

Wavefront-Constrained Passive Obscured Object Detection

Zhiwen Zheng^{1,2*}, Yiwei Ouyang^{1,2*}, Zhao Huang³, Tao Zhang¹, Xiaoshuai Zhang^{4†}, Huiyu Zhou⁵,
Wenwen Tang^{6†}, Shaowei Jiang^{1†}, Jin Liu^{1,2†}, Xingru Huang^{1,2†}

¹School of Communication Engineering, Hangzhou Dianzi University, Hangzhou 310018, China

²Weidian Corporation, Beijing, 100015, China

³Department of Computer and Information Science, Northumbria University, NE1 8ST Newcastle upon Tyne, UK

⁴Faculty of Information Science and Engineering, Ocean University of China, 266404, Qingdao, China

⁵School of Informatics, University of Leicester, Leicester LE1 7RH, UK

⁶Johns Hopkins University, Baltimore, MD, 21218, USA

zhiwen.zheng@hdu.edu.cn, yiwei.ouyang@hdu.edu.cn, zhao.huang@northumbria.ac.uk, tzhang@hdu.edu.cn,
x.zhang@ouc.edu.cn, hz143@leicester.ac.uk, wtang18@jhu.edu, jiangsw@hdu.edu.cn, jinliu@hdu.edu.cn,
xingru.huang@hdu.edu.cn

Abstract

Accurately localizing and segmenting obscured objects from faint light patterns beyond the field of view is highly challenging due to multiple scattering and medium-induced perturbations. Most existing methods, based on real-valued modeling or local convolutional operations, are inadequate for capturing the underlying physics of coherent light propagation. Moreover, under low signal-to-noise conditions, these methods often converge to non-physical solutions, severely compromising the stability and reliability of the observation. To address these challenges, we propose a novel physics-driven Wavefront Propagating Compensation Network (WavePCNet) to simulate wavefront propagation and enhance the perception of obscured objects. This WavePCNet integrates the Tri-Phase Wavefront Complex-Propagation Reprojection (TriWCP) to incorporate complex amplitude transfer operators to precisely constrain coherent propagation behavior, along with a momentum memory mechanism to effectively suppress the accumulation of perturbations. Additionally, a High-frequency Cross-layer Compensation Enhancement is introduced to construct frequency-selective pathways with multi-scale receptive fields and dynamically models structural consistency across layers, further boosting the model’s robustness and interpretability under complex environmental conditions. Extensive experiments conducted on four physically collected datasets demonstrate that WavePCNet consistently outperforms state-of-the-art methods across both accuracy and robustness.

Code —

<https://github.com/IMOP-lab/WavePCNet-Pytorch>

Introduction

The Non-line-of-sight (NLOS) targets detection has broad real-world applications, including blind-spot collision avoidance in autonomous vehicles, security surveillance

*Equal Contribution.

†Corresponding Authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

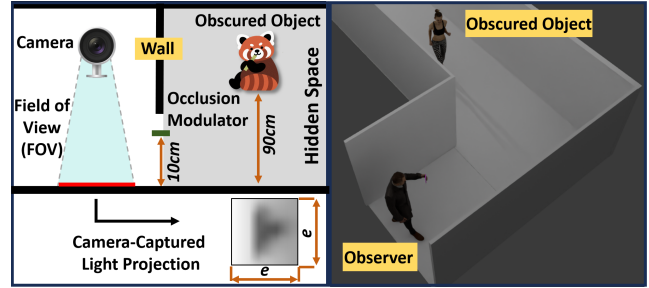


Figure 1: Typical scenario underlying the passive OOD regime, in which the target, entirely occluded by an opaque intervening barrier that prevents any direct radiative conveyance to the sensor, is inferred solely through the diffuse, low-intensity speckle distribution scattered onto the visible relay surface.

through occluding structures, and post-disaster search and rescue scenarios, where these targets are often hidden around corners or fully obscured by opaque barriers (Zhao et al. 2023). These applications pose a challenging question “Is it possible to localize a target hidden behind a wall solely by relying on the faint speckle field generated by ambient light scattering, without any active illumination?”

Accordingly, the work formulates the Obscured Object Detection (OOD) task: from a measured surface-intensity distribution $I_w(\mathbf{x})$ of exceptionally low signal-to-noise ratio, determine both a semantic mask $M(\mathbf{x})$ and a geometric estimate \mathbf{p} of the out of view object. This passive regimen circumvents spectral-allocation and regulatory burdens imposed by emitters, permits direct adoption of commodity imaging assemblies, and thus promises economical deployment (Beckus, Tamasan, and Atia 2019) (Boger-Lombard and Katz 2019). Figure 1 depicts the canonical arrangement: an object entirely eclipsed by an opaque screen, no direct photons admitted to the sensor; yet, from the residual speckle upon the observable wall, the observer infers its

position and extent.

Although OOD affords emitter free operation, cost efficiency and deployment simplicity, it exacts inherently feeble observations and entails convoluted physics: multiple reflections and medium roughness rapidly decorrelate phase along propagation paths, impose spectral bandlimitation and deplete high-frequency content, such that only low-SNR, phase mismatched speckle appears on accessible surfaces (Khawaja et al. 2019). Models built on real-valued convolution kernels cannot represent the phase perturbations and path-integral effects intrinsic to coherent scattering (Rotter and Gigan 2017). Hence two principal obstacles arise: first, in occluded environments, convoluted propagation trajectories and severe wavefront distortion may drive the inversion to settle in unphysical or multi-valued solution spaces (May et al. 2021); second, indirect reflection induces energy decay and high frequency detail loss, manifesting as blurred edges and fractured structural features (Cao, Mosk, and Rotter 2022).

To overcome these challenges, we propose WavePCNet, a novel framework tailored for inferring obscured objects from a single speckle field, which consists of Tri-Phase Wavefront Complex-Propagation Reprojection (TriWCP) and High frequency Cross-layer Compensation. The TriWCP, grounded in Fresnel diffraction, forges a three-stage complex-domain optimisation pathway; a momentum-retentive update confines the trajectory to a physically admissible low-rank sub-space, curbing multiplicity driven by wavefront distortion. The Cross-layer Compensation, via multi-scale frequency-selective filtering, amplifies residual high-frequency cues, while a semantics-guided attention stream aligns structural detail across layers, thus conjointly repairing blurred edges and fragmented semantics.

