0.5810	17.29	0.6344	0.5834
NAFNet	18.16	0.6587	0.5759	18.15	0.6583	0.5762	18.13	0.6575	0.5767	18.11	0.6568	0.5771	18.06	0.6554	0.5784
Restormer	17.37	0.6438	0.5676	17.38	0.6440	0.5677	17.37	0.6441	0.5677	17.32	0.6434	0.5695	17.25	0.6415	0.5716
Uformer	17.55	0.6470	0.6142	17.54	0.6473	0.6147	17.54	0.6473	0.6144	17.51	0.6468	0.6140	17.51	0.6465	0.6134
MIRNetv2	17.68	0.6559	0.6119	17.68	0.6559	0.6128	17.67	0.6560	0.6134	17.67	0.6560	0.6142	17.67	0.6560	0.6148
AST	17.46	0.6452	0.6158	17.29	0.6444	0.6187	17.10	0.6424	0.6210	16.74	0.6395	0.6260	16.54	0.6360	0.6271
ConvIR	17.64	0.6514	0.6013	17.64	0.6517	0.6022	17.63	0.6519	0.6031	17.62	0.6518	0.6038	17.60	0.6516	0.6046
DeepSN-Net	10.87	0.4440	0.7390	10.87	0.4451	0.7404	10.88	0.4460	0.7416	10.88	0.4465	0.7421	10.88	0.4466	0.7419
SFHformer	10.01	0.4620	0.7598	10.01	0.4629	0.7618	10.01	0.4638	0.7640	10.01	0.4645	0.7659	10.01	0.4650	0.7675
Models for Related Degradations															
DocRes	17.53	0.6472	0.5794	17.52	0.6466	0.5794	17.50	0.6449	0.5803	17.46	0.6431	0.5829	17.45	0.6410	0.5840
Retinexformer	17.92	0.6556	0.5924	17.90	0.6547	0.5930	17.87	0.6536	0.5938	17.80	0.6514	0.5953	17.75	0.6498	0.5962
HVI-CIDNet	17.77	0.6555	0.5978	17.68	0.6521	0.6022	17.65	0.6505	0.6042	17.62	0.6497	0.6053	17.55	0.6474	0.6080
Ada-Deblur	17.75	0.6577	0.5878	17.70	0.6566	0.5899	17.62	0.6543	0.5933	17.48	0.6509	0.5986	17.33	0.6470	0.6032
DEA-Net	17.73	0.6579	0.5879	17.55	0.6513	0.5917	17.17	0.6406	0.6026	16.50	0.6286	0.6169	15.70	0.6136	0.6345
DehazeFormer	17.77	0.6494	0.6074	17.74	0.6484	0.6085	17.67	0.6473	0.6099	17.61	0.6462	0.6114	17.55	0.6448	0.6122
Pixel2Pixel	11.14	0.4502	0.6917	11.14	0.4501	0.6929	11.13	0.4498	0.6943	11.13	0.4500	0.6956	11.13	0.4502	0.6967
DarkIR	4.46	0.0409	0.9361	4.46	0.0409	0.9378	4.46	0.0410	0.9397	4.46	0.0410	0.9418	4.46	0.0410	0.9437
EBLTR	**18.71**	**0.6665**	**0.4864**	**18.67**	**0.6652**	**0.4877**	**18.55**	**0.6604**	**0.4925**	**18.39**	**0.6545**	**0.4980**	**18.24**	**0.6494**	**0.5044**

Table B19 confirms the Game result against all baselines at every jitter level. EBLTR leads PSNR and LPIPS across the sequence, showing the advantage of spatially organized observability as the screen trace moves across the wall image. Evidence location can change while the recovery target remains stable, so estimating where useful evidence resides supports a consistent screen reconstruction. Table B19 provides the quantitative foundation for Figure B25, linking the visible layout stability to the complete fidelity, structure, and perceptual comparison.

Table B20: Complete Lock comparison under camera jitter. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Jitter sigma (pixels) 0			Jitter sigma (pixels) 0.1			Jitter sigma (pixels) 0.2			Jitter sigma (pixels) 0.3			Jitter sigma (pixels) 0.4		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4787	0.7221	11.82	0.4793	0.7230	11.82	0.4798	0.7239	11.82	0.4802	0.7248	11.82	0.4805	0.7256
U-Net	12.18	0.4874	0.7618	12.17	0.4874	0.7641	12.14	0.4872	0.7663	12.13	0.4870	0.7680	12.11	0.4869	0.7690
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.17	0.5729	0.6455	15.16	0.5740	0.6451	15.16	0.5749	0.6449	15.16	0.5755	0.6450	15.16	0.5761	0.6449
NAFNet	15.24	0.5619	0.6637	15.24	0.5627	0.6639	15.24	0.5634	0.6641	15.24	0.5640	0.6644	15.24	0.5644	0.6646
Restormer	16.75	0.6070	0.6329	16.71	0.6075	0.6345	16.64	0.6081	0.6352	16.60	0.6075	0.6359	16.56	0.6075	0.6363
Uformer	15.15	0.5651	0.6744	15.16	0.5661	0.6746	15.17	0.5668	0.6748	15.17	0.5673	0.6752	15.17	0.5675	0.6753
MIRNetv2	15.21	0.6042	0.6575	15.21	0.6044	0.6583	15.21	0.6046	0.6590	15.21	0.6048	0.6597	15.21	0.6049	0.6601
AST	15.07	0.5365	0.6920	15.07	0.5372	0.6924	15.07	0.5379	0.6929	15.07	0.5383	0.6935	15.07	0.5384	0.6938
ConvIR	9.72	0.2906	0.7792	9.72	0.2909	0.7799	9.72	0.2911	0.7806	9.72	0.2912	0.7815	9.72	0.2914	0.7822
DeepSN-Net	12.84	0.4228	0.7711	12.84	0.4233	0.7722	12.84	0.4237	0.7733	12.83	0.4240	0.7745	12.83	0.4243	0.7752
SFHformer	12.46	0.5095	0.7404	12.46	0.5099	0.7413	12.46	0.5103	0.7422	12.46	0.5107	0.7431	12.46	0.5110	0.7436
Models for Related Degradations															
DocRes	16.36	0.6094	0.6467	16.36	0.6110	0.6463	16.36	0.6112	0.6465	16.35	0.6115	0.6464	16.32	0.6108	0.6479
Retinexformer	14.60	0.5586	0.6506	14.60	0.5603	0.6508	14.60	0.5617	0.6512	14.61	0.5628	0.6519	14.61	0.5636	0.6524
HVI-CIDNet	15.61	0.5758	0.6714	15.64	0.5813	0.6761	15.63	0.5833	0.6778	15.62	0.5847	0.6790	15.61	0.5840	0.6804
Ada-Deblur	16.27	0.6086	0.6608	16.20	0.6061	0.6612	16.10	0.6030	0.6616	16.02	0.6009	0.6630	15.88	0.5960	0.6641
DEA-Net	15.96	0.6069	0.6741	15.92	0.6052	0.6755	15.84	0.6025	0.6779	15.75	0.5998	0.6813	15.53	0.5941	0.6859
DehazeFormer	14.87	0.5375	0.6687	14.87	0.5383	0.6686	14.87	0.5390	0.6686	14.86	0.5396	0.6685	14.86	0.5398	0.6683
Pixel2Pixel	12.41	0.5178	0.7302	12.41	0.5180	0.7307	12.41	0.5182	0.7312	12.41	0.5183	0.7317	12.41	0.5184	0.7320
DarkIR	8.79	0.0949	0.8040	8.79	0.0950	0.8046	8.79	0.0950	0.8053	8.79	0.0950	0.8060	8.79	0.0950	0.8066
EBLTR	**17.16**	**0.6295**	**0.6298**	**17.04**	**0.6277**	**0.6308**	**16.85**	**0.6245**	**0.6329**	**16.63**	**0.6202**	**0.6370**	**16.58**	**0.6191**	**0.6380**

Table B20 extends the same spatial-tracking result to Lock. EBLTR achieves the highest PSNR and SSIM throughout the jitter sequence, showing that its advantage carries from dense game structure to broad lock-screen regions. The shared direction across both content categories supports recovery from shifted wall responses: LTOR recomputes observability, and EBL uses it to bound refinement of the same screen target.

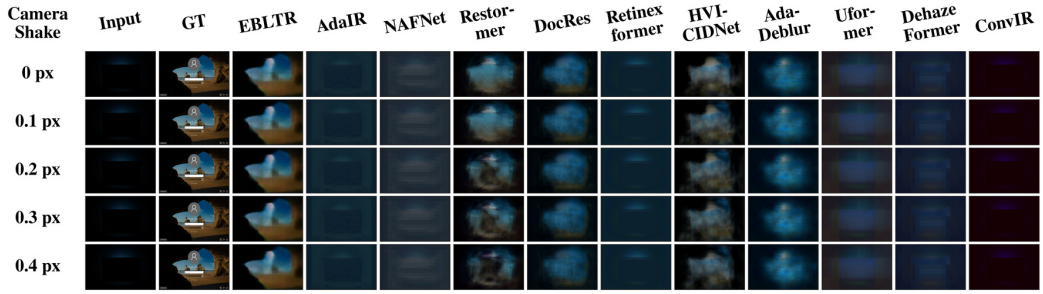

Figure B26: Representative Lock reconstructions under the camera-jitter sequence.

Figure B26 supplies the complementary Lock sequence. The colour bands, regional divisions, and main outline remain identifiable as the wall response shifts, revealing stable recovered content across the tested jitter conditions. Together with the full Lock table, the figure establishes both visible screen structure and metric advantage, strengthening the feasibility of ordinary-camera recovery under the spatial perturbation.

B.2.5 Horizontal Viewing Angle

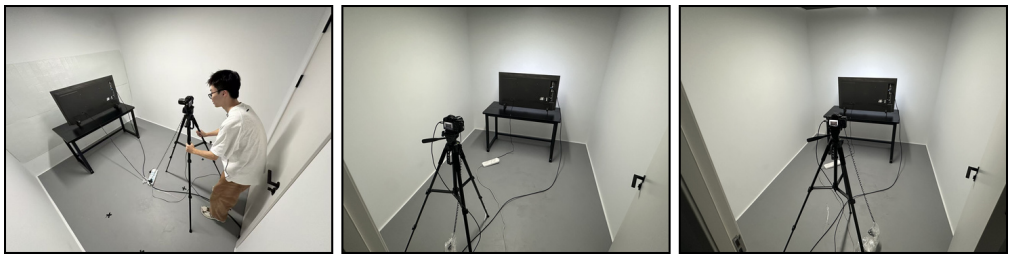

Figure B27: Acquisition views illustrating changes in horizontal observation angle.

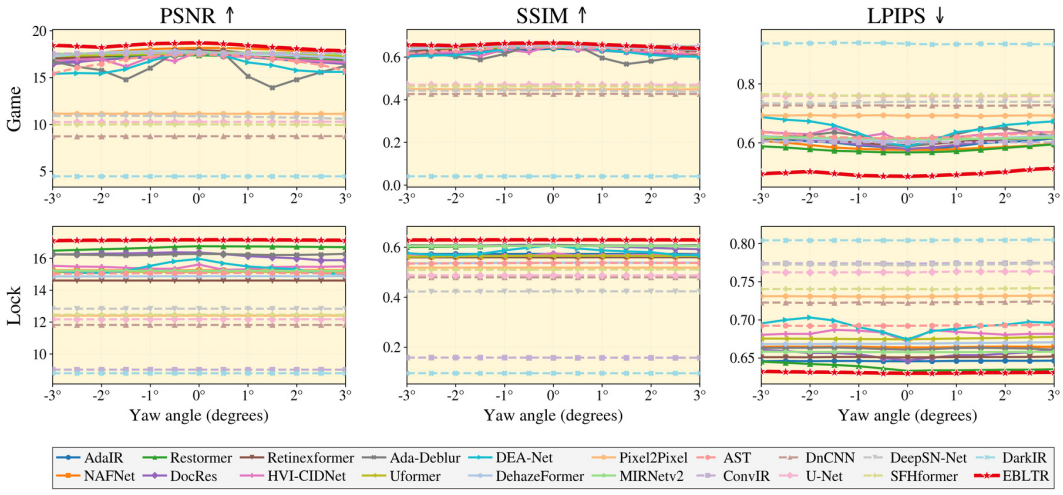

Figure B28: Game and Lock robustness curves under the horizontal-viewing-angle sequence.

EBLTR maintains the highest PSNR and lowest LPIPS across all thirteen horizontal viewing angles on both content categories in Figure B28. Screen recovery therefore remains effective as the observed spatial arrangement of the wall response changes. The result highlights the spatial component of LTOR. The viewing-angle sequence changes the observed arrangement of colour and boundary cues; LTOR represents each transformed wall response, and EBL uses its observability to set local correction limits.