The main contributions are summarized as follow:

- Introduce WavePCNet, a novel physics-driven framework for passive detection of obscured objects beyond the line of sight. By unifying complex-amplitude propagation with frequency–structure compensation, WavePCNet achieves state-of-the-art accuracy and robustness on our physically collected datasets, even under extreme occlusion and low-SNR conditions.
- Propose the TriWCP module, which pioneers a three-phase complex-domain optimization strategy grounded in Fresnel diffraction. It incorporates an OTF-consistent complex propagation operator and a momentum memory mechanism to guide the solution within a physically valid, low-rank subspace, effectively mitigating ambiguity from wavefront distortion.
- Design the Cross-Layer Compensation module, which integrates multi-scale frequency-selective filtering with semantics-guided attention to enhance both high-frequency detail restoration and cross-layer semantic coherence, enabling more precise reconstruction of fragmented object structures.

Related Work

Visible-domain object detection comprises multiple sub-tasks, notably Salient Object Detection (SOD) (Wei et al.

2024) and Camouflaged Object Detection (COD) (Ruan et al. 2025). Models such as U2-Net (Qin et al. 2020), PFPN (Wang et al. 2020), F3Net (Wei, Wang, and Huang 2020) represent typical SOD pipelines that adopt multi-scale aggregation and edge refinement, raising localisation fidelity within salient extents; COD, confronting background homogeneity and boundary ambiguity, invokes frequency domain representation, contextual fusion, prompt guidance, as implemented in FAPNet (Zhou et al. 2022), FNet (Zhong et al. 2022), VSCODE (Luo et al. 2024), enlarges representational capacity. Generative counterparts, CamoDiffusion (Chen, Sun, and Lin 2024), Camoformer (Yin et al. 2024), together with branched backbones DC-Net (Zhu, Qin, and Elsaddik 2025), ZoomNeXt (Pang et al. 2024), concentrate on camouflage edge modelling and detail amplification; RefCOD (Zhang et al. 2025) appends linguistic cues, improving multi-instance semantic disentanglement under concealment. Progress centres on fusion, edge adjustment, spectrum modelling; unaddressed remains coherent, phase-coupled indirect transport, hence full occlusion outside the field of view escapes current visible-domain formulations.

Non-line-of-sight (NLOS) imaging supplies an alternative trajectory: indirect radiative transport characterised, hidden geometry or radiance inferred. Transient schemes, e.g., NeTF (Shen et al. 2021), utilise time-of-flight sampling; coherent formulations, Liu et al. (Liu et al. 2019), inject phase-propagation analysis for strong scattering; millimetre-wave radar localises obscured objects through penetration (Scheiner et al. 2020), infrared thermography retrieves outlines via thermal radiative fields (Yuan et al. 2024). Structural accuracy emerges, yet specialised hardware, environmental stability, cost burdens persist; focus settles on physical-field regression, not semantic discrimination. Hybrid networks embedding physical priors mitigate expense, still confine modelling to signal synthesis, low SNR structural recognition and boundary restoration remaining unresolved.

Method

The study proposed WavePCNet (Figure 2), a novel physics-driven framework designed to infer obscured objects from a single speckle field in passive, non-line-of-sight scenarios. Specifically, it integrates two key components: (1) Tri-Phase Wavefront Complex-Propagation Reprojection (TriWCP), grounded in Fresnel diffraction, which constructs a three-stage optimization pathway in the complex domain. A momentum-retentive update constrains the solution trajectory within a physically valid low-rank subspace, effectively mitigating ambiguities introduced by wavefront distortion. (2) High-Frequency Cross-Layer Compensation, which enhances residual high-frequency cues via multi-scale frequency-selective filtering, while a semantics-guided attention stream aligns structural details across layers. Together, these components enable robust recovery of blurred edges and fragmented semantics, significantly improving detection stability and accuracy under complex occlusions.

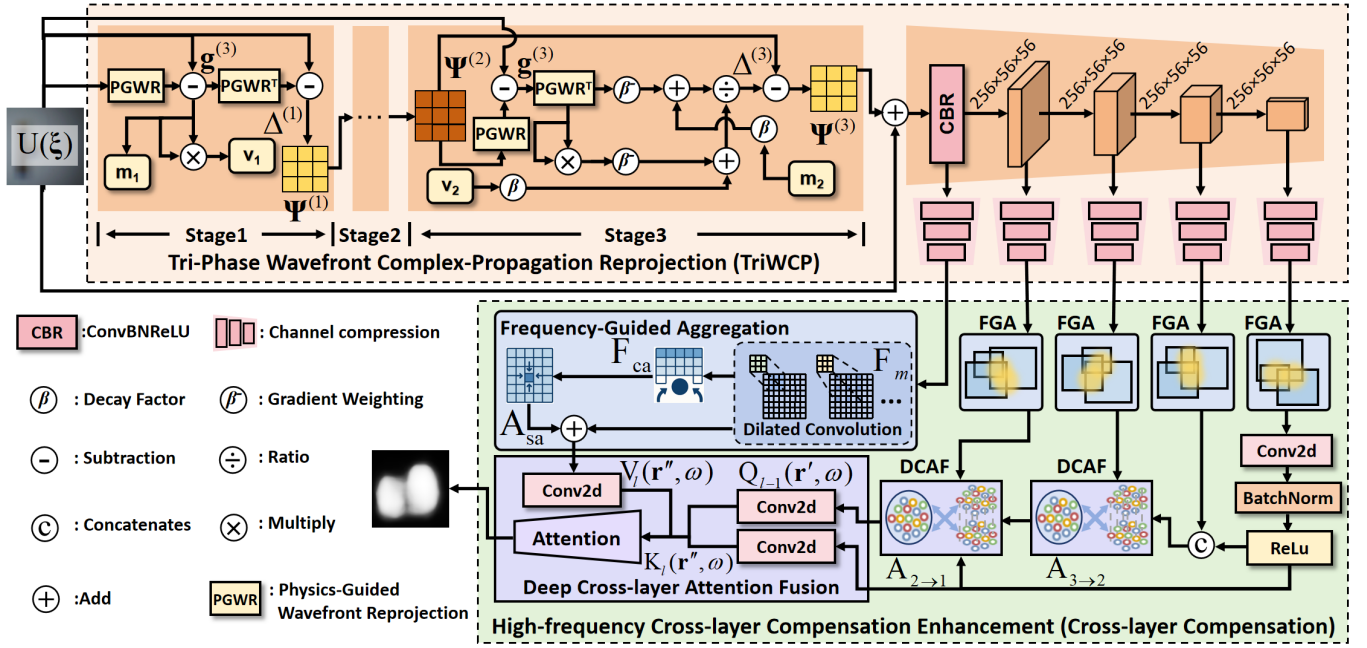


Figure 2: Schematic representation of WavePCNet architecture, wherein a bifurcated topology integrates (i) a physically constrained optimisation pipeline grounded in complex-amplitude propagation (TriWCP), enforcing low-rank spectral admissibility via manifold-constrained trajectory reprojection; and (ii) a cross-layer compensatory stream that augments attenuated high-frequency constituents through semantics-guided, frequency-selective enhancement.

Tri-Phase Wavefront Complex-Propagation Reprojection

Propagation of light manifests intrinsic wave behavior and coherence; its advancement adheres to Fresnel diffraction as encapsulated by the complex-amplitude transfer function, thereby confining field evolution to a physically admissible, low-dimensional subspace. To enforce this constraint, we introduce a Tri-Phase Wavefront Complex-Propagation Reprojection module, which comprises a three-stage reconstruction pipeline that explicitly embeds the complex-valued propagation operator, restricts the optimisation trajectory to the coherent low-rank subspace, and reprojects obscured object image structure from the complex wavefront. Such reprojection mitigates unphysical solutions and solution multiplicity. A momentum memory mechanism accumulates integrated responses of historical propagation paths, attenuating high frequency noise and stochastic perturbations. The pipeline is illustrated in Figure 3.

Under the high frequency scalar approximation, free-space complex amplitude propagation conforms to Fresnel diffraction theory, such that the free propagation process remains strictly governed by the physically permissible propagation operator. Denoting the source plane complex amplitude by $\mathcal{U}(\xi)$, the propagation distance by z and the observation plane field by $\mathcal{E}(\rho)$, the Hertz-vector wave equation under scalar approximation yields the complex-propagation integral:

$$\mathcal{E}(\rho) = \frac{e^{ikz}}{i\lambda z} \iint_{\mathbb{R}^2} \mathcal{U}(\xi) \exp\left(\frac{ik}{2z}\|\rho - \xi\|^2\right) d\xi. \quad (1)$$

Let $k = 2\pi/\lambda$ denote the wavenumber, i.e., the propagation constant; λ the wavelength; $\xi, \rho \in \mathbb{R}^2$ the transverse spatial coordinates; the quadratic term $\|\rho - \xi\|^2$ the squared transverse separation, representing the projected path-length difference. The integral thus manifests propagation as a space invariant, convolution type superposition that aggregates global contributions of the source field. Passing to the Fourier domain yields:

$$\begin{aligned} \mathcal{E}(\rho) &= \mathcal{F}^{-1}\{H_F(\mathbf{f}) \cdot \mathcal{F}\{\mathcal{U}(\xi)\}\}, \\ H_F(\mathbf{f}) &= \exp(-i\pi\lambda z\|\mathbf{f}\|^2), \end{aligned} \quad (2)$$

where \mathcal{F} and \mathcal{F}^{-1} denote the two-dimensional Fourier transform and its inverse, \mathbf{f} the spatial-frequency variable with unit m^{-1} . The multiplicative factor $H_F(\mathbf{f})$, known as the complex amplitude transfer function or optical transfer function (OTF), is a pure phase term that effects the Fresnel quadratic-phase modulation of the spectrum.

Under the Fresnel transfer formulation, free space advancement acts as a complex linear operator that couples a passband frequency filter with quadratic phase modulation. Define, for any $\mathcal{U} \in L^2(\mathbb{R}^2, \mathbb{C})$, the propagation operator \mathcal{T}_z by:

$$\mathcal{T}_z[\mathcal{U}] = \mathcal{F}^{-1}[\chi_{\Omega_z}(\mathbf{f}) \cdot \exp(-i\pi\lambda z\|\mathbf{f}\|^2) \cdot \mathcal{F}\{\mathcal{U}\}], \quad (3)$$

where χ_{Ω_z} denotes the indicator of the passband $\Omega_z := \mathbf{f} \in \mathbb{R}^2 \mid \|\mathbf{f}\|^2 \leq 1/(\lambda z)$; frequencies outside Ω_z are suppressed and, within this model, treated as non-propagating.

The physically admissible fields consequently reside in a low-rank complex subspace:

$$\mathbb{S}_z := \left\{ \psi \in L^2(\mathbb{R}^2, \mathbb{C}) \mid \psi = \mathcal{F}^{-1}[\chi_{\Omega_z} \cdot \hat{\psi}] \right\}. \quad (4)$$

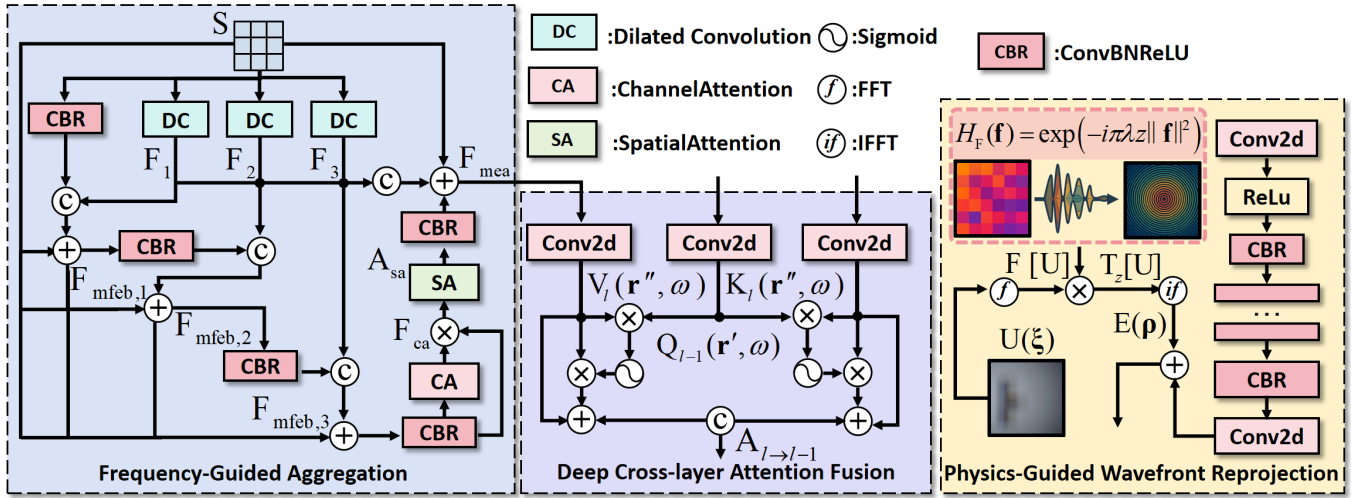


Figure 3: Architectural composition of the proposed modules. Left: Cross-layer Compensation integrates multi-receptive field frequency-selective filtering and semantics-guided structural alignment to recover attenuated high-frequency cues. Center: spatial-frequency attention projection across layers enhances structural consistency. Right: Tri-Phase Complex-Propagation Reprojection constrains optimisation within a physically admissible low-rank subspace via embedded Fresnel transfer operator and memory-augmented updates.