Table B21: Complete Game comparison under horizontal viewing angle. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	-3.0°			-2.5°			-2.0°			-1.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	8.72	0.4263	0.7265	8.73	0.4268	0.7268	8.73	0.4270	0.7265	8.73	0.4272	0.7252
U-Net	10.25	0.4699	0.7593	10.25	0.4699	0.7588	10.25	0.4700	0.7591	10.25	0.4704	0.7595
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	16.50	0.6072	0.6147	16.70	0.6125	0.6107	16.91	0.6173	0.6047	17.04	0.6218	0.6006
NAFNet	17.21	0.6320	0.6064	17.41	0.6377	0.6003	17.66	0.6445	0.5927	17.89	0.6508	0.5841
Restormer	17.02	0.6324	0.5882	17.13	0.6348	0.5842	17.22	0.6379	0.5782	17.28	0.6407	0.5729
Uformer	17.25	0.6391	0.6170	17.35	0.6414	0.6156	17.40	0.6426	0.6140	17.45	0.6438	0.6125
MIRNetv2	17.53	0.6495	0.6187	17.54	0.6502	0.6188	17.56	0.6514	0.6183	17.60	0.6531	0.6166
AST	15.43	0.6158	0.6359	16.05	0.6250	0.6295	16.45	0.6312	0.6250	16.97	0.6384	0.6194
ConvIR	17.56	0.6486	0.6050	17.58	0.6492	0.6041	17.58	0.6493	0.6038	17.59	0.6496	0.6032
DeepSN-Net	10.91	0.4419	0.7342	10.92	0.4417	0.7336	10.92	0.4417	0.7331	10.92	0.4423	0.7337
SFHformer	9.97	0.4616	0.7649	9.98	0.4625	0.7649	9.99	0.4628	0.7631	10.00	0.4620	0.7601
Models for Related Degradations												
DocRes	16.43	0.6029	0.6232	16.64	0.6098	0.6178	16.92	0.6208	0.6086	17.06	0.6264	0.6028
Retinexformer	17.00	0.6278	0.6222	17.09	0.6292	0.6191	17.25	0.6328	0.6150	17.47	0.6401	0.6092
HVI-CIDNet	16.71	0.6157	0.6376	16.93	0.6194	0.6319	16.95	0.6238	0.6294	16.17	0.6050	0.6503
Ada-Deblur	16.74	0.6276	0.6095	16.16	0.6141	0.6199	15.79	0.6028	0.6261	14.77	0.5880	0.6357
DEA-Net	15.40	0.6023	0.6875	15.48	0.6068	0.6784	15.43	0.6084	0.6731	15.91	0.6199	0.6589
DehazeFormer	17.24	0.6358	0.6226	17.51	0.6422	0.6178	17.68	0.6467	0.6134	17.76	0.6492	0.6107
Pixel2Pixel	11.15	0.4495	0.6932	11.15	0.4488	0.6924	11.14	0.4481	0.6938	11.14	0.4495	0.6920
DarkIR	4.46	0.0408	0.9373	4.46	0.0408	0.9370	4.46	0.0409	0.9378	4.46	0.0409	0.9392
EBLTR	**18.44**	**0.6573**	**0.4951**	**18.35**	**0.6560**	**0.4989**	**18.23**	**0.6526**	**0.5027**	**18.42**	**0.6582**	**0.4961**

Table B21 (continued): -1.0° to 0.5°

Method	-1.0°			-0.5°			0.0°			0.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	8.73	0.4276	0.7264	8.74	0.4275	0.7269	8.74	0.4273	0.7259	8.74	0.4274	0.7263
U-Net	10.26	0.4702	0.7595	10.27	0.4707	0.7583	10.29	0.4715	0.7581	10.28	0.4706	0.7575
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	17.18	0.6264	0.5917	17.32	0.6327	0.5851	17.44	0.6394	0.5754	17.33	0.6344	0.5798
NAFNet	18.04	0.6552	0.5794	18.11	0.6573	0.5772	18.16	0.6587	0.5759	18.16	0.6584	0.5763
Restormer	17.33	0.6414	0.5707	17.38	0.6433	0.5679	17.37	0.6438	0.5676	17.33	0.6410	0.5684
Uformer	17.49	0.6453	0.6125	17.52	0.6464	0.6136	17.55	0.6470	0.6142	17.57	0.6468	0.6140
MIRNetv2	17.64	0.6546	0.6147	17.66	0.6554	0.6133	17.68	0.6559	0.6119	17.68	0.6560	0.6121
AST	17.25	0.6423	0.6167	17.42	0.6446	0.6156	17.46	0.6452	0.6158	17.41	0.6439	0.6171
ConvIR	17.62	0.6505	0.6021	17.64	0.6513	0.6015	17.64	0.6514	0.6013	17.63	0.6512	0.6014
DeepSN-Net	10.91	0.4435	0.7359	10.89	0.4441	0.7378	10.87	0.4440	0.7390	10.84	0.4433	0.7395
SFHformer	10.01	0.4622	0.7604	10.01	0.4622	0.7606	10.01	0.4620	0.7598	10.00	0.4629	0.7612
Models for Related Degradations												
DocRes	17.25	0.6345	0.5953	17.44	0.6421	0.5861	17.53	0.6472	0.5794	17.44	0.6431	0.5847
Retinexformer	17.68	0.6466	0.6015	17.87	0.6533	0.5941	17.92	0.6556	0.5924	17.84	0.6527	0.5958
HVI-CIDNet	17.18	0.6342	0.6184	16.76	0.6223	0.6312	17.77	0.6555	0.5978	17.23	0.6324	0.6189
Ada-Deblur	16.01	0.6142	0.6247	17.64	0.6522	0.5943	17.75	0.6577	0.5878	17.61	0.6517	0.5963
DEA-Net	16.75	0.6356	0.6326	17.49	0.6525	0.6018	17.73	0.6579	0.5879	17.35	0.6502	0.6081
DehazeFormer	17.81	0.6506	0.6087	17.83	0.6508	0.6075	17.77	0.6494	0.6074	17.64	0.6464	0.6092
Pixel2Pixel	11.14	0.4492	0.6927	11.14	0.4482	0.6936	11.14	0.4502	0.6917	11.14	0.4485	0.6923
DarkIR	4.46	0.0409	0.9391	4.46	0.0409	0.9385	4.46	0.0409	0.9361	4.46	0.0409	0.9336
EBLTR	**18.62**	**0.6641**	**0.4893**	**18.66**	**0.6655**	**0.4875**	**18.71**	**0.6665**	**0.4864**	**18.61**	**0.6638**	**0.4885**

Table B21 (continued): 1.0° to 3.0°

Method	1.0°			1.5°			2.0°			2.5°			3.0°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.74	0.4275	0.7257	8.74	0.4273	0.7256	8.74	0.4271	0.7260	8.74	0.4268	0.7267	8.75	0.4265	0.7277
U-Net	10.27	0.4703	0.7583	10.28	0.4704	0.7586	10.29	0.4702	0.7583	10.29	0.4700	0.7582	10.28	0.4699	0.7594
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.20	0.6297	0.5881	16.98	0.6227	0.5985	16.78	0.6169	0.6030	16.77	0.6163	0.6067	16.47	0.6062	0.6151
NAFNet	18.08	0.6554	0.5788	17.97	0.6514	0.5820	17.85	0.6474	0.5858	17.65	0.6426	0.5905	17.42	0.6379	0.5966
Restormer	17.26	0.6385	0.5713	17.16	0.6359	0.5768	17.00	0.6311	0.5829	16.88	0.6273	0.5886	16.79	0.6243	0.5947
Uformer	17.55	0.6462	0.6138	17.52	0.6456	0.6138	17.45	0.6448	0.6148	17.37	0.6439	0.6157	17.28	0.6429	0.6168
MIRNetv2	17.67	0.6556	0.6128	17.65	0.6545	0.6135	17.62	0.6530	0.6146	17.57	0.6511	0.6166	17.51	0.6493	0.6187
AST	17.29	0.6408	0.6196	17.02	0.6354	0.6236	16.72	0.6309	0.6276	16.35	0.6256	0.6316	15.70	0.6173	0.6359
ConvIR	17.62	0.6508	0.6021	17.61	0.6505	0.6021	17.62	0.6506	0.6016	17.63	0.6508	0.6015	17.62	0.6509	0.6017
DeepSN-Net	10.81	0.4418	0.7401	10.76	0.4404	0.7395	10.71	0.4395	0.7392	10.66	0.4381	0.7390	10.60	0.4369	0.7393
SFHformer	9.99	0.4629	0.7613	9.98	0.4620	0.7601	9.97	0.4615	0.7609	9.96	0.4609	0.7616	9.94	0.4601	0.7615
Models for Related Degradations															
DocRes	17.28	0.6367	0.5932	17.05	0.6272	0.6034	16.82	0.6195	0.6121	16.62	0.6114	0.6191	16.44	0.6046	0.6258
Retinexformer	17.63	0.6445	0.6025	17.44	0.6388	0.6086	17.27	0.6338	0.6123	17.15	0.6308	0.6149	17.00	0.6274	0.6197
HVI-CIDNet	17.23	0.6314	0.6196	17.07	0.6249	0.6277	16.85	0.6245	0.6314	16.76	0.6182	0.6358	16.66	0.6128	0.6361
Ada-Deblur	15.14	0.5956	0.6293	13.93	0.5667	0.6482	14.78	0.5811	0.6495	15.58	0.5983	0.6352	16.24	0.6134	0.6217
DEA-Net	16.62	0.6331	0.6361	16.31	0.6232	0.6453	15.77	0.6091	0.6604	15.60	0.6031	0.6673	15.60	0.6009	0.6728
DehazeFormer	17.54	0.6443	0.6117	17.45	0.6420	0.6145	17.35	0.6391	0.6177	17.22	0.6365	0.6205	17.05	0.6330	0.6240
Pixel2Pixel	11.13	0.4479	0.6924	11.13	0.4487	0.6907	11.13	0.4479	0.6922	11.13	0.4477	0.6930	11.13	0.4485	0.6924
DarkIR	4.46	0.0409	0.9346	4.46	0.0409	0.9359	4.46	0.0409	0.9353	4.46	0.0409	0.9346	4.46	0.0408	0.9341
EBLTR	**18.42**	**0.6582**	**0.4926**	**18.25**	**0.6534**	**0.4963**	**18.09**	**0.6489**	**0.5017**	**17.93**	**0.6432**	**0.5096**	**17.86**	**0.6404**	**0.5138**

Table B21 confirms the Game result against every baseline and horizontal angle. EBLTR leads PSNR and LPIPS throughout, showing that explicit observability organization provides a stronger route from angle-dependent wall evidence to the hidden game's regional layout. The complete factor matrix establishes a consistent dataset-level advantage. Each angle changes the observation, while the evidence-budgeted design continues to recover the target's major structure more faithfully across the viewing-angle comparison.

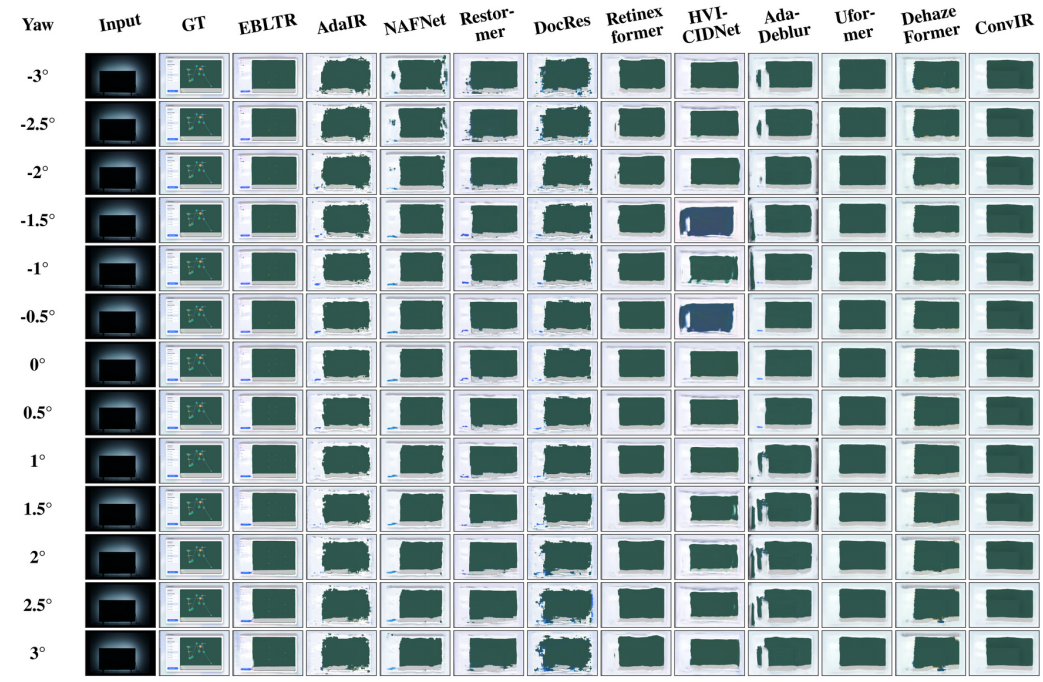

Figure B29: Representative Game reconstructions under the horizontal-viewing-angle sequence.

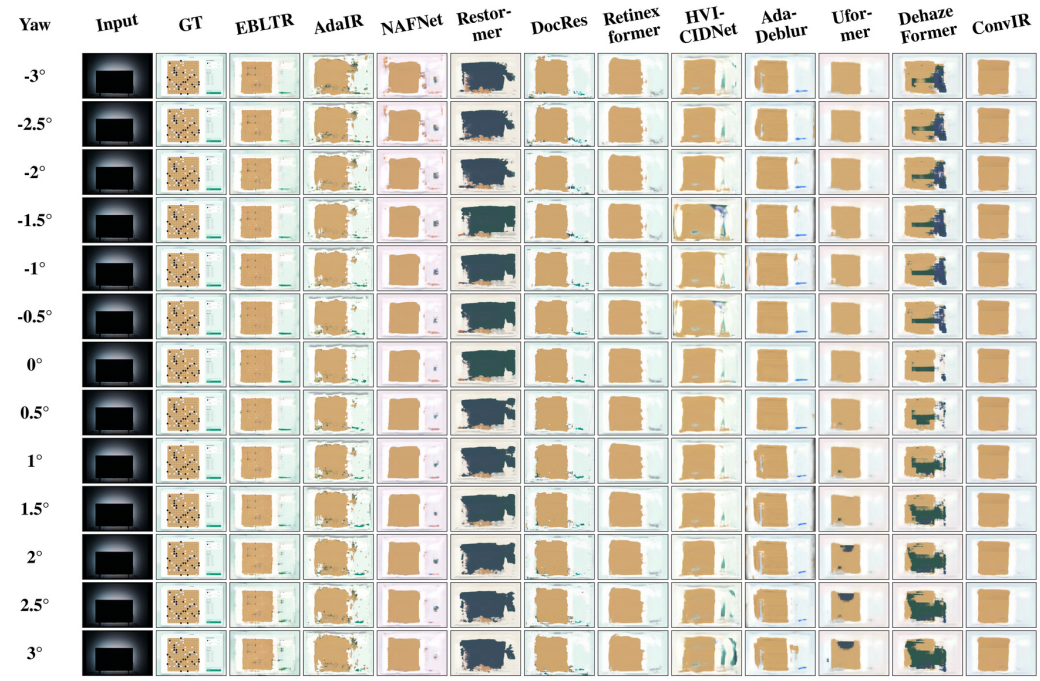

Figure B30: Representative Game reconstruction comparison under the horizontal-viewpoint sequence.

The green playing area in Figure B29 remains distinct from the surrounding interface across the horizontal viewpoints. Its recovered placement links angle-dependent wall observations to a consistent target layout. Table B21 anchors Figure B29 in the complete Game comparison, connecting the representative layout preservation to PSNR, SSIM, and LPIPS across all angles.

These are exactly the content components expected to survive diffuse transport most reliably. Their preservation shows LTOR selecting broad and boundary evidence at each angle and EBL refining the regions supported by that updated observation. Supplementary Figure B30 varies horizontal viewpoint for one Game example from negative to positive angles while holding the target content fixed. EBLTR retains the broad board placement across the displayed sequence, providing a focused geometry visualization that complements the complete quantitative comparison in Table B21.

Table B22: Complete Lock comparison under horizontal viewing angle. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	-3°			-2.5°			-2°			-1.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	11.82	0.4789	0.7225	11.82	0.4789	0.7225	11.82	0.4789	0.7220	11.82	0.4789	0.7221
U-Net	12.17	0.4871	0.7622	12.17	0.4871	0.7621	12.17	0.4870	0.7621	12.17	0.4870	0.7620
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	15.16	0.5727	0.6456	15.16	0.5725	0.6460	15.16	0.5726	0.6456	15.16	0.5726	0.6460
NAFNet	15.24	0.5607	0.6646	15.24	0.5609	0.6648	15.24	0.5612	0.6646	15.24	0.5613	0.6647
Restormer	16.48	0.6003	0.6451	16.52	0.6013	0.6436	16.56	0.6022	0.6420	16.60	0.6031	0.6408
Uformer	15.13	0.5614	0.6753	15.14	0.5622	0.6753	15.14	0.5631	0.6751	15.15	0.5639	0.6749
MIRNetv2	15.21	0.6042	0.6581	15.21	0.6042	0.6580	15.21	0.6042	0.6578	15.21	0.6042	0.6579
AST	15.02	0.5338	0.6921	15.03	0.5344	0.6922	15.04	0.5350	0.6919	15.05	0.5355	0.6921
ConvIR	9.01	0.1576	0.7742	9.01	0.1576	0.7741	9.01	0.1576	0.7738	9.01	0.1575	0.7738
DeepSN-Net	12.84	0.4222	0.7726	12.84	0.4223	0.7726	12.84	0.4224	0.7721	12.84	0.4225	0.7721
SFHformer	12.46	0.5086	0.7402	12.46	0.5089	0.7403	12.46	0.5090	0.7402	12.46	0.5091	0.7403
Models for Related Degradations												
DocRes	16.24	0.6041	0.6604	16.25	0.6056	0.6590	16.29	0.6070	0.6565	16.30	0.6063	0.6565
Retinexformer	14.60	0.5590	0.6508	14.60	0.5591	0.6509	14.60	0.5591	0.6507	14.60	0.5590	0.6511
HVI-CIDNet	15.53	0.5712	0.6803	15.48	0.5696	0.6814	15.46	0.5682	0.6815	15.40	0.5648	0.6868
Ada-Deblur	16.24	0.6077	0.6631	16.22	0.6065	0.6631	16.21	0.6061	0.6632	16.19	0.6059	0.6636
DEA-Net	15.24	0.5760	0.6951	15.09	0.5696	0.6998	15.07	0.5686	0.7029	15.20	0.5730	0.6989
DehazeFormer	14.86	0.5352	0.6681	14.86	0.5356	0.6684	14.86	0.5361	0.6682	14.87	0.5366	0.6683
Pixel2Pixel	12.40	0.5177	0.7308	12.40	0.5177	0.7309	12.40	0.5177	0.7307	12.40	0.5178	0.7306
DarkIR	8.79	0.0950	0.8040	8.79	0.0950	0.8041	8.79	0.0950	0.8042	8.79	0.0950	0.8042
EBLTR	**17.11**	**0.6285**	**0.6323**	**17.12**	**0.6287**	**0.6318**	**17.13**	**0.6289**	**0.6314**	**17.14**	**0.6290**	**0.6312**

Table B22 (continued): -1° to 0.5°

Method	-1°			-0.5°			0°			0.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	11.82	0.4788	0.7228	11.82	0.4788	0.7225	11.82	0.4787	0.7221	11.82	0.4787	0.7227
U-Net	12.18	0.4870	0.7622	12.18	0.4872	0.7623	12.18	0.4874	0.7618	12.18	0.4871	0.7627
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	15.16	0.5725	0.6460	15.17	0.5727	0.6456	15.17	0.5729	0.6455	15.16	0.5724	0.6457
NAFNet	15.24	0.5616	0.6647	15.24	0.5618	0.6640	15.24	0.5619	0.6637	15.24	0.5621	0.6637
Restormer	16.66	0.6045	0.6390	16.72	0.6062	0.6361	16.75	0.6070	0.6329	16.75	0.6066	0.6336
Uformer	15.15	0.5646	0.6747	15.15	0.5650	0.6745	15.15	0.5651	0.6744	15.15	0.5652	0.6748
MIRNetv2	15.21	0.6042	0.6578	15.21	0.6042	0.6575	15.21	0.6042	0.6575	15.21	0.6042	0.6575
AST	15.06	0.5359	0.6924	15.07	0.5363	0.6921	15.07	0.5365	0.6920	15.07	0.5366	0.6924
ConvIR	9.01	0.1575	0.7739	9.01	0.1575	0.7738	9.01	0.1575	0.7736	9.01	0.1574	0.7739
DeepSN-Net	12.84	0.4226	0.7724	12.84	0.4227	0.7720	12.84	0.4228	0.7711	12.84	0.4229	0.7719
SFHformer	12.46	0.5094	0.7405	12.46	0.5094	0.7404	12.46	0.5095	0.7404	12.46	0.5100	0.7400
Models for Related Degradations												
DocRes	16.34	0.6083	0.6546	16.37	0.6097	0.6497	16.36	0.6094	0.6467	16.27	0.6061	0.6502
Retinexformer	14.60	0.5591	0.6512	14.60	0.5589	0.6509	14.60	0.5586	0.6506	14.60	0.5588	0.6509
HVI-CIDNet	15.34	0.5652	0.6857	15.34	0.5658	0.6832	15.61	0.5758	0.6714	15.27	0.5660	0.6863
Ada-Deblur	16.21	0.6061	0.6628	16.25	0.6074	0.6615	16.27	0.6086	0.6608	16.26	0.6078	0.6614
DEA-Net	15.53	0.5866	0.6902	15.81	0.5984	0.6837	15.96	0.6069	0.6741	15.69	0.5947	0.6851
DehazeFormer	14.87	0.5371	0.6685	14.87	0.5373	0.6684	14.87	0.5375	0.6687	14.87	0.5378	0.6689
Pixel2Pixel	12.41	0.5178	0.7304	12.41	0.5178	0.7302	12.41	0.5178	0.7302	12.41	0.5178	0.7305
DarkIR	8.79	0.0950	0.8041	8.79	0.0950	0.8041	8.79	0.0949	0.8040	8.79	0.0949	0.8042
EBLTR	**17.15**	**0.6292**	**0.6307**	**17.16**	**0.6294**	**0.6300**	**17.16**	**0.6295**	**0.6298**	**17.16**	**0.6295**	**0.6298**

Table B22 (continued): 1° to 3°

Method	1°			1.5°			2°			2.5°			3°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4788	0.7230	11.82	0.4788	0.7230	11.82	0.4789	0.7232	11.82	0.4789	0.7238	11.82	0.4789	0.7237
U-Net	12.18	0.4869	0.7631	12.17	0.4868	0.7633	12.17	0.4869	0.7633	12.17	0.4868	0.7633	12.17	0.4868	0.7634
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.16	0.5726	0.6461	15.16	0.5725	0.6462	15.16	0.5724	0.6463	15.16	0.5726	0.6462	15.16	0.5727	0.6464
NAFNet	15.24	0.5622	0.6641	15.24	0.5621	0.6645	15.24	0.5621	0.6647	15.24	0.5623	0.6646	15.24	0.5623	0.6644
Restormer	16.74	0.6065	0.6337	16.73	0.6059	0.6340	16.72	0.6054	0.6344	16.71	0.6052	0.6345	16.70	0.6048	0.6350
Uformer	15.15	0.5652	0.6754	15.15	0.5650	0.6758	15.14	0.5648	0.6762	15.13	0.5644	0.6769	15.11	0.5639	0.6775
MIRNetv2	15.21	0.6042	0.6576	15.21	0.6042	0.6576	15.21	0.6042	0.6578	15.21	0.6042	0.6580	15.21	0.6042	0.6580
AST	15.07	0.5367	0.6925	15.07	0.5367	0.6924	15.06	0.5367	0.6926	15.06	0.5367	0.6930	15.05	0.5366	0.6934
ConvIR	9.01	0.1573	0.7743	9.01	0.1573	0.7743	9.01	0.1572	0.7744	9.01	0.1572	0.7748	9.01	0.1571	0.7747
DeepSN-Net	12.84	0.4230	0.7724	12.83	0.4229	0.7724	12.83	0.4229	0.7728	12.83	0.4229	0.7736	12.83	0.4228	0.7737
SFHformer	12.46	0.5102	0.7402	12.46	0.5102	0.7406	12.46	0.5105	0.7410	12.46	0.5105	0.7413	12.45	0.5104	0.7414
Models for Related Degradations															
DocRes	16.13	0.6028	0.6535	16.06	0.6003	0.6536	15.99	0.5973	0.6556	15.87	0.5939	0.6573	15.88	0.5916	0.6607
Retinexformer	14.60	0.5591	0.6510	14.60	0.5592	0.6511	14.60	0.5594	0.6511	14.60	0.5594	0.6512	14.60	0.5593	0.6518
HVI-CIDNet	15.31	0.5672	0.6841	15.45	0.5683	0.6822	15.47	0.5691	0.6803	15.46	0.5686	0.6814	15.49	0.5710	0.6816
Ada-Deblur	16.23	0.6063	0.6630	16.20	0.6058	0.6632	16.21	0.6056	0.6627	16.23	0.6064	0.6623	16.28	0.6082	0.6608
DEA-Net	15.49	0.5869	0.6881	15.38	0.5828	0.6917	15.31	0.5797	0.6934	15.05	0.5691	0.6971	15.02	0.5676	0.6959
DehazeFormer	14.88	0.5380	0.6693	14.88	0.5378	0.6696	14.88	0.5378	0.6697	14.87	0.5379	0.6701	14.87	0.5378	0.6704
Pixel2Pixel	12.41	0.5178	0.7308	12.41	0.5178	0.7310	12.41	0.5178	0.7312	12.41	0.5178	0.7314	12.41	0.5178	0.7317
DarkIR	8.79	0.0949	0.8046	8.79	0.0949	0.8045	8.79	0.0949	0.8044	8.79	0.0949	0.8049	8.79	0.0948	0.8047
EBLTR	**17.16**	**0.6294**	**0.6301**	**17.15**	**0.6291**	**0.6306**	**17.15**	**0.6291**	**0.6307**	**17.14**	**0.6289**	**0.6309**	**17.13**	**0.6289**	**0.6311**

The Game and Lock tables jointly show that a changing reflected pattern still contains recoverable structure. EBLTR exploits this structure through the same LTOR–EBL mechanism across two distinct content organizations.

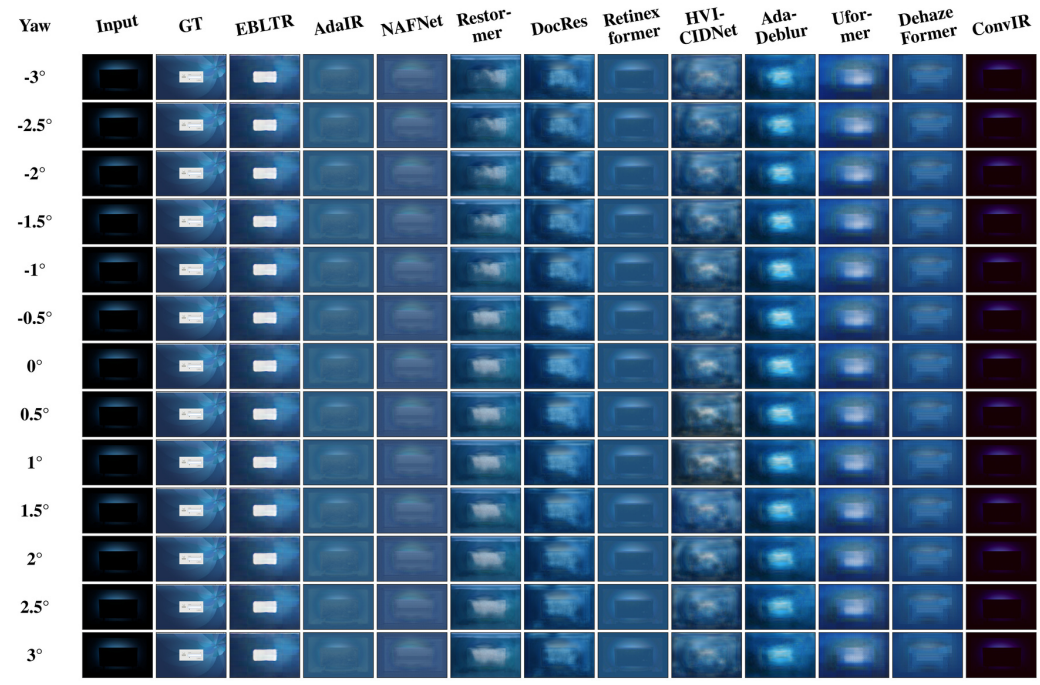

Figure B31: Representative Lock reconstructions under the horizontal-viewing-angle sequence.

Figure B31 supplies the complementary Lock sequence. The colour bands, regional divisions, and main outline remain recognizable across the horizontal viewpoints, making the recovered low-frequency and boundary evidence explicit.

Table B22 carries the horizontal-angle result to Lock. EBLTR preserves the strongest reconstruction comparison across the viewpoints, extending the geometry-aware evidence advantage to broad lock-screen bands and boundaries.

Combined with Table B22, the sequence supports preliminary hidden-screen recovery throughout the tested angular changes and ties the visible result to the complete dataset-level comparison.

B.2.6 VERTICAL VIEWING ANGLE



Figure B32: Acquisition views illustrating the camera placements associated with changes in vertical observation angle.

Vertical viewing-angle changes likewise redistribute the reflected observation. LTOR supplies a fresh ordering of the available evidence at each angle, and EBL translates that ordering into an observability-bounded local correction that keeps the recovered screen structure coherent.

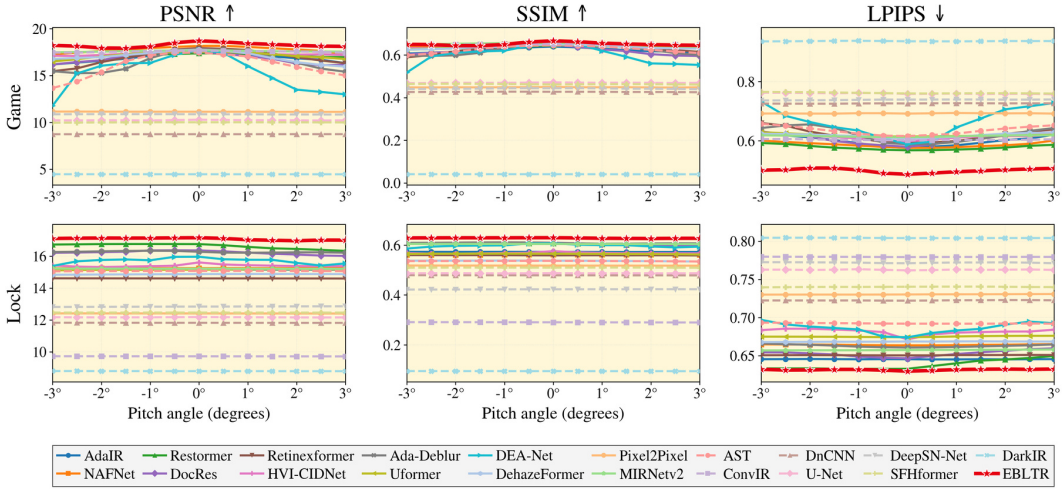

Figure B33: Game and Lock robustness curves under the vertical-viewing-angle sequence.

The vertical-angle study in Figure B33 extends this advantage to a second viewing direction: EBLTR leads PSNR and LPIPS throughout the sequence on both content categories. Together, the angular studies connect changing wall views to recovery of a stable screen target.