In effect, its effective dimensionality scales with the area of Ω_z and is strictly lower-rank within $L^2(\mathbb{R}^2)$. Any perturbation $\delta\mathcal{U} \notin \mathbb{S}_z$ is removed by propagation. Hence we project search directions onto \mathbb{S}_z via the orthogonal projector:

$$\delta\mathcal{U}_{\text{phys}} = \mathcal{P}_{\mathbb{S}_z}[\delta\mathcal{U}] = \mathcal{T}_z^\dagger \mathcal{T}_z[\delta\mathcal{U}], \quad (5)$$

with $\mathcal{P}_{\mathbb{S}_z}$ the orthogonal projection onto \mathbb{S}_z and \mathcal{T}_z^\dagger the adjoint of \mathcal{T}_z . The discrete Fourier domain derivation appears in Appendix. Constraining updates by \mathcal{T}_z thus restricts the reachable set to \mathbb{S}_z , automatically rejecting directions inconsistent with the propagation physics.

In realistic optical systems, the complex field $\Psi^{(z)}(\mathbf{r})$ at any position \mathbf{r} arises from the aggregate superposition of multiple diffraction, interference and scattering trajectories, which collectively form an integral propagation–response operator spanning the full history of propagation.

By introducing a generalized propagator \mathcal{A}_z^\dagger , which encapsulates all orders of diffraction, interference and scattering, and by projecting gradients onto the physically admissible subspace \mathbb{S}_z , do we prevent non-physical perturbations from accumulating during optimization. To that end, momentum and energy-spectrum terms are formed via exponential weighting:

$$\mathbf{m}^{(t)}(\mathbf{r}) := \sum_{\tau=1}^t \beta_1^{t-\tau} \cdot \mathcal{A}_z^\dagger \left[\mathbf{g}^{(\tau)}(\mathbf{r}) - \Pi_{\mathbb{S}_z}[\mathbf{g}^{(\tau)}](\mathbf{r}) \right], \quad (6)$$

where the decay factor $\beta_1^{t-\tau}$ attenuates early perturbations, imparting an “inertial memory” of the field’s interference structure and thus enhancing convergence stability and directional robustness in complex media.

$$\mathbf{v}^{(t)}(\mathbf{r}) := \sum_{\tau=1}^t \beta_2^{t-\tau} \cdot \left| \mathcal{A}_z^\dagger \left[\mathbf{g}^{(\tau)}(\mathbf{r}) - \Pi_{\mathbb{S}_z}[\mathbf{g}^{(\tau)}](\mathbf{r}) \right] \right|^2, \quad (7)$$

this energy spectrum term capturing the spatial accumulation of perturbations along propagation paths, modulates momentum amplitude and suppresses unstable propagation modes via square-root scaling. Complete derivation in Appendix.

To attenuate residual high spatial frequency perturbations while maintaining the spatial continuity mandated by physical propagation, a convolutional smoothing corrector is inserted into the momentum update: the previous-step residual $\delta^{(t-1)}(\mathbf{r}) = \mathbf{g}^{(t-1)}(\mathbf{r}) - \Pi_{\mathbb{S}_z}[\mathbf{g}^{(t-1)}](\mathbf{r})$ passes through a low-pass kernel $\mathcal{K}(\mathbf{r} - \mathbf{r}')$, yielding a locally averaged disturbance which biases the momentum direction away from non-physical oscillatory components and, hence, prevents their admission into the update trajectory. The update at iteration t is defined as:

$$\Delta^{(t)}(\mathbf{r}) = \frac{\mathbf{m}^{(t)}(\mathbf{r}) - \iint_{\mathbb{R}^2} \mathcal{K}(\mathbf{r} - \mathbf{r}') \cdot \delta^{(t-1)}(\mathbf{r}') \, d\mathbf{r}'}{\sqrt{\mathbf{v}^{(t)}(\mathbf{r})} + \varepsilon}. \quad (8)$$

Starting from the initial complex field $\Psi^{(0)}(\mathbf{r})$, three such iterations accumulate to $\Psi^{(3)}(\mathbf{r})$; enforcing the optical constraint then projects the result onto the feasible set \mathbb{S}_z , giving the final field:

$$\Psi_{\text{final}}(\mathbf{r}) = \Pi_{\mathbb{S}_z} \left\{ \Psi^{(0)}(\mathbf{r}) - \sum_{t=1}^3 \gamma_t \cdot \Delta^{(t)}(\mathbf{r}) \right\}, \quad (9)$$

where γ_t denotes the step size at iteration t , and $\Delta^{(t)}(\mathbf{r})$ is the update increment defined above.

High-frequency Cross-layer Compensation Enhancement

Indirect observation entails multibounce transport through intermediary surfaces; along such energy paths the measured field undergoes severe attenuation, spectral smoothing, and detail depletion. In the hidden-target region, multipath loss, non-ideal reflection, and ambient noise jointly degrade edge sharpness; high-frequency geometric components, already faint, become difficult to resolve, which constrains structural-fidelity recovery. To mitigate this bottleneck, Cross-layer Compensation forms two complementary streams with multi-receptive-field design and cross-level structural consistency: a frequency-selective branch that extracts residual high-frequency evidence via variably dilated band-pass kernels, avoiding over-smoothing inherent to purely spatial convolution; a semantic-anchoring branch that uses high-level global semantics as anchors to align low-level spatial detail, thereby suppressing noise-induced spurious boundaries. The pipeline is illustrated in Figure 3.

During propagation the optical signal experiences repeated reflections and frequency-dependent material attenuation; cumulative loss grows with the reflection count K . Yet high-frequency components, less susceptible to absorption, may retain discernible traces in the final observation despite low energy. Mathematical modeling and frequency domain analysis appear in Appendix.