Horizontal and vertical angle studies test the same recovery mechanism along two viewing directions. LTOR derives a new observability ordering from each wall image, and EBL converts that ordering into local correction capacity while preserving the hidden screen's principal organization.

Table B23: Complete Game comparison under vertical viewing angle. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	-3.0°			-2.5°			-2.0°			-1.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	8.74	0.4261	0.7255	8.74	0.4268	0.7252	8.73	0.4270	0.7251	8.74	0.4274	0.7250
U-Net	10.22	0.4678	0.7624	10.21	0.4683	0.7620	10.23	0.4696	0.7612	10.22	0.4694	0.7604
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	16.18	0.6068	0.6262	16.38	0.6105	0.6178	16.57	0.6158	0.6112	16.80	0.6218	0.6021
NAFNet	17.27	0.6373	0.6004	17.45	0.6405	0.5957	17.61	0.6439	0.5916	17.78	0.6480	0.5878
Restormer	16.75	0.6251	0.5922	16.86	0.6276	0.5888	17.05	0.6313	0.5822	17.16	0.6346	0.5766
Uformer	16.49	0.6269	0.6283	16.74	0.6302	0.6253	17.03	0.6342	0.6213	17.28	0.6386	0.6174
MIRNetv2	17.48	0.6484	0.6182	17.51	0.6507	0.6172	17.55	0.6525	0.6159	17.61	0.6540	0.6144
AST	13.67	0.5988	0.6571	14.35	0.6062	0.6501	15.36	0.6163	0.6419	16.55	0.6287	0.6322
ConvIR	17.46	0.6489	0.6054	17.44	0.6483	0.6058	17.50	0.6488	0.6048	17.56	0.6496	0.6032
DeepSN-Net	10.86	0.4401	0.7361	10.87	0.4409	0.7363	10.88	0.4414	0.7368	10.88	0.4424	0.7371
SFHformer	9.96	0.4640	0.7650	9.98	0.4647	0.7645	9.99	0.4652	0.7641	10.00	0.4647	0.7630
Models for Related Degradations												
DocRes	16.18	0.6021	0.6311	16.44	0.6100	0.6211	16.63	0.6157	0.6142	17.12	0.6299	0.5987
Retinexformer	15.46	0.5890	0.6597	15.76	0.6000	0.6507	16.43	0.6175	0.6307	16.92	0.6301	0.6167
HVI-CIDNet	17.11	0.6321	0.6205	17.09	0.6327	0.6205	17.19	0.6351	0.6175	17.32	0.6379	0.6122
Ada-Deblur	15.46	0.5976	0.6428	15.25	0.5956	0.6517	15.28	0.5984	0.6553	15.72	0.6102	0.6432
DEA-Net	11.78	0.5180	0.7323	15.22	0.5928	0.6839	16.05	0.6103	0.6631	16.31	0.6163	0.6459
DehazeFormer	16.78	0.6253	0.6234	16.88	0.6275	0.6226	16.96	0.6313	0.6209	17.20	0.6371	0.6179
Pixel2Pixel	11.15	0.4481	0.6919	11.16	0.4478	0.6921	11.15	0.4484	0.6917	11.14	0.4482	0.6916
DarkIR	4.47	0.0410	0.9359	4.47	0.0410	0.9360	4.47	0.0410	0.9367	4.47	0.0410	0.9376
EBLTR	**18.22**	**0.6498**	**0.5003**	**18.14**	**0.6489**	**0.5018**	**17.94**	**0.6437**	**0.5076**	**17.92**	**0.6436**	**0.5076**

Table B23 (continued): -1.0° to 0.5°

Method	-1.0°			-0.5°			0.0°			0.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	8.74	0.4279	0.7259	8.74	0.4278	0.7263	8.74	0.4273	0.7259	8.74	0.4275	0.7265
U-Net	10.21	0.4693	0.7613	10.25	0.4702	0.7607	10.29	0.4715	0.7581	10.26	0.4704	0.7587
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	17.10	0.6291	0.5916	17.36	0.6355	0.5794	17.44	0.6394	0.5754	17.34	0.6355	0.5798
NAFNet	17.93	0.6517	0.5834	18.05	0.6545	0.5791	18.16	0.6587	0.5759	18.16	0.6590	0.5758
Restormer	17.20	0.6372	0.5730	17.31	0.6405	0.5693	17.37	0.6438	0.5676	17.37	0.6431	0.5680
Uformer	17.45	0.6434	0.6146	17.53	0.6461	0.6141	17.55	0.6470	0.6142	17.52	0.6455	0.6147
MIRNetv2	17.65	0.6555	0.6131	17.67	0.6559	0.6125	17.68	0.6559	0.6119	17.67	0.6558	0.6129
AST	17.15	0.6381	0.6240	17.42	0.6437	0.6173	17.46	0.6452	0.6158	17.25	0.6435	0.6195
ConvIR	17.60	0.6508	0.6022	17.64	0.6515	0.6014	17.64	0.6514	0.6013	17.63	0.6511	0.6016
DeepSN-Net	10.88	0.4438	0.7381	10.87	0.4442	0.7390	10.87	0.4440	0.7390	10.86	0.4445	0.7391
SFHformer	10.01	0.4638	0.7619	10.01	0.4628	0.7609	10.01	0.4620	0.7598	10.01	0.4624	0.7597
Models for Related Degradations												
DocRes	17.33	0.6383	0.5903	17.46	0.6435	0.5845	17.53	0.6472	0.5794	17.34	0.6380	0.5874
Retinexformer	17.25	0.6383	0.6080	17.77	0.6508	0.5955	17.92	0.6556	0.5924	17.86	0.6531	0.5943
HVI-CIDNet	17.46	0.6421	0.6075	17.60	0.6466	0.6025	17.77	0.6555	0.5978	17.55	0.6487	0.6043
Ada-Deblur	16.85	0.6354	0.6174	17.59	0.6541	0.5932	17.75	0.6577	0.5878	17.76	0.6580	0.5890
DEA-Net	16.35	0.6200	0.6340	17.20	0.6427	0.6085	17.73	0.6579	0.5879	17.42	0.6502	0.5972
DehazeFormer	17.39	0.6386	0.6162	17.69	0.6467	0.6102	17.77	0.6494	0.6074	17.65	0.6480	0.6093
Pixel2Pixel	11.14	0.4486	0.6930	11.14	0.4492	0.6929	11.14	0.4502	0.6917	11.14	0.4486	0.6927
DarkIR	4.47	0.0410	0.9379	4.46	0.0409	0.9371	4.46	0.0409	0.9361	4.46	0.0409	0.9360
EBLTR	**18.10**	**0.6494**	**0.5009**	**18.50**	**0.6609**	**0.4902**	**18.71**	**0.6665**	**0.4864**	**18.59**	**0.6625**	**0.4905**

Table B23 (continued): 1.0° to 3.0°

Method	1.0°			1.5°			2.0°			2.5°			3.0°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.74	0.4273	0.7269	8.74	0.4265	0.7266	8.74	0.4264	0.7262	8.74	0.4259	0.7264	8.74	0.4253	0.7263
U-Net	10.24	0.4694	0.7593	10.24	0.4701	0.7582	10.22	0.4693	0.7584	10.20	0.4683	0.7585	10.22	0.4692	0.7570
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.22	0.6297	0.5846	17.05	0.6226	0.5928	16.84	0.6175	0.6023	16.69	0.6163	0.6077	16.32	0.6074	0.6204
NAFNet	18.05	0.6537	0.5782	17.91	0.6486	0.5816	17.78	0.6445	0.5846	17.63	0.6400	0.5898	17.28	0.6313	0.6006
Restormer	17.29	0.6414	0.5700	17.21	0.6391	0.5726	17.11	0.6368	0.5768	16.97	0.6324	0.5828	16.96	0.6310	0.5861
Uformer	17.43	0.6434	0.6157	17.28	0.6417	0.6167	17.07	0.6402	0.6188	16.94	0.6389	0.6202	16.77	0.6373	0.6214
MIRNetv2	17.64	0.6547	0.6147	17.59	0.6526	0.6168	17.58	0.6510	0.6177	17.56	0.6493	0.6182	17.56	0.6479	0.6183
AST	17.04	0.6392	0.6241	16.48	0.6323	0.6333	15.91	0.6242	0.6413	15.43	0.6163	0.6473	15.01	0.6089	0.6521
ConvIR	17.60	0.6505	0.6022	17.58	0.6499	0.6026	17.57	0.6494	0.6027	17.52	0.6489	0.6035	17.44	0.6486	0.6048
DeepSN-Net	10.86	0.4436	0.7394	10.85	0.4420	0.7393	10.84	0.4416	0.7388	10.84	0.4408	0.7389	10.84	0.4391	0.7394
SFHformer	10.01	0.4630	0.7602	10.01	0.4632	0.7605	10.01	0.4630	0.7605	10.00	0.4622	0.7601	10.00	0.4613	0.7594
Models for Related Degradations															
DocRes	17.05	0.6265	0.5987	16.80	0.6179	0.6071	16.28	0.6038	0.6222	16.11	0.5973	0.6264	16.13	0.5952	0.6281
Retinexformer	17.73	0.6473	0.5987	17.36	0.6366	0.6100	16.95	0.6288	0.6226	16.60	0.6222	0.6315	16.31	0.6146	0.6382
HVI-CIDNet	17.39	0.6430	0.6105	17.27	0.6392	0.6147	17.23	0.6376	0.6155	17.26	0.6374	0.6156	17.19	0.6339	0.6204
Ada-Deblur	17.56	0.6504	0.5957	16.92	0.6365	0.6060	16.37	0.6237	0.6183	15.74	0.6088	0.6332	15.41	0.5996	0.6427
DEA-Net	16.00	0.6196	0.6447	14.74	0.5912	0.6751	13.49	0.5605	0.7061	13.25	0.5578	0.7164	12.99	0.5532	0.7275
DehazeFormer	17.38	0.6442	0.6133	16.86	0.6379	0.6189	16.42	0.6331	0.6240	16.14	0.6310	0.6262	15.95	0.6299	0.6273
Pixel2Pixel	11.14	0.4485	0.6938	11.13	0.4488	0.6929	11.12	0.4478	0.6924	11.12	0.4472	0.6929	11.12	0.4486	0.6916
DarkIR	4.46	0.0408	0.9359	4.46	0.0408	0.9363	4.46	0.0407	0.9370	4.46	0.0407	0.9369	4.46	0.0406	0.9369
EBLTR	**18.43**	**0.6553**	**0.4950**	**18.33**	**0.6519**	**0.4981**	**18.22**	**0.6487**	**0.5017**	**18.14**	**0.6465**	**0.5045**	**18.11**	**0.6454**	**0.5065**

Table B23 confirms the vertical-angle result for Game across all methods and factors. EBLTR leads PSNR and LPIPS throughout, showing that its playing-area and regional recovery remains closely tied to the changing wall evidence. The vertical-viewpoint comparison shows that coupling wall-evidence representation with bounded refinement preserves a stronger correspondence between the changing observation and the fixed screen target.

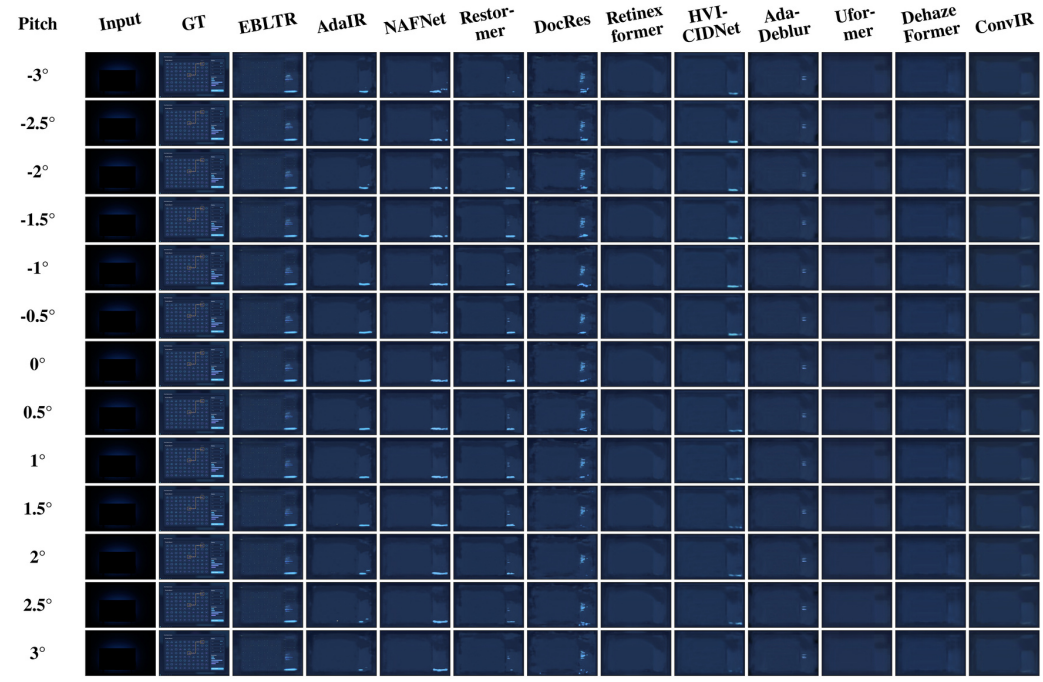

Figure B34: Representative Game reconstructions under the vertical-viewing-angle sequence.

In the dark-screen example in Figure B34, the playing area and interface arrangement provide the visible recovery landmarks. Their continuity in the EBLTR sequence illustrates retention of screen-

level organization as the vertical wall view changes. Broad layout remains the foundation, and principal boundaries receive refinement where the angle-dependent wall response supports them.

Table B24: Complete Lock comparison under vertical viewing angle. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	-3°			-2.5°			-2°			-1.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	11.82	0.4792	0.7224	11.82	0.4791	0.7224	11.82	0.4790	0.7224	11.82	0.4790	0.7224
U-Net	12.17	0.4869	0.7627	12.16	0.4870	0.7627	12.16	0.4871	0.7623	12.16	0.4869	0.7627
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	15.17	0.5730	0.6455	15.17	0.5728	0.6456	15.17	0.5726	0.6458	15.17	0.5727	0.6456
NAFNet	15.24	0.5608	0.6650	15.24	0.5610	0.6649	15.24	0.5611	0.6646	15.24	0.5613	0.6644
Restormer	16.73	0.6070	0.6335	16.74	0.6068	0.6334	16.76	0.6069	0.6335	16.76	0.6073	0.6330
Uformer	15.10	0.5644	0.6750	15.13	0.5651	0.6750	15.14	0.5654	0.6749	15.14	0.5655	0.6749
MIRNetv2	15.22	0.6045	0.6570	15.22	0.6045	0.6571	15.22	0.6044	0.6571	15.22	0.6044	0.6572
AST	15.05	0.5356	0.6933	15.05	0.5358	0.6931	15.06	0.5359	0.6929	15.06	0.5362	0.6927
ConvIR	9.72	0.2912	0.7797	9.72	0.2911	0.7796	9.72	0.2910	0.7795	9.72	0.2909	0.7795
DeepSN-Net	12.81	0.4217	0.7724	12.82	0.4219	0.7724	12.82	0.4220	0.7723	12.83	0.4223	0.7721
SFHformer	12.46	0.5093	0.7396	12.46	0.5095	0.7398	12.46	0.5097	0.7401	12.46	0.5098	0.7401
Models for Related Degradations												
DocRes	16.21	0.6081	0.6545	16.23	0.6084	0.6544	16.25	0.6089	0.6531	16.28	0.6092	0.6517
Retinexformer	14.61	0.5596	0.6507	14.61	0.5596	0.6508	14.61	0.5594	0.6508	14.60	0.5595	0.6507
HVI-CIDNet	15.42	0.5677	0.6836	15.34	0.5653	0.6854	15.31	0.5653	0.6853	15.34	0.5665	0.6844
Ada-Deblur	16.26	0.6089	0.6662	16.27	0.6091	0.6651	16.30	0.6098	0.6644	16.33	0.6106	0.6633
DEA-Net	15.39	0.5860	0.6970	15.67	0.5937	0.6910	15.76	0.5969	0.6882	15.79	0.5979	0.6864
DehazeFormer	14.83	0.5374	0.6679	14.84	0.5375	0.6680	14.85	0.5375	0.6682	14.85	0.5377	0.6683
Pixel2Pixel	12.41	0.5180	0.7300	12.41	0.5180	0.7301	12.41	0.5179	0.7301	12.41	0.5179	0.7301
DarkIR	8.79	0.0951	0.8044	8.79	0.0951	0.8045	8.79	0.0951	0.8045	8.79	0.0950	0.8045
EBLTR	**17.11**	**0.6286**	**0.6323**	**17.13**	**0.6289**	**0.6313**	**17.14**	**0.6291**	**0.6315**	**17.13**	**0.6291**	**0.6321**

Table B24 (continued): -1° to 0.5°

Method	-1°			-0.5°			0°			0.5°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines												
DnCNN	11.82	0.4790	0.7224	11.82	0.4788	0.7222	11.82	0.4787	0.7221	11.82	0.4789	0.7223
U-Net	12.16	0.4864	0.7633	12.17	0.4869	0.7624	12.18	0.4874	0.7618	12.17	0.4869	0.7622
General-Purpose and Multi-Task Image Restoration Frameworks												
AdaIR	15.17	0.5728	0.6454	15.17	0.5726	0.6456	15.17	0.5729	0.6455	15.16	0.5726	0.6452
NAFNet	15.24	0.5616	0.6640	15.24	0.5619	0.6638	15.24	0.5619	0.6637	15.24	0.5622	0.6637
Restormer	16.76	0.6073	0.6329	16.76	0.6069	0.6329	16.75	0.6070	0.6329	16.68	0.6053	0.6365
Uformer	15.15	0.5656	0.6746	15.15	0.5655	0.6744	15.15	0.5651	0.6744	15.16	0.5650	0.6746
MIRNetv2	15.21	0.6044	0.6572	15.21	0.6043	0.6574	15.21	0.6042	0.6575	15.21	0.6042	0.6576
AST	15.07	0.5364	0.6924	15.07	0.5366	0.6922	15.07	0.5365	0.6920	15.07	0.5365	0.6920
ConvIR	9.72	0.2908	0.7794	9.72	0.2907	0.7792	9.72	0.2906	0.7792	9.72	0.2906	0.7792
DeepSN-Net	12.83	0.4226	0.7719	12.83	0.4228	0.7713	12.84	0.4228	0.7711	12.84	0.4229	0.7717
SFHformer	12.46	0.5097	0.7401	12.46	0.5096	0.7402	12.46	0.5095	0.7404	12.46	0.5097	0.7404
Models for Related Degradations												
DocRes	16.36	0.6103	0.6489	16.36	0.6101	0.6475	16.36	0.6094	0.6467	16.30	0.6072	0.6496
Retinexformer	14.60	0.5594	0.6507	14.60	0.5589	0.6506	14.60	0.5586	0.6506	14.60	0.5590	0.6506
HVI-CIDNet	15.37	0.5678	0.6831	15.40	0.5690	0.6810	15.61	0.5758	0.6714	15.48	0.5707	0.6780
Ada-Deblur	16.34	0.6108	0.6619	16.33	0.6103	0.6612	16.27	0.6086	0.6608	16.22	0.6065	0.6610
DEA-Net	15.74	0.5987	0.6845	15.96	0.6063	0.6752	15.96	0.6069	0.6741	15.80	0.6013	0.6799
DehazeFormer	14.86	0.5379	0.6684	14.87	0.5378	0.6685	14.87	0.5375	0.6687	14.88	0.5373	0.6686
Pixel2Pixel	12.41	0.5178	0.7300	12.41	0.5178	0.7301	12.41	0.5178	0.7302	12.40	0.5178	0.7304
DarkIR	8.79	0.0950	0.8043	8.79	0.0950	0.8041	8.79	0.0949	0.8040	8.79	0.0949	0.8040
EBLTR	**17.13**	**0.6291**	**0.6321**	**17.15**	**0.6294**	**0.6311**	**17.16**	**0.6295**	**0.6298**	**17.11**	**0.6285**	**0.6308**

Table B24 (continued): 1° to 3°

Method	1°			1.5°			2°			2.5°			3°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4788	0.7227	11.82	0.4786	0.7227	11.81	0.4787	0.7229	11.81	0.4786	0.7229	11.81	0.4783	0.7226
U-Net	12.17	0.4868	0.7629	12.16	0.4872	0.7625	12.16	0.4868	0.7626	12.15	0.4866	0.7629	12.16	0.4872	0.7624
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.16	0.5723	0.6452	15.16	0.5720	0.6454	15.15	0.5722	0.6453	15.15	0.5722	0.6455	15.15	0.5723	0.6456
NAFNet	15.24	0.5621	0.6639	15.24	0.5617	0.6642	15.24	0.5616	0.6645	15.24	0.5613	0.6648	15.24	0.5608	0.6651
Restormer	16.58	0.6035	0.6401	16.49	0.6012	0.6437	16.45	0.6004	0.6451	16.39	0.5987	0.6471	16.33	0.5971	0.6498
Uformer	15.17	0.5647	0.6753	15.18	0.5642	0.6757	15.18	0.5639	0.6760	15.18	0.5634	0.6761	15.17	0.5622	0.6761
MIRNetv2	15.21	0.6041	0.6577	15.21	0.6040	0.6578	15.21	0.6040	0.6579	15.20	0.6039	0.6581	15.20	0.6038	0.6582
AST	15.07	0.5361	0.6920	15.07	0.5356	0.6920	15.07	0.5354	0.6921	15.06	0.5348	0.6922	15.06	0.5341	0.6920
ConvIR	9.72	0.2904	0.7793	9.71	0.2902	0.7793	9.71	0.2901	0.7794	9.71	0.2900	0.7794	9.71	0.2898	0.7794
DeepSN-Net	12.84	0.4230	0.7721	12.85	0.4230	0.7721	12.85	0.4231	0.7723	12.85	0.4231	0.7719	12.85	0.4230	0.7711
SFHformer	12.46	0.5098	0.7405	12.46	0.5099	0.7405	12.46	0.5098	0.7403	12.46	0.5094	0.7400	12.46	0.5091	0.7399
Models for Related Degradations															
DocRes	16.22	0.6044	0.6522	16.15	0.6020	0.6546	16.10	0.6004	0.6565	16.04	0.5991	0.6586	16.01	0.5985	0.6597
Retinexformer	14.60	0.5591	0.6509	14.60	0.5588	0.6509	14.60	0.5591	0.6510	14.60	0.5589	0.6511	14.60	0.5584	0.6510
HVI-CIDNet	15.43	0.5685	0.6802	15.40	0.5670	0.6810	15.38	0.5671	0.6817	15.37	0.5668	0.6818	15.37	0.5666	0.6841
Ada-Deblur	16.21	0.6054	0.6614	16.22	0.6056	0.6623	16.24	0.6064	0.6630	16.24	0.6071	0.6637	16.22	0.6068	0.6649
DEA-Net	15.77	0.5995	0.6831	15.77	0.5994	0.6853	15.58	0.5944	0.6908	15.41	0.5906	0.6947	15.54	0.5937	0.6926
DehazeFormer	14.88	0.5367	0.6688	14.88	0.5360	0.6689	14.88	0.5356	0.6687	14.88	0.5349	0.6686	14.87	0.5340	0.6685
Pixel2Pixel	12.40	0.5177	0.7304	12.40	0.5176	0.7305	12.40	0.5175	0.7306	12.40	0.5175	0.7307	12.40	0.5174	0.7308
DarkIR	8.79	0.0949	0.8040	8.79	0.0949	0.8040	8.78	0.0948	0.8041	8.78	0.0948	0.8041	8.78	0.0947	0.8041
EBLTR	**17.03**	**0.6272**	**0.6320**	**16.99**	**0.6266**	**0.6324**	**16.97**	**0.6262**	**0.6326**	**17.01**	**0.6269**	**0.6322**	**17.00**	**0.6270**	**0.6327**

Table B24 extends the vertical-angle result to Lock. EBLTR leads the reconstruction comparisons across the sequence, showing that its geometry-aware evidence route also preserves broad screen regions and principal divisions.

The complementary content type is important: the same mechanism serves detailed game-board organization and large lock-screen bands, supporting a common wall-reflected recovery principle across the acquisition scenes.



Figure B35: Representative Lock reconstructions under the vertical-viewing-angle sequence.

Figure B35 supplies the complementary Lock sequence. The recovered colour bands, regional divisions, and main outline remain identifiable as vertical viewpoint changes redistribute the wall response.

Together with Table B24, the figure shows that EBLTR preserves the strongest connection between the observed light spot and the hidden lock screen along the tested vertical-angle range.

B.2.7 Horizontal Camera Position

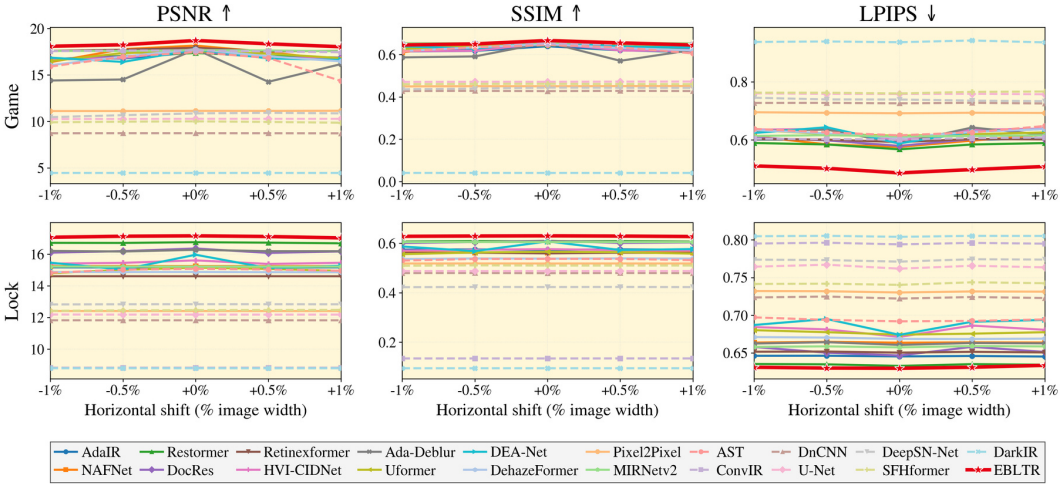

Figure B36: Game and Lock robustness curves under the horizontal-position offsets.

Horizontal camera displacement changes the position and perspective of the wall response in the captured image. The screen content remains fixed, making this a direct test of whether reconstruction follows the screen structure as its observed evidence shifts. EBLTR leads PSNR and LPIPS at every horizontal offset on Game and Lock in Figure B36. The principal screen divisions remain recoverable even as their evidence shifts within the captured wall image. The position sequence directly exercises LTOR's local evidence field. For each shifted input, LTOR recomputes its frequency and multilevel representation; EBL then uses the resulting observability to set local correction limits.

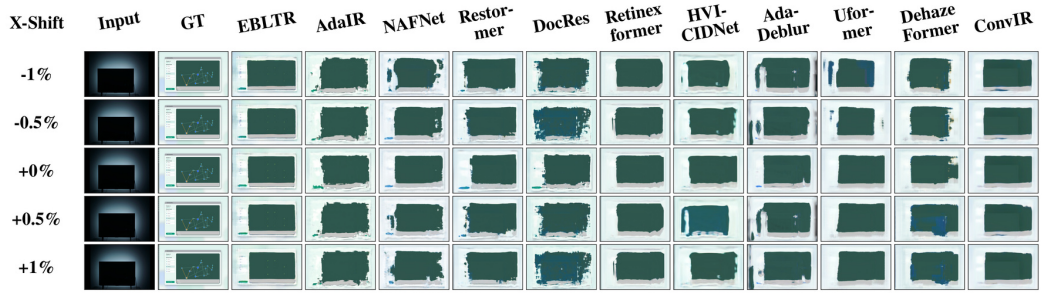

Figure B37: Representative Game reconstructions under the horizontal-position offsets.

The green playing area and surrounding interface in Figure B37 retain their overall arrangement as the observed wall pattern moves horizontally. The result connects spatially shifting evidence to stable screen-domain structure.

Table B25 gives this representative observation its dataset-level foundation. Together they show that ordinary-camera wall capture continues to support preliminary screen recovery throughout the tested horizontal positions. Table B25 anchors Figure B37 in PSNR, SSIM, and LPIPS across the complete Game comparison, linking visible layout retention to the quantitative trajectory.

Table B25 confirms the Game result across all methods and camera offsets. EBLTR leads PSNR and LPIPS throughout, showing that its evidence-aware reconstruction preserves more of the playing

area and regional layout as the observation location changes. The complete comparison links that advantage to method design. The screen target stays organized in its own domain, while LTOR represents the transported evidence in each shifted wall observation and EBL turns that evidence into a matched reconstruction budget.