To explicitly target these components, we define a set of frequency-selective filters with spatial locality and adaptive receptive fields. Let $\mathcal{S}(\mathbf{r}, t)$ denote the input spatiotemporal signal, where $\mathbf{r} \in \mathbb{R}^2$ is spatial position and $t \in \mathbb{R}$ is time. Each branch m uses a learnable band-pass kernel \mathcal{G}_m with spatial dilation δ_m and scale σ_m , modulated by frequency response \mathcal{W}_m (see Appendix for details).

For each filter branch with dilation $m \in \{1, 3, 5\}$, we define a Gaussian convolution operator that provides localized smoothing:

$$\mathcal{K}_\sigma[f](\mathbf{r}, t) := \int_{\mathbb{R}^2} f(\mathbf{r}', t) \cdot \exp\left(-\frac{\|\mathbf{r} - \mathbf{r}'\|^2}{2\sigma^2}\right) d\mathbf{r}', \quad (10)$$

where $\mathbf{r}, \mathbf{r}' \in \mathbb{R}^2$ are spatial coordinates, $\sigma \in \mathbb{R}^+$ controls the degree of spatial spread, regulating sensitivity to local variation.

Filtered features are:

$$\mathcal{F}_m(\mathbf{r}, t) = \iint_{\mathbb{R}^2 \times \mathbb{R}} \mathcal{G}_m(\mathbf{r} - \mathbf{r}', \omega'; \delta_m) \cdot \mathcal{Q}_m(\mathbf{r}', t, \omega') d\mathbf{r}' d\omega', \quad (11)$$

where $\mathcal{Q}_m(\mathbf{r}', t, \omega') := \mathcal{S}(\mathbf{r}', t) \cdot \mathcal{A}_m(\mathbf{r}', \omega')$, \mathcal{S} is the input signal, $\mathcal{A}_m(\mathbf{r}', \omega') = \sigma(\Theta_m \cdot \mathcal{F}_{\omega'}[\mathcal{S}(\mathbf{r}', t)])$, $\Theta_m \in \mathbb{R}^{C \times C}$ are learnable parameters, and $\sigma(\cdot)$ is ReLU. Fused multi-branch features:

$$\mathcal{F}_{\text{fused}}(\mathbf{r}, t) = \sigma\left(\mathcal{K}_{\sigma_{\text{fuse}}}\left[\sum_{m=1}^3 \mathcal{T}_m \cdot \mathcal{F}_m\right]\right), \quad (12)$$

where $\mathcal{T}_m \in \mathbb{R}^{C \times C}$ are fusion weights, $\sigma_{\text{fuse}} \in \mathbb{R}^+$.

The outputs of all branches are fused to form a unified frequency-enhanced representation:

$$\mathcal{F}_{\text{mfeb},k}(\mathbf{r}, t) = \mathcal{K}_{\sigma_k}\left[\sum_{m=1}^k \mathcal{F}_m + \mathcal{H}_k + \mathcal{S}\right], \quad (13)$$

where $\mathcal{H}_k(\mathbf{r}, t)$ is the output of the k -th auxiliary convolution and $\sigma_k \in \mathbb{R}^+$.

We then apply channel-wise attention to reweight semantic channels by their spatial importance:

$$\begin{aligned} \mathcal{C}(\mathbf{r}, c) &= \sigma(\mathcal{K}_{\sigma_c}[\mathcal{F}_{\text{mfeb}} \cdot \Phi_c]), \\ \mathcal{F}_{\text{ca}}(\mathbf{r}, t) &= \sum_{c=1}^C \mathcal{C}(\mathbf{r}, c) \cdot \mathcal{F}_{\text{mfeb}}(\mathbf{r}, t, c), \end{aligned} \quad (14)$$

where $\Phi_c \in \mathbb{R}^{C \times C}$, $\sigma_c \in \mathbb{R}^+$. Next, edge-aware filtering is applied by combining low-level and high-level paths:

$$\begin{aligned} \mathcal{F}_{\text{edge},\text{in}} &= \mathcal{K}_{\sigma_{\text{edge}}}[\mathcal{F}_{\text{fused}} + \mathcal{F}_{\text{ca}}], \\ \mathcal{E}(\mathbf{r}, t) &= \mathcal{K}_{\sigma_{\text{ep}}}[\Psi_{\text{edge}} \cdot \mathcal{F}_{\text{edge},\text{in}}], \end{aligned} \quad (15)$$

where $\Psi_{\text{edge}} \in \mathbb{R}^{C \times 1}$, $\sigma_{\text{edge}}, \sigma_{\text{ep}} \in \mathbb{R}^+$.

We further include a spatial attention map to emphasize salient regions:

$$\mathcal{A}_{\text{sa}}(\mathbf{r}, t) = \sigma(\mathcal{K}_{\sigma_{\text{sa}}}[\mathcal{F}_{\text{ca}} \cdot \Phi_{\text{sa}}]), \quad (16)$$

where $\Phi_{\text{sa}} \in \mathbb{R}^{C \times 1}$, $\sigma_{\text{sa}} \in \mathbb{R}^+$.

We generate the saliency-enhanced features by combining all enhanced features and saliency cues:

$$\mathcal{F}_{\text{mea}}(\mathbf{r}, t) = \mathcal{K}_{\sigma_{\text{mea}}}\left[\left(\mathcal{A}_{\text{sa}} + \sigma(\mathcal{E})\right) \cdot \mathcal{F}_{\text{ca}} + \sum_{k=1}^3 \mathcal{F}_{\text{mfeb},k} + \mathcal{S}\right], \quad (17)$$

where $\sigma_{\text{mea}} \in \mathbb{R}^+$. Final saliency convolution:

$$\mathcal{F}_{\text{out}}(\mathbf{r}, t) = \mathcal{K}_{\sigma_{\text{sal}}}[\Psi_{\text{sal}} \cdot \mathcal{F}_{\text{mea}}], \quad (18)$$

where $\Psi_{\text{sal}} \in \mathbb{R}^{C \times C}$, $\sigma_{\text{sal}} \in \mathbb{R}^+$.