Table B25: Complete Game comparison under horizontal camera position. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Horizontal shift (% width) -1			Horizontal shift (% width) -0.5			Horizontal shift (% width) 0			Horizontal shift (% width) 0.5			Horizontal shift (% width) 1		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.73	0.4277	0.7272	8.74	0.4286	0.7275	8.74	0.4273	0.7259	8.73	0.4286	0.7277	8.72	0.4277	0.7253
U-Net	10.26	0.4708	0.7602	10.27	0.4717	0.7585	10.29	0.4715	0.7581	10.28	0.4723	0.7587	10.28	0.4726	0.7570
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	16.75	0.6149	0.6015	16.78	0.6171	0.6011	17.44	0.6394	0.5754	16.89	0.6199	0.6000	16.68	0.6119	0.6052
NAFNet	16.37	0.6108	0.6128	17.78	0.6475	0.5846	18.16	0.6587	0.5759	17.47	0.6414	0.5976	16.57	0.6133	0.6181
Restormer	16.70	0.6278	0.5894	17.04	0.6364	0.5844	17.37	0.6438	0.5676	17.02	0.6381	0.5843	16.87	0.6313	0.5891
Uformer	16.49	0.6260	0.6342	17.36	0.6439	0.6205	17.55	0.6470	0.6142	17.30	0.6422	0.6194	16.81	0.6329	0.6258
MIRNetv2	17.58	0.6536	0.6137	17.58	0.6532	0.6162	17.68	0.6559	0.6119	17.64	0.6548	0.6136	17.55	0.6529	0.6161
AST	15.91	0.6152	0.6380	16.95	0.6345	0.6258	17.46	0.6452	0.6158	16.83	0.6345	0.6258	14.36	0.6057	0.6479
ConvIR	17.56	0.6497	0.6042	17.63	0.6516	0.6029	17.64	0.6514	0.6013	17.53	0.6493	0.6063	17.45	0.6471	0.6082
DeepSN-Net	10.46	0.4358	0.7452	10.68	0.4397	0.7393	10.87	0.4440	0.7390	10.93	0.4433	0.7352	10.87	0.4429	0.7332
SFHformer	9.91	0.4600	0.7626	9.98	0.4617	0.7624	10.01	0.4620	0.7598	9.95	0.4608	0.7654	9.86	0.4603	0.7665
Models for Related Degradations															
DocRes	16.70	0.6185	0.6074	16.98	0.6259	0.5985	17.53	0.6472	0.5794	16.86	0.6250	0.6027	16.88	0.6229	0.6027
Retinexformer	17.60	0.6447	0.6027	17.71	0.6481	0.6012	17.92	0.6556	0.5924	17.56	0.6454	0.6033	17.59	0.6462	0.6018
HVI-CIDNet	16.65	0.6156	0.6352	16.82	0.6177	0.6379	17.77	0.6555	0.5978	16.79	0.6192	0.6346	16.79	0.6167	0.6391
Ada-Deblur	14.40	0.5867	0.6316	14.51	0.5912	0.6355	17.75	0.6577	0.5878	14.27	0.5705	0.6428	16.17	0.6221	0.6174
DEA-Net	16.86	0.6394	0.6238	16.40	0.6282	0.6436	17.73	0.6579	0.5879	16.78	0.6381	0.6280	16.62	0.6330	0.6374
DehazeFormer	16.05	0.6164	0.6343	16.98	0.6344	0.6225	17.77	0.6494	0.6074	17.01	0.6307	0.6278	16.46	0.6200	0.6373
Pixel2Pixel	11.12	0.4494	0.6952	11.12	0.4507	0.6932	11.14	0.4502	0.6917	11.14	0.4510	0.6934	11.15	0.4512	0.6928
DarkIR	4.47	0.0410	0.9371	4.46	0.0410	0.9387	4.46	0.0409	0.9361	4.46	0.0409	0.9420	4.46	0.0409	0.9355
EBLTR	**18.09**	**0.6454**	**0.5107**	**18.25**	**0.6503**	**0.5024**	**18.71**	**0.6665**	**0.4864**	**18.35**	**0.6552**	**0.4981**	**18.03**	**0.6452**	**0.5087**

Table B26 carries the camera-position result to Lock. EBLTR leads the reconstruction results throughout the offsets, extending recovery under spatially shifted wall observations from the dense game board to broad lock-screen bands. The agreement between both content categories shows that the method tracks transported structure across changing spatial patterns. LTOR derives observability from the current wall response, and EBL converts it into screen-domain refinement.

Table B26: Complete Lock comparison under horizontal camera position. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Horizontal shift (% width) -1			Horizontal shift (% width) -0.5			Horizontal shift (% width) 0			Horizontal shift (% width) 0.5			Horizontal shift (% width) 1		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4791	0.7237	11.82	0.4794	0.7249	11.82	0.4787	0.7221	11.82	0.4794	0.7242	11.82	0.4790	0.7229
U-Net	12.18	0.4875	0.7644	12.18	0.4873	0.7670	12.18	0.4874	0.7618	12.18	0.4874	0.7657	12.18	0.4875	0.7635
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.17	0.5738	0.6463	15.16	0.5747	0.6464	15.17	0.5729	0.6455	15.16	0.5749	0.6461	15.16	0.5735	0.6453
NAFNet	15.24	0.5620	0.6638	15.24	0.5626	0.6648	15.24	0.5619	0.6637	15.24	0.5629	0.6641	15.24	0.5620	0.6639
Restormer	16.71	0.6064	0.6355	16.71	0.6083	0.6350	16.75	0.6070	0.6329	16.72	0.6082	0.6353	16.68	0.6063	0.6346
Uformer	14.79	0.5555	0.6803	15.01	0.5623	0.6776	15.15	0.5651	0.6744	15.07	0.5634	0.6757	14.94	0.5576	0.6776
MIRNetv2	15.21	0.6043	0.6582	15.21	0.6045	0.6588	15.21	0.6042	0.6575	15.21	0.6045	0.6592	15.21	0.6043	0.6588
AST	14.82	0.5297	0.6973	15.02	0.5359	0.6939	15.07	0.5365	0.6920	15.03	0.5363	0.6926	14.94	0.5312	0.6945
ConvIR	8.82	0.1345	0.7951	8.82	0.1346	0.7962	8.82	0.1345	0.7940	8.82	0.1346	0.7960	8.82	0.1346	0.7949
DeepSN-Net	12.82	0.4221	0.7738	12.83	0.4232	0.7732	12.84	0.4228	0.7711	12.83	0.4232	0.7743	12.83	0.4225	0.7738
SFHformer	12.42	0.5086	0.7416	12.45	0.5104	0.7419	12.46	0.5095	0.7404	12.47	0.5095	0.7437	12.47	0.5096	0.7425
Models for Related Degradations															
DocRes	16.09	0.6004	0.6580	16.19	0.6046	0.6505	16.36	0.6094	0.6467	16.08	0.6007	0.6582	16.17	0.6029	0.6515
Retinexformer	14.60	0.5598	0.6518	14.60	0.5614	0.6520	14.60	0.5586	0.6506	14.60	0.5614	0.6516	14.60	0.5598	0.6508
HVI-CIDNet	15.40	0.5697	0.6844	15.45	0.5732	0.6814	15.61	0.5757	0.6714	15.38	0.5720	0.6864	15.45	0.5702	0.6808
Ada-Deblur	16.20	0.6053	0.6626	16.17	0.6048	0.6644	16.27	0.6086	0.6608	16.19	0.6053	0.6633	16.19	0.6053	0.6628
DEA-Net	15.46	0.5864	0.6872	15.03	0.5677	0.6951	15.96	0.6069	0.6741	15.11	0.5721	0.6915	15.26	0.5760	0.6938
DehazeFormer	14.87	0.5357	0.6712	14.87	0.5381	0.6705	14.87	0.5375	0.6687	14.86	0.5384	0.6686	14.86	0.5375	0.6691
Pixel2Pixel	12.41	0.5179	0.7322	12.41	0.5180	0.7319	12.41	0.5178	0.7302	12.40	0.5180	0.7316	12.40	0.5179	0.7313
DarkIR	8.79	0.0949	0.8047	8.79	0.0949	0.8053	8.79	0.0949	0.8040	8.79	0.0950	0.8052	8.79	0.0950	0.8051
EBLTR	**17.07**	**0.6275**	**0.6312**	**17.13**	**0.6288**	**0.6301**	**17.16**	**0.6295**	**0.6298**	**17.11**	**0.6285**	**0.6312**	**17.03**	**0.6266**	**0.6337**

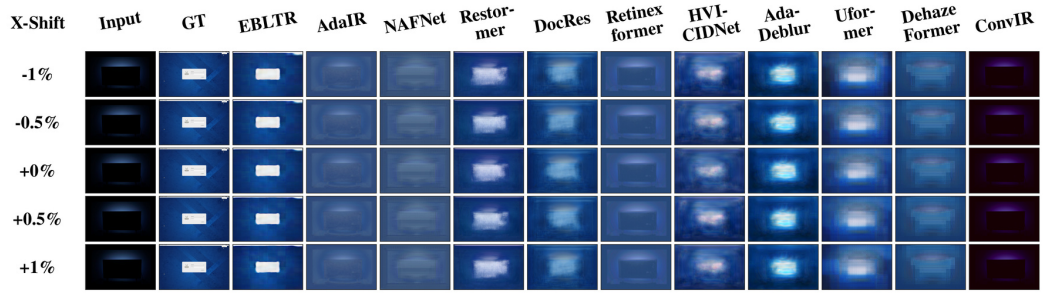

Figure B38: Representative Lock reconstructions under the horizontal-position offsets.

Figure B38 supplies the complementary Lock sequence. The colour bands, regional divisions, and main outline remain recoverable as camera position shifts the observed wall pattern. Paired with the complete Lock table, the sequence demonstrates both visible structure and leading dataset-level agreement across the tested offsets, reinforcing spatially adaptive evidence budgeting.

B.2.8 Wall Texture

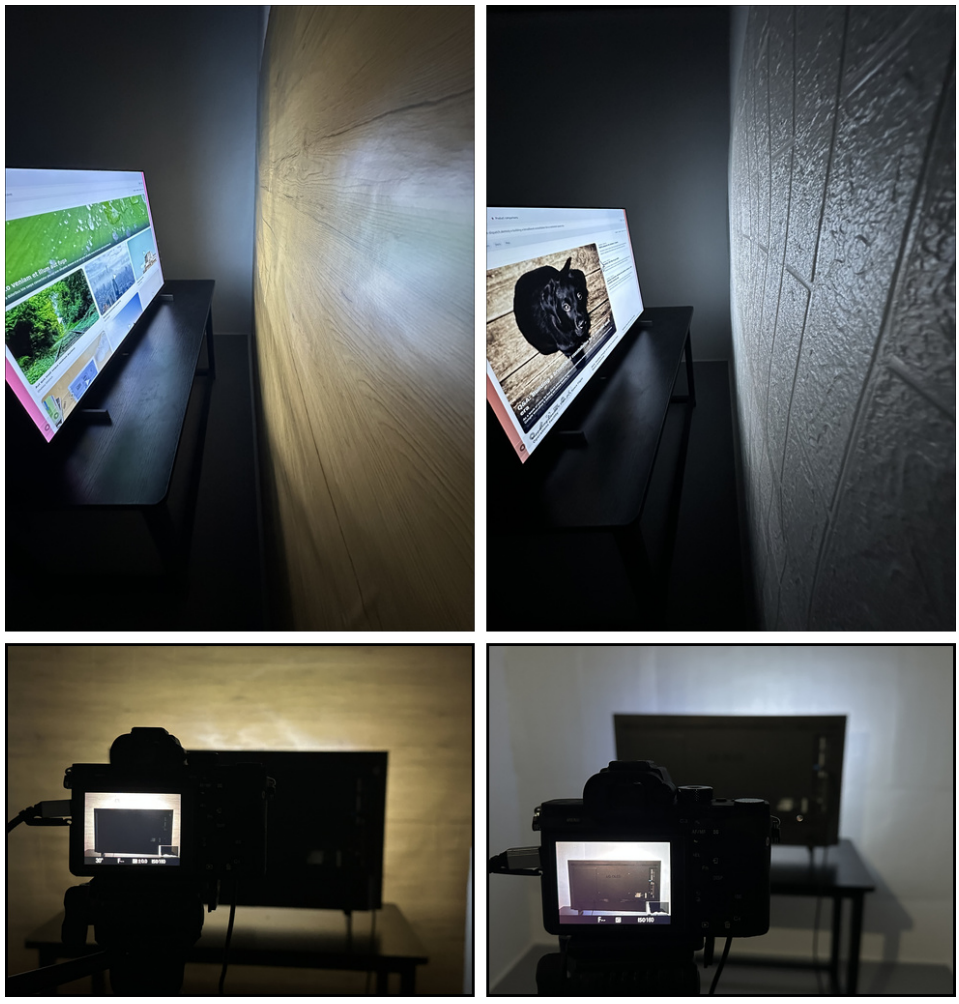

Figure B39: Wall-material photographs illustrating the surface variation motivating the controlled texture study. The left image shows the experimental setup with a wooden wall. The right image shows the experimental setup with a tiled wall.

Wall texture introduces additional surface appearance that competes with the screen information transported through the reflected light, making it more difficult to determine which spatial structures originate from the hidden display and which are caused by the wall itself. This setting therefore

directly tests the value of evidence selection in the proposed framework. LTOR prioritizes screen-relevant structure within mixed wall responses by organizing the observable spatial and frequency cues, while EBL uses this evidence estimate to control the amount of refinement applied to different regions. EBL assigns larger correction limits to higher-observability regions and keeps lower-observability regions closer to the conservative estimate, guiding recovery of the major screen layout and principal contours amid texture-induced wall responses.

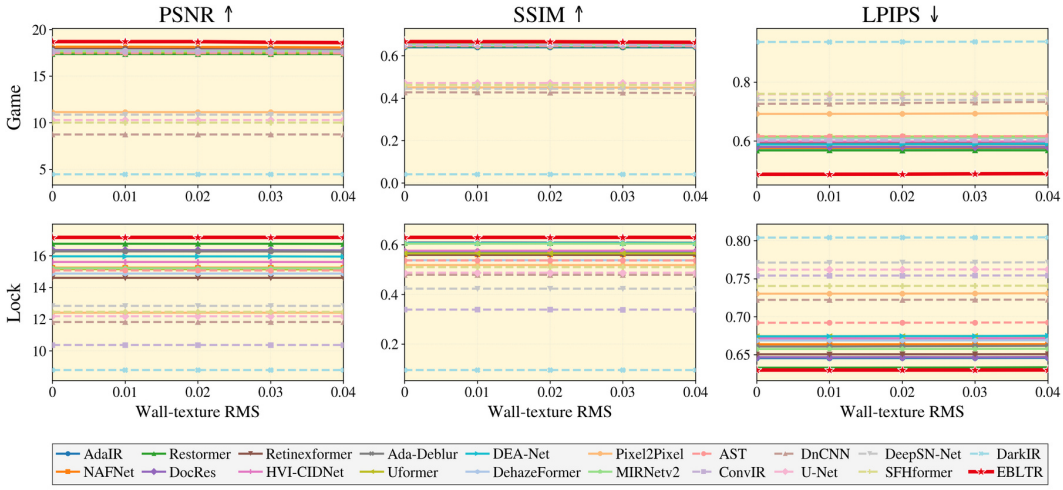

Figure B40: Game and Lock robustness curves under the wall-texture strengths.