We introduce cross-layer attention to inject high-level feedback into low-level features via a similarity-weighted aggregation. Cross-layer attention weight and response:

$$\alpha_l(\mathbf{r}, \mathbf{r}', \omega) = \sigma\left(\frac{\int_{\mathbb{R}} \mathcal{Q}_{l-1}(\mathbf{r}, \omega') \cdot \mathcal{K}_l(\mathbf{r}', \omega') d\omega'}{\sqrt{\int_{\mathbb{R}^2} \int_{\mathbb{R}} \|\mathcal{K}_l(\mathbf{r}', \omega')\|^2 d\omega' d\mathbf{r}'}}\right), \quad (19)$$

$$\begin{aligned} \mathcal{A}_{l \rightarrow l-1}(\mathbf{r}, t) &= \int_{\mathbb{R}^2} \int_{\mathbb{R}} \alpha_l(\mathbf{r}, \mathbf{r}', \omega) \cdot \mathcal{V}_l(\mathbf{r}', \omega) \\ &\quad \cdot \exp\left(-\frac{\|\mathbf{r} - \mathbf{r}'\|^2}{2\sigma_{\text{attn}}^2}\right) d\omega d\mathbf{r}', \end{aligned} \quad (20)$$

where $\mathcal{Q}_{l-1}, \mathcal{K}_l, \mathcal{V}_l \in \mathbb{C}$, and $\sigma_{\text{attn}} \in \mathbb{R}^+$.

Finally, the fused output incorporating both attention and saliency enhancement is:

$$\mathcal{F}_{\text{final}}(\mathbf{r}, t) = \mathcal{K}_{\sigma_{\text{final}}}[\mathcal{A}_{l \rightarrow l-1} + \mathcal{F}_{\text{out}}], \quad (21)$$

where $\sigma_{\text{final}} \in \mathbb{R}^+$.

The complete derivation is provided in Appendix. This framework compensates for low-frequency loss, leverages semantic priors, and enhances stability in low signal-to-noise conditions.




Figure 4: Visual comparison of detection results across representative methods on occluded-target datasets.

Method	OOD-DUTS						OOD-HKU					
	Max-F	Mean-F	Fbw	MAE	SM	EM	Max-F	Mean-F	Fbw	MAE	SM	EM
UNet (Ronneberger, Fischer, and Brox 2015)	0.645	0.602	0.446	0.220	0.674	0.785	0.515	0.499	0.370	0.233	0.608	0.748
LDF (Wei et al. 2020)	0.649	0.604	0.455	0.214	0.677	0.784	0.522	0.501	0.388	0.220	0.611	0.747
F3Net (Wei, Wang, and Huang 2020)	0.646	0.603	0.451	0.215	0.676	0.783	0.514	0.498	0.379	0.227	0.608	0.744
GCPANet (Chen et al. 2020)	0.646	0.600	0.456	0.211	0.677	0.781	0.520	0.501	0.383	0.221	0.612	0.751
EDN (Wu et al. 2022)	0.647	0.600	0.456	0.214	0.676	0.780	0.531	0.514	0.390	0.211	0.614	0.760
PFSNet (Ma, Xia, and Li 2021)	0.649	0.602	0.446	0.220	0.673	0.782	0.510	0.492	0.374	0.219	0.604	0.748
MENet (Wang et al. 2023)	0.638	0.602	0.433	0.225	0.672	0.786	0.512	0.494	0.355	0.231	0.603	0.750
DC-Net (Zhu, Qin, and Elsaddik 2025)	0.484	0.434	0.302	0.369	0.472	0.658	0.413	0.369	0.265	0.290	0.515	0.683
RefCOD (Zhang et al. 2025)	0.670	0.626	0.505	0.179	0.697	0.802	0.506	0.489	0.373	0.226	0.602	0.742
ZoomNeXt (Pang et al. 2024)	0.672	0.629	0.495	0.186	0.700	0.802	0.532	0.510	0.389	0.224	0.615	0.746
FAPNet (Zhou et al. 2022)	0.648	0.604	0.459	0.207	0.680	0.785	0.520	0.504	0.387	0.213	0.610	0.756
FDNet (Zhong et al. 2022)	0.652	0.604	0.446	0.221	0.674	0.783	0.518	0.499	0.379	0.233	0.610	0.745
Our	0.692	0.654	0.550	0.154	0.708	0.823	0.564	0.545	0.439	0.170	0.610	0.785

Table 1: Performance comparison on OOD-DUTS and OOD-HKU datasets under six evaluation metrics.

Experiments

We evaluate WavePCNet on four physically constructed datasets: OOD-DUTS, OOD-HKU, OOD-ECSSD, and OOD-DUT-OMRON, which are derived from the corresponding public datasets DUTS(Wang et al. 2017), HKU-IS(Li and Yu 2015), ECSSD(Yan et al. 2013), and DUT-OMRON(Yang et al. 2013) by photographing printed scenes at our experimental site. The evaluation is conducted using six standard metrics: Maximum F-measure (Max-F), Mean F-measure (Mean-F), Weighted F-measure (Fbw), Mean Absolute Error (MAE), Structure-measure (S-measure), and Enhanced-alignment Measure (E-measure).

WavePCNet achieves state-of-the-art performance on both the OOD-DUTS and OOD-HKU datasets, as shown in Table 1. Notably, on OOD-DUTS, it achieves a Max-F score of 0.692, which is a 4.0% absolute improvement over the best existing method (RefCOD, 0.670), and reduces

MAE to 0.154, the lowest among all methods, demonstrating superior boundary precision and noise resilience. Similarly, on OOD-HKU, WavePCNet yields the highest Mean-F (0.545) and Fbw (0.439), and the lowest MAE (0.170), outperforming the second-best method (ZoomNeXt) by 1.6% in Mean-F and 2.3% in Fbw. Table 2 shows the results on the OOD-ECSSD and OOD-DUT-OMRON datasets. On OOD-ECSSD, WavePCNet achieves the highest scores in all metrics, including Max-F (0.666) and Fbw (0.533), and achieves the lowest MAE (0.159), an absolute improvement of 5.1% in Max-F and reduction of 0.081 in MAE, compared to the next best method (RefCOD). On OOD-DUT-OMRON, WavePCNet maintains top performance with a Max-F of 0.471, outperforming the best prior method (ZoomNeXt, 0.456) by 1.5%, and achieves the lowest MAE (0.169) and highest SM (0.603). Furthermore, Figure 4 shows the comparison of the detection results. Compared to prior ap-