Wall texture introduces spatial structure that competes with the screen-derived response. EBLTR leads all three reconstruction metrics at every texture strength on Game and Lock in Figure B40, while the visual comparisons retain the target's main regions and divisions. The result supports the value of organizing wall responses by their relevance to screen recovery. This perturbation targets LTOR's central advantage. Texture energy can be visually strong while carrying limited value for the hidden layout; LTOR evaluates it alongside frequency ordering and multilevel context, and EBL spends detail on the regions that remain connected to screen content.

Table B27: Complete Game comparison under wall texture. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Texture RMS 0.0			Texture RMS 0.01			Texture RMS 0.02			Texture RMS 0.03			Texture RMS 0.04		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.74	0.4273	0.7259	8.74	0.4270	0.7266	8.74	0.4263	0.7284	8.74	0.4252	0.7304	8.74	0.4238	0.7325
U-Net	10.29	0.4715	0.7581	10.29	0.4715	0.7582	10.29	0.4714	0.7584	10.29	0.4714	0.7584	10.28	0.4714	0.7586
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.44	0.6394	0.5754	17.44	0.6394	0.5753	17.43	0.6394	0.5755	17.43	0.6393	0.5754	17.44	0.6393	0.5757
NAFNet	18.16	0.6587	0.5759	18.16	0.6586	0.5761	18.14	0.6582	0.5765	18.12	0.6576	0.5770	18.08	0.6566	0.5776
Restormer	17.37	0.6438	0.5676	17.37	0.6438	0.5676	17.37	0.6438	0.5676	17.36	0.6436	0.5679	17.36	0.6437	0.5678
Uformer	17.55	0.6470	0.6142	17.55	0.6469	0.6141	17.53	0.6466	0.6143	17.53	0.6466	0.6146	17.51	0.6462	0.6148
MIRNetv2	17.68	0.6559	0.6119	17.68	0.6559	0.6120	17.68	0.6558	0.6120	17.68	0.6559	0.6120	17.67	0.6557	0.6122
AST	17.46	0.6452	0.6158	17.45	0.6453	0.6157	17.46	0.6452	0.6158	17.45	0.6451	0.6157	17.43	0.6445	0.6162
ConvIR	17.64	0.6514	0.6013	17.64	0.6514	0.6013	17.64	0.6512	0.6016	17.64	0.6514	0.6015	17.64	0.6513	0.6020
DeepSN-Net	10.87	0.4440	0.7390	10.87	0.4440	0.7390	10.87	0.4439	0.7391	10.87	0.4439	0.7393	10.87	0.4438	0.7393
SFHformer	10.01	0.4620	0.7598	10.01	0.4620	0.7598	10.01	0.4620	0.7599	10.01	0.4619	0.7600	10.01	0.4619	0.7602
Models for Related Degradations															
DocRes	17.53	0.6472	0.5794	17.53	0.6472	0.5794	17.52	0.6471	0.5793	17.52	0.6470	0.5795	17.52	0.6471	0.5794
Retinexformer	17.92	0.6556	0.5924	17.92	0.6555	0.5926	17.90	0.6548	0.5931	17.87	0.6542	0.5939	17.83	0.6529	0.5948
HVI-CIDNet	17.77	0.6555	0.5978	17.77	0.6555	0.5979	17.77	0.6554	0.5979	17.77	0.6555	0.5980	17.78	0.6555	0.5978
Ada-Deblur	17.75	0.6577	0.5878	17.75	0.6576	0.5879	17.75	0.6576	0.5880	17.75	0.6576	0.5881	17.76	0.6579	0.5881
DEA-Net	17.73	0.6579	0.5879	17.72	0.6574	0.5884	17.70	0.6572	0.5889	17.70	0.6567	0.5891	17.66	0.6553	0.5909
DehazeFormer	17.77	0.6494	0.6074	17.77	0.6491	0.6077	17.75	0.6488	0.6082	17.73	0.6483	0.6086	17.69	0.6475	0.6097
Pixel2Pixel	11.14	0.4502	0.6917	11.14	0.4501	0.6918	11.14	0.4498	0.6923	11.13	0.4494	0.6930	11.12	0.4489	0.6938
DarkIR	4.46	0.0409	0.9361	4.46	0.0409	0.9362	4.46	0.0409	0.9365	4.46	0.0409	0.9369	4.46	0.0409	0.9374
EBLTR	**18.71**	**0.6665**	**0.4864**	**18.70**	**0.6662**	**0.4866**	**18.70**	**0.6660**	**0.4868**	**18.63**	**0.6645**	**0.4882**	**18.60**	**0.6632**	**0.4890**

Table B27 confirms the Game result across every model and texture strength. EBLTR leads the reconstruction results throughout the sequence, showing stronger recovery of the game layout in the presence of competing wall appearance.

The full comparison turns texture handling into a relative advantage of evidence organization. Wall structure and screen structure enter the same camera image, while LTOR and EBL preserve the stronger route from transported evidence to the hidden board.

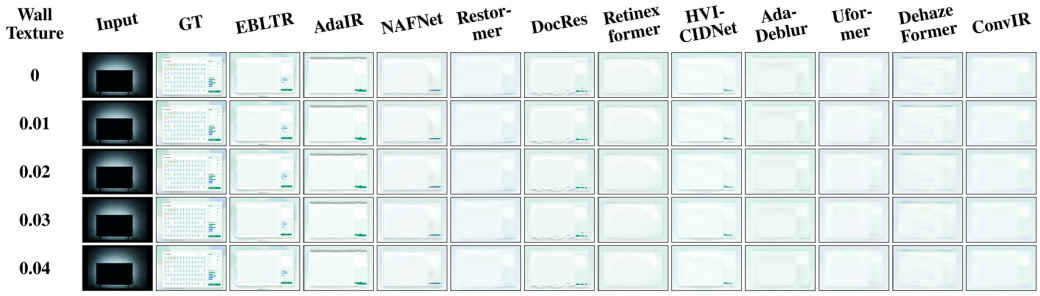

Figure B41: Representative Game reconstructions under the wall-texture sequence.

The light-background game layout remains separated from its side panel as wall texture changes in Figure B41. This example links the reconstruction advantage to preservation of target organization in the presence of competing surface structure.

The representative sequence demonstrates why the task benefits from LTOR. Surface patterns compete with faint screen cues, and the evidence-centred pathway carries the target's broad regions and principal boundaries into the output more consistently.

Table B27 anchors Figure B41 in the complete Game PSNR, SSIM, and LPIPS comparison, tying visual separation of wall and screen structure to the dataset-level trajectory.

Table B28: Complete Lock comparison under wall texture. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Texture RMS 0.0			Texture RMS 0.01			Texture RMS 0.02			Texture RMS 0.03			Texture RMS 0.04		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.82	0.4787	0.7221	11.82	0.4787	0.7221	11.82	0.4787	0.7222	11.82	0.4786	0.7223	11.82	0.4785	0.7224
U-Net	12.18	0.4874	0.7618	12.18	0.4874	0.7618	12.18	0.4874	0.7619	12.18	0.4873	0.7621	12.18	0.4873	0.7622
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.17	0.5729	0.6455	15.17	0.5729	0.6455	15.17	0.5729	0.6456	15.17	0.5728	0.6457	15.17	0.5729	0.6458
NAFNet	15.24	0.5619	0.6637	15.24	0.5619	0.6637	15.24	0.5619	0.6637	15.24	0.5619	0.6638	15.24	0.5618	0.6638
Restormer	16.75	0.6070	0.6329	16.75	0.6070	0.6330	16.75	0.6071	0.6331	16.75	0.6067	0.6332	16.75	0.6068	0.6335
Uformer	15.15	0.5651	0.6744	15.15	0.5652	0.6743	15.15	0.5651	0.6745	15.14	0.5650	0.6746	15.13	0.5650	0.6744
MIRNetv2	15.21	0.6042	0.6575	15.21	0.6042	0.6575	15.21	0.6042	0.6576	15.21	0.6042	0.6577	15.21	0.6042	0.6578
AST	15.07	0.5365	0.6920	15.07	0.5365	0.6920	15.07	0.5365	0.6921	15.07	0.5364	0.6922	15.07	0.5364	0.6925
ConvIR	10.37	0.3385	0.7540	10.37	0.3385	0.7540	10.37	0.3385	0.7541	10.37	0.3384	0.7541	10.37	0.3384	0.7542
DeepSN-Net	12.84	0.4228	0.7711	12.84	0.4228	0.7711	12.84	0.4228	0.7712	12.84	0.4228	0.7713	12.84	0.4228	0.7715
SFHformer	12.46	0.5095	0.7404	12.46	0.5094	0.7404	12.46	0.5094	0.7405	12.46	0.5094	0.7406	12.46	0.5093	0.7408
Models for Related Degradations															
DocRes	16.36	0.6094	0.6467	16.36	0.6094	0.6467	16.35	0.6091	0.6468	16.35	0.6091	0.6467	16.33	0.6080	0.6469
Retinexformer	14.60	0.5586	0.6506	14.60	0.5586	0.6506	14.60	0.5586	0.6506	14.60	0.5586	0.6506	14.60	0.5586	0.6507
HVI-CIDNet	15.61	0.5758	0.6714	15.60	0.5758	0.6715	15.60	0.5757	0.6715	15.60	0.5758	0.6716	15.60	0.5758	0.6717
Ada-Deblur	16.27	0.6086	0.6608	16.27	0.6085	0.6608	16.27	0.6084	0.6611	16.27	0.6082	0.6614	16.26	0.6084	0.6618
DEA-Net	15.96	0.6069	0.6741	15.96	0.6068	0.6742	15.96	0.6068	0.6744	15.96	0.6068	0.6745	15.95	0.6066	0.6749
DehazeFormer	14.87	0.5375	0.6687	14.87	0.5375	0.6687	14.87	0.5375	0.6687	14.87	0.5374	0.6688	14.87	0.5373	0.6690
Pixel2Pixel	12.41	0.5178	0.7302	12.41	0.5178	0.7303	12.41	0.5178	0.7304	12.40	0.5177	0.7305	12.40	0.5177	0.7308
DarkIR	8.79	0.0949	0.8040	8.79	0.0949	0.8040	8.79	0.0949	0.8041	8.79	0.0949	0.8042	8.79	0.0949	0.8043
EBLTR	**17.16**	**0.6295**	**0.6298**	**17.16**	**0.6295**	**0.6298**	**17.16**	**0.6295**	**0.6298**	**17.16**	**0.6296**	**0.6299**	**17.16**	**0.6294**	**0.6299**

Table B28 provides the numerical companion for Lock and records EBLTR's leading reconstruction performance across the texture strengths. The result extends evidence selection to a screen dominated by broad colour bands and layout divisions.

The Lock comparison complements Game by exposing a different structural scale. In both cases, LTOR prioritizes transported screen organization over competing wall appearance, and EBL keeps refinement matched to that evidence.

Together with Table B28, the figure shows that the evidence-budgeted route preserves both a leading metric direction and recognizable hidden content throughout the tested wall-texture sequence.

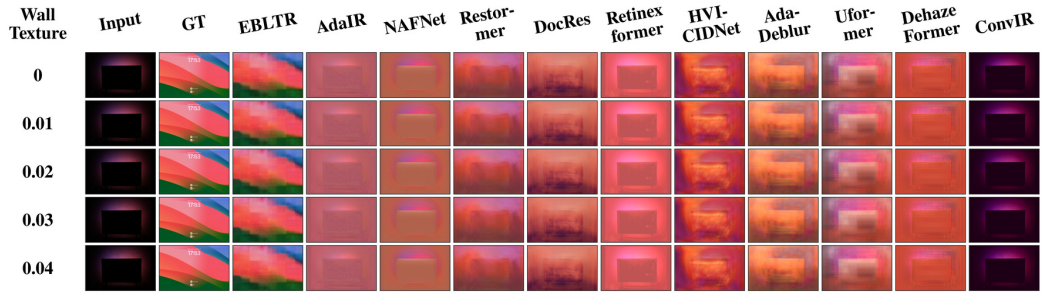

Figure B42: Representative Lock reconstruction comparison under five wall-texture strengths.

Supplementary Figure B42 provides a focused Lock comparison under smooth wall texture. The EBLTR column retains the broad three-band colour layout across the displayed strengths more consistently, while several baseline outputs carry wall-boundary structure into the result or merge the target into broad colour regions.

The example makes LTOR's relative advantage concrete: observable screen bands are selected from a structured relay surface, then EBL allocates refinement around the recovered regional divisions. The resulting output remains closer to the hidden lock-screen organization throughout the displayed sequence.

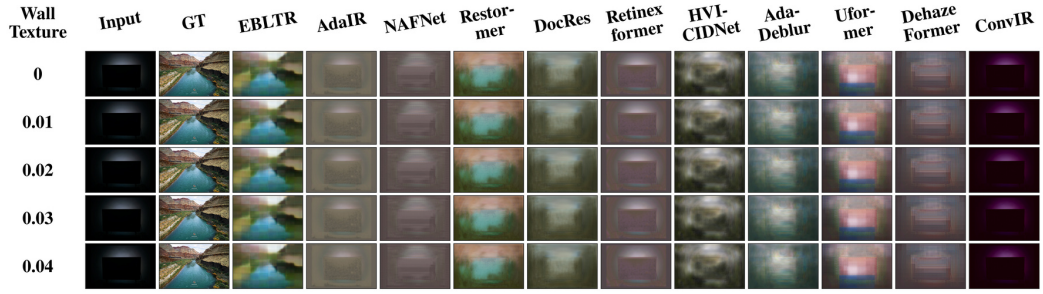

Figure B43: Representative Lock reconstructions under the wall-texture sequence.

Figure B43 supplies the complementary Lock sequence. EBLTR retains the lock-screen colour bands, regional divisions, and main outline across the texture strengths, making the recovered screen evidence visible against the changing wall surface. As additional texture is introduced, the observation contains more competing spatial structure, yet the dominant screen organization remains recognizable.

The factor matrix links wall texture directly to recoverable content. Across the settings, EBLTR continues to preserve the broad lock-screen layout and principal boundaries, indicating that LTOR can organize transport-related evidence while EBL limits refinement to regions supported by the observation.

B.2.9 WALL REFLECTANCE

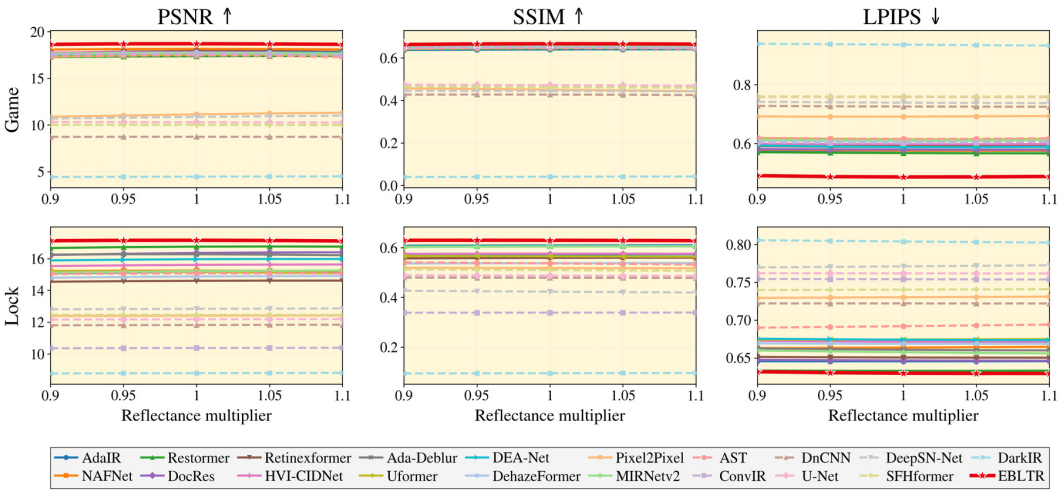

Figure B44: Game and Lock robustness curves under the wall-reflectance factors.