Method	Publication	OOD-ECSSD						OOD-DUT-OMRON					
		Max-F	Mean-F	Fbw	MAE	SM	EM	Max-F	Mean-F	Fbw	MAE	SM	EM
UNet	MICCAI 2015	0.604	0.557	0.384	0.258	0.628	0.721	0.436	0.402	0.292	0.203	0.583	0.671
LDF	CVPR 2020	0.610	0.565	0.401	0.240	0.636	0.726	0.447	0.407	0.312	0.191	0.588	0.673
F3Net	AAAI 2020	0.608	0.566	0.406	0.237	0.641	0.726	0.435	0.400	0.303	0.195	0.583	0.67
GCPANet	AAAI 2020	0.600	0.56	0.394	0.243	0.635	0.723	0.436	0.399	0.305	0.195	0.584	0.670
EDN	TIP 2022	0.618	0.562	0.435	0.223	0.649	0.723	0.464	0.422	0.326	0.193	0.597	0.681
PFSNet	AAAI 2021	0.612	0.565	0.427	0.232	0.644	0.726	0.447	0.411	0.314	0.192	0.592	0.677
MENet	CVPR 2023	0.605	0.557	0.381	0.241	0.622	0.724	0.442	0.404	0.297	0.207	0.585	0.676
DC-Net	Pattern Recognit. 2024	0.466	0.401	0.302	0.402	0.449	0.595	0.346	0.312	0.218	0.279	0.508	0.628
RefCOD	TPAMI 2025	0.615	0.569	0.435	0.226	0.646	0.727	0.421	0.387	0.288	0.220	0.571	0.656
ZoomNeXt	PAMI2024	0.605	0.546	0.415	0.245	0.630	0.707	0.449	0.408	0.310	0.195	0.590	0.672
FAPNet	TIP2022	0.603	0.559	0.400	0.242	0.638	0.723	0.455	0.411	0.327	0.194	0.593	0.673
FDNet	CVPR2022	0.603	0.554	0.384	0.255	0.627	0.720	0.441	0.402	0.301	0.200	0.585	0.667
Our	-	0.666	0.634	0.533	0.159	0.667	0.787	0.471	0.433	0.341	0.169	0.603	0.695

Table 2: Performance comparison on OOD-ECSSD and OOD-DUT-OMRON datasets under six evaluation metrics.

proaches, WavePCNet consistently produces results more accurate and closer to the ground truth, with sharper boundaries, clearer object shapes, and fewer artifacts. While methods like U-Net, F3Net, and GCPANet often generate blurred or incomplete outputs, and others like DC-Net and MENet fail under heavy occlusion, WavePCNet preserves fine structure and semantic integrity even in challenging scenes.

Overall, real-valued, locality-biased convolutions in existing architectures, constrained by their value-domain modeling, respond poorly to phase aliasing and non-specular diffusion, hence information loss in low-SNR, multiply scattered fields. By contrast, TriWCP in WavePCNet introduces complex-domain propagation with a three-stage optimisation and low-rank spectral constraint; update trajectories remain confined to a physics-consistent propagating subspace, expected to improve recovery of signals from occluded regions. Additionally, Cross-layer Compensation, coupling multi-scale frequency-selective filtering with cross-layer attention guidance, extracts residual high-frequency responses under strong scattering backgrounds and aligns structural semantics across levels, thereby providing finer edge resolution and stronger context awareness for obscured objects.

TriWCP	Compen	Max-F	Mean-F	Fbw	MAE	SM	EM
✗	✓	0.610	0.557	0.424	0.237	0.636	0.715
✓	✗	0.477	0.429	0.284	0.273	0.535	0.682
MobileNet	✓	0.598	0.566	0.454	0.186	0.603	0.745
VGG	✓	0.607	0.569	0.482	0.186	0.618	0.746
ShuffleNetV2	✓	0.592	0.555	0.452	0.196	0.620	0.743
EfficientNet-B0	✓	0.587	0.549	0.452	0.195	0.606	0.734
✓	UNet	0.622	0.583	0.475	0.190	0.641	0.749
✓	FPN	0.616	0.575	0.453	0.195	0.635	0.747
✓	HRNet	0.634	0.592	0.480	0.184	0.654	0.757
✓	Transformer	0.608	0.570	0.431	0.225	0.639	0.734
✓	✓	0.666	0.634	0.533	0.159	0.667	0.787

Table 3: Ablation configurations evaluating the isolated and joint contributions of TriWCP and Cross-layer Compensation (Compen) under indirect, multiply-scattered conditions.

Ablation Studies

To verify the independent efficacy, and the joint contribution of TriWCP and Cross-layer Compensation under indirect, multiply scattered propagation, three experiment groups are configured, the results are shown in Table 3.

Two ablations (removed TriWCP or Compensation respectively): the full model achieves the best performance, with a Max-F of 0.666, Mean-F of 0.634, and the lowest MAE of 0.159. Compared to using only Compensation, the full model improves Max-F by 5.6% and reduces MAE by 0.078, while compared to using only TriWCP, it shows an even larger gain of 18.9% in Max-F and 0.114 reduction in MAE, highlighting the importance of both modules.

Compen with real-valued convolutional backbones(including MobileNet and VGG): VGG achieves the best results with a Max-F of 0.607 and Mean-F of 0.569, while EfficientNet-B0 performs the worst, with a lower Max-F (0.587) and the highest MAE (0.195). Although VGG slightly outperforms others, all backbone-only variants significantly underperform compared to the full WavePCNet, achieving up to a 9.1% increase in Max-F (over EfficientNet-B0) and a 0.036 reduction in MAE (over MobileNet and VGG). These results highlight the performance of TriWCP.

TriWCP with classical fusion frameworks (e.g., UNet and FPN): HRNet performs best (Max-F: 0.634, MAE: 0.184), compared to the full WavePCNet, which achieves a Max-F of 0.666 and MAE of 0.159, the improvement reaches up to a 5.8% increase in Max-F (over Transformer) and 0.066 reduction in MAE. Even over the strongest baseline (HRNet), WavePCNet still improves Max-F by 3.2% and reduces MAE by 0.025.

Overall, the comparative analyses across various backbone configurations demonstrate the critical importance of TriWCP and Cross-layer Compensation beyond simply selecting a strong backbone, and the synergy between physics-guided coherent propagation and frequency-aware structural enhancement is key to robust obscured object detection under passive, non-line-of-sight conditions.