Changing reflectance varies the strength of the transported screen signal. EBLTR leads all three metrics at every relative reflectance factor on Game and Lock in Figure B44, and the reconstructions preserve the dominant screen layout. This complements the texture study by testing signal strength alongside competing surface structure.

Reflectance changes evidence magnitude throughout the image, making adaptive budgeting especially relevant. LTOR evaluates the revised band and local responses, and EBL scales refinement through the resulting observability field, preserving the screen's broad organization across the tested sequence.

Table B29: Complete Game comparison under wall reflectance. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Relative reflectance gain 0.9			Relative reflectance gain 0.95			Relative reflectance gain 1.0			Relative reflectance gain 1.05			Relative reflectance gain 1.1		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	8.73	0.4272	0.7278	8.74	0.4275	0.7266	8.74	0.4273	0.7259	8.73	0.4265	0.7254	8.73	0.4251	0.7250
U-Net	10.31	0.4740	0.7583	10.30	0.4727	0.7583	10.29	0.4715	0.7581	10.27	0.4703	0.7577	10.26	0.4692	0.7574
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	17.40	0.6374	0.5770	17.42	0.6386	0.5760	17.44	0.6394	0.5754	17.44	0.6400	0.5750	17.44	0.6402	0.5749
NAFNet	18.08	0.6552	0.5795	18.14	0.6576	0.5775	18.16	0.6587	0.5759	18.14	0.6587	0.5755	18.10	0.6579	0.5760
Restormer	17.29	0.6413	0.5704	17.34	0.6428	0.5687	17.37	0.6438	0.5676	17.39	0.6446	0.5668	17.39	0.6451	0.5666
Uformer	17.32	0.6430	0.6186	17.47	0.6454	0.6157	17.55	0.6470	0.6142	17.53	0.6469	0.6140	17.46	0.6452	0.6155
MIRNetv2	17.67	0.6547	0.6135	17.68	0.6555	0.6126	17.68	0.6559	0.6119	17.67	0.6558	0.6117	17.64	0.6554	0.6121
AST	17.33	0.6425	0.6194	17.42	0.6444	0.6168	17.46	0.6452	0.6158	17.41	0.6447	0.6161	17.27	0.6429	0.6180
ConvIR	17.58	0.6509	0.6031	17.63	0.6517	0.6016	17.64	0.6514	0.6013	17.59	0.6502	0.6026	17.48	0.6480	0.6054
DeepSN-Net	10.69	0.4447	0.7419	10.79	0.4445	0.7402	10.87	0.4440	0.7390	10.94	0.4433	0.7381	11.00	0.4425	0.7374
SFHformer	10.02	0.4646	0.7604	10.02	0.4632	0.7600	10.01	0.4620	0.7598	10.00	0.4607	0.7598	9.99	0.4596	0.7599
Models for Related Degradations															
DocRes	17.48	0.6453	0.5814	17.51	0.6466	0.5800	17.53	0.6472	0.5794	17.53	0.6473	0.5791	17.53	0.6471	0.5793
Retinexformer	17.72	0.6455	0.5987	17.86	0.6521	0.5946	17.92	0.6556	0.5924	17.92	0.6563	0.5923	17.84	0.6554	0.5938
HVI-CIDNet	17.77	0.6547	0.5990	17.78	0.6553	0.5982	17.77	0.6555	0.5978	17.76	0.6555	0.5978	17.75	0.6552	0.5979
Ada-Deblur	17.64	0.6532	0.5909	17.72	0.6562	0.5888	17.75	0.6577	0.5878	17.76	0.6583	0.5877	17.73	0.6580	0.5883
DEA-Net	17.55	0.6523	0.5936	17.64	0.6552	0.5902	17.73	0.6579	0.5879	17.76	0.6593	0.5874	17.72	0.6590	0.5889
DehazeFormer	17.62	0.6444	0.6099	17.73	0.6479	0.6080	17.77	0.6494	0.6074	17.73	0.6489	0.6083	17.64	0.6466	0.6111
Pixel2Pixel	10.88	0.4561	0.6922	11.02	0.4534	0.6914	11.14	0.4502	0.6917	11.24	0.4468	0.6925	11.31	0.4434	0.6934
DarkIR	4.43	0.0398	0.9394	4.45	0.0404	0.9377	4.46	0.0409	0.9361	4.48	0.0414	0.9346	4.50	0.0419	0.9332
EBLTR	**18.64**	**0.6629**	**0.4909**	**18.70**	**0.6657**	**0.4875**	**18.71**	**0.6665**	**0.4864**	**18.68**	**0.6660**	**0.4867**	**18.64**	**0.6646**	**0.4881**

Table B29 confirms the Game result across all baselines and reflectance factors. EBLTR leads the reconstruction comparisons throughout, showing that its evidence-aware pathway makes the strongest use of the available reflected light at each level.

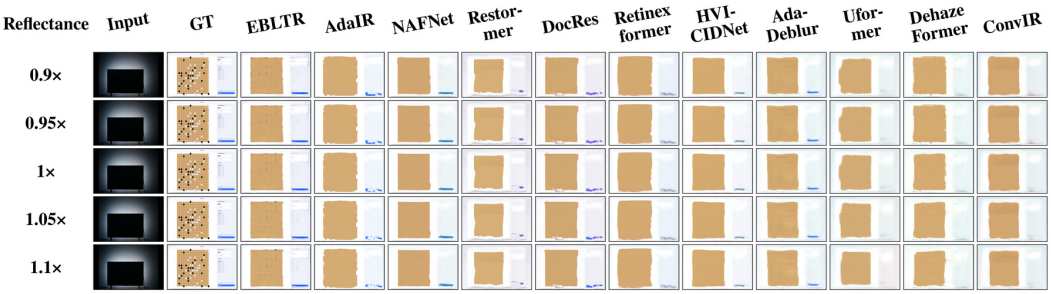

Figure B45: Representative Game reconstructions under the wall-reflectance sequence.

The board region and adjacent interface panel in Figure B45 remain recognizable as reflected-light strength changes. Their continuity shows that the wall response still carries usable evidence for the target's principal divisions.

The figure illustrates the method's full progression: reflected intensity supplies the available evidence, LTOR organizes its spatial and frequency structure, and EBL sets the local correction limits from that evidence. This allows the dominant game layout and regional boundaries to remain stable as the reflected signal changes.

Table B29 anchors Figure B45 in the complete Game comparison, connecting visible recovery of board structure to PSNR, SSIM, and LPIPS across all reflectance factors.

Taken together, Figures B43 and B45 show robustness to two different wall variations: texture changes the surface structure, while reflectance changes the signal strength. In both cases, EBLTR continues to recover the dominant hidden-screen layout and principal contours by organizing observable evidence and allocating refinement accordingly.

Table B30: Complete Lock comparison under wall reflectance. PSNR↑/SSIM↑/LPIPS↓. Best in red, second best in blue.

Method	Relative reflectance gain 0.9			Relative reflectance gain 0.95			Relative reflectance gain 1			Relative reflectance gain 1.05			Relative reflectance gain 1.1		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Classical CNN Baselines															
DnCNN	11.80	0.4794	0.7221	11.81	0.4791	0.7221	11.82	0.4787	0.7221	11.83	0.4783	0.7220	11.84	0.4780	0.7220
U-Net	12.16	0.4888	0.7620	12.17	0.4881	0.7619	12.18	0.4874	0.7618	12.19	0.4867	0.7617	12.19	0.4859	0.7616
General-Purpose and Multi-Task Image Restoration Frameworks															
AdaIR	15.08	0.5714	0.6457	15.13	0.5723	0.6455	15.17	0.5729	0.6455	15.19	0.5733	0.6456	15.21	0.5735	0.6457
NAFNet	15.21	0.5642	0.6630	15.24	0.5630	0.6633	15.24	0.5619	0.6637	15.24	0.5608	0.6641	15.22	0.5597	0.6647
Restormer	16.67	0.6066	0.6333	16.72	0.6069	0.6331	16.75	0.6070	0.6329	16.76	0.6068	0.6330	16.75	0.6065	0.6332
Uformer	15.19	0.5662	0.6743	15.18	0.5657	0.6743	15.15	0.5651	0.6744	15.11	0.5644	0.6746	15.07	0.5635	0.6750
MIRNetv2	15.11	0.6022	0.6598	15.17	0.6033	0.6586	15.21	0.6042	0.6575	15.24	0.6048	0.6567	15.25	0.6051	0.6561
AST	15.03	0.5408	0.6901	15.06	0.5387	0.6910	15.07	0.5365	0.6920	15.07	0.5343	0.6931	15.07	0.5320	0.6941
ConvIR	10.35	0.3382	0.7544	10.36	0.3383	0.7542	10.37	0.3385	0.7540	10.38	0.3386	0.7539	10.39	0.3388	0.7537
DeepSN-Net	12.81	0.4264	0.7697	12.82	0.4246	0.7704	12.84	0.4228	0.7711	12.85	0.4210	0.7718	12.86	0.4193	0.7725
SFHformer	12.47	0.5123	0.7399	12.46	0.5109	0.7401	12.46	0.5095	0.7404	12.46	0.5080	0.7406	12.45	0.5067	0.7409
Models for Related Degradations															
DocRes	16.23	0.6075	0.6475	16.30	0.6086	0.6470	16.36	0.6094	0.6467	16.39	0.6099	0.6465	16.40	0.6101	0.6464
Retinexformer	14.54	0.5564	0.6514	14.58	0.5577	0.6509	14.60	0.5586	0.6506	14.61	0.5593	0.6504	14.61	0.5597	0.6504
HVI-CIDNet	15.53	0.5752	0.6727	15.58	0.5756	0.6720	15.61	0.5758	0.6714	15.62	0.5757	0.6711	15.63	0.5754	0.6710
Ada-Deblur	16.25	0.6084	0.6626	16.27	0.6086	0.6616	16.27	0.6086	0.6608	16.25	0.6082	0.6603	16.21	0.6075	0.6600
DEA-Net	15.89	0.6050	0.6755	15.93	0.6061	0.6747	15.96	0.6069	0.6741	15.97	0.6073	0.6737	15.97	0.6073	0.6735
DehazeFormer	14.80	0.5363	0.6693	14.84	0.5370	0.6690	14.87	0.5375	0.6687	14.89	0.5377	0.6686	14.90	0.5378	0.6687
Pixel2Pixel	12.39	0.5187	0.7294	12.40	0.5182	0.7298	12.41	0.5178	0.7302	12.41	0.5174	0.7307	12.42	0.5169	0.7311
DarkIR	8.77	0.0939	0.8056	8.78	0.0944	0.8048	8.79	0.0949	0.8040	8.80	0.0954	0.8032	8.81	0.0959	0.8026
EBLTR	**17.13**	**0.6295**	**0.6318**	**17.15**	**0.6297**	**0.6306**	**17.16**	**0.6295**	**0.6298**	**17.14**	**0.6289**	**0.6295**	**17.11**	**0.6281**	**0.6295**

Table B30 extends the reflectance result to Lock and retains every model at each factor. EBLTR leads the reconstruction comparison, showing that the same evidence-aware mechanism preserves broad lock-screen structure as reflected-light strength changes.

The paired Game and Lock results strengthen the task-level conclusion. Across detailed and coarse content organizations, a wall reflection continues to carry useful brightness and layout evidence, and EBLTR turns that evidence into the strongest reconstruction.

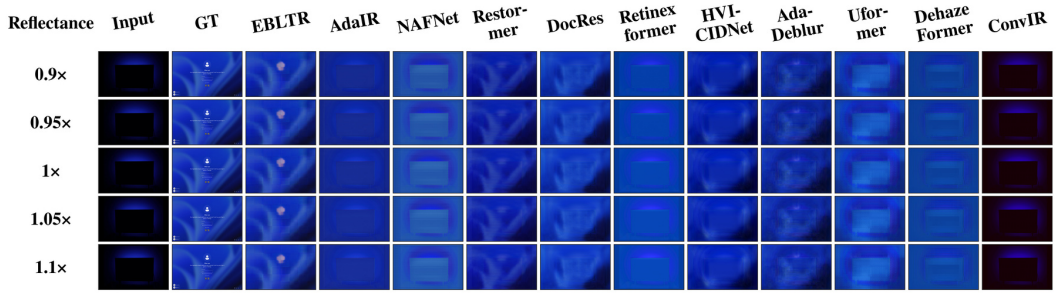

Figure B46: Representative Lock reconstructions under the wall-reflectance sequence.

Figure B46 supplies the complementary Lock sequence. The colour bands, regional divisions, and main outline remain identifiable across the reflectance factors, visualizing the coarse screen structure retained in the wall response.

Together with Table B30, the figure connects that recognizable structure to a leading dataset-level trajectory and reinforces ordinary-camera wall observation as a viable route to preliminary hidden-screen recovery under the tested reflectance changes.

The allocation studies in Tables A5–A6 and the robustness results in Figures B14–B46 connect local refinement decisions to visible screen recovery. Across changes in illumination, sampling, noise, viewing geometry, and wall appearance, EBLTR preserves screen organization by linking added detail to wall-derived evidence.

B.3 Evidence Chain Across the Complete Experimental Supplement

The diagnostic studies connect the recovery of hidden screen content to two linked decisions: which wall evidence to use and how much refinement it should permit. Equal-area occlusions reveal spatially uneven reconstruction influence (Figure B8). The representation comparisons connect LTOR to fewer erroneous boundaries, including in low-observability regions (Table B7), and this advantage persists under local brightness, texture, and high-frequency perturbations (Table B9). With LTOR fixed, the learning and allocation studies show how evidence-aligned correction improves the use of the same representation (Tables A6 and B10). Together, these experiments support a common principle: the amount and placement of refinement should follow the evidence preserved in the wall response.

The robustness studies put this principle into changing observations. Illumination and reflectance vary signal strength, geometry changes the observed arrangement, and noise and wall texture introduce competing responses. Across these changes, EBLTR continues to recover the screen's dominant colour organization and regional layout, with local structure following the available evidence. Together, these results support the preliminary feasibility of recovering hidden screen content from ordinary-camera wall observations, with surviving wall evidence guiding refinement locations and bounding local correction magnitudes.