Method	Params (M) ↓	FPS ↑	MAE ↓
UNet	23.95	186.270	0.258
LDF	25.15	116.433	0.240
F3Net	25.54	107.250	0.237
GCPANet	67.06	94.307	0.243
EDN	42.90	65.445	0.223
PFSNet	31.18	93.118	0.232
MENet	27.83	26.986	0.241
DC-Net	284.11	30.346	0.402
RefCOD	24.98	109.188	0.226
ZoomNeXt	28.46	38.646	0.245
FAPNet	24.84	91.110	0.242
FDNet	24.10	178.444	0.255
WavePCNet (Ours)	26.05	39.428	0.159

Table 4: Comparison of computational/storage complexity (Params), runtime throughput (FPS), and reconstruction fidelity (MAE).

Complexity and Performance Trade-off

As shown in Table 4, WavePCNet achieves a favorable complexity–speed–accuracy trade-off: it uses 26.05M parameters—comparable to UNet (23.95M) and F3Net (25.54M), and far smaller than GCPANet (67.06M) or DC-Net (284.11M). It runs at 39.428 FPS, meeting real-time needs (slower than UNet/FDNet but faster than MENet, 26.986 FPS, and comparable to SSN, 37.821 FPS). Most importantly, it delivers the best accuracy with an MAE of 0.159, outperforming EDN (0.223) and RefCOD (0.226), attributable to the proposed physics-driven TriWCP and cross-layer frequency compensation.

Conclusion

This study pioneers the definition and validation of the passive post-corner OOD task. Operating without active illumination and relying solely on faint ambient-scattered returns, WavePCNet enforces coherent propagation consistency in the complex domain, augmented by cross-layer frequency compensation, to jointly localize and segment targets hidden beyond the line of sight. The TriWCP module constrains the optimization process to a physically meaningful low-rank subspace, mitigating instability caused by wavefront distortion. Meanwhile, the Cross-layer Compensation module amplifies residual high-frequency features and restores boundary fidelity through semantics-guided structural alignment. Comprehensive evaluations and ablation studies across four physically collected datasets validate the robustness and complementarity of the proposed components, highlighting their effectiveness in tackling the unique challenges of passive OOD.

Acknowledgments

This work was supported in part by the China Postdoctoral Science Foundation (General Program) under

Grant 2024M760716 and the Special Financial Grant from the China Postdoctoral Science Foundation under Grant 2025T180962; in part by the Fundamental Research Funds for the Provincial Universities of Zhejiang under Grant GK259909299001-026; in part by the Excellent Young Scientists Fund Program (Overseas) of Shandong Province under Grant 2025HWYQ-033; and in part by the Young Scientists Fund of the National Natural Science Foundation of China under Grant 62506103 and 42501545.

References

- Beckus, A.; Tamasan, A.; and Atia, G. K. 2019. Multi-modal non-line-of-sight passive imaging. *IEEE Transactions on Image Processing*, 28(7): 3372–3382.
- Boger-Lombard, J.; and Katz, O. 2019. Passive optical time-of-flight for non line-of-sight localization. *Nature communications*, 10(1): 3343.
- Cao, H.; Mosk, A. P.; and Rotter, S. 2022. Shaping the propagation of light in complex media. *Nature Physics*, 18(9): 994–1007.
- Chen, Z.; Sun, K.; and Lin, X. 2024. CamoDiffusion: Camouflaged object detection via conditional diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1272–1280.
- Chen, Z.; Xu, Q.; Cong, R.; and Huang, Q. 2020. Global context-aware progressive aggregation network for salient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, 10599–10606.
- Khawaja, W.; Guvenc, I.; Matolak, D. W.; Fiebig, U.-C.; and Schneckenburger, N. 2019. A survey of air-to-ground propagation channel modeling for unmanned aerial vehicles. *IEEE Communications Surveys & Tutorials*, 21(3): 2361–2391.
- Li, G.; and Yu, Y. 2015. Visual Saliency Based on Multiscale Deep Features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Liu, X.; Guillén, I.; La Manna, M.; Nam, J. H.; Reza, S. A.; Huu Le, T.; Jarabo, A.; Gutierrez, D.; and Velten, A. 2019. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature*, 572(7771): 620–623.
- Luo, Z.; Liu, N.; Zhao, W.; Yang, X.; Zhang, D.; Fan, D.-P.; Khan, F.; and Han, J. 2024. Vscope: General visual salient and camouflaged object detection with 2d prompt learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17169–17180.
- Ma, M.; Xia, C.; and Li, J. 2021. Pyramidal feature shrinking for salient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, 2311–2318.
- May, M. A.; Barré, N.; Kummer, K. K.; Kress, M.; Ritsch-Marte, M.; and Jesacher, A. 2021. Fast holographic scattering compensation for deep tissue biological imaging. *Nature Communications*, 12(1): 4340.
- Pang, Y.; Zhao, X.; Xiang, T.-Z.; Zhang, L.; and Lu, H. 2024. Zoomnext: A unified collaborative pyramid network for camouflaged object detection. *IEEE transactions on pattern analysis and machine intelligence*, 46(12): 9205–9220.

- Qin, X.; Zhang, Z.; Huang, C.; Dehghan, M.; Zaiane, O. R.; and Jagersand, M. 2020. U2-Net: Going deeper with nested U-structure for salient object detection. *Pattern Recognition*, 106: 107404.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Rotter, S.; and Gigan, S. 2017. Light fields in complex media: Mesoscopic scattering meets wave control. *Reviews of Modern Physics*, 89(1): 015005.
- Ruan, J.; Yuan, W.; Lin, Z.; Liao, N.; Li, Z.; Xiong, F.; Liu, T.; and Fu, Y. 2025. Mm-camobj: A comprehensive multimodal dataset for camouflaged object scenarios. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 6740–6748.
- Scheiner, N.; Kraus, F.; Wei, F.; Phan, B.; Mannan, F.; Appenrodt, N.; Ritter, W.; Dickmann, J.; Dietmayer, K.; Sick, B.; et al. 2020. Seeing around street corners: Non-line-of-sight detection and tracking in-the-wild using doppler radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2068–2077.
- Shen, S.; Wang, Z.; Liu, P.; Pan, Z.; Li, R.; Gao, T.; Li, S.; and Yu, J. 2021. Non-line-of-sight imaging via neural transient fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(7): 2257–2268.
- Wang, B.; Chen, Q.; Zhou, M.; Zhang, Z.; Jin, X.; and Gai, K. 2020. Progressive feature polishing network for salient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, 12128–12135.
- Wang, L.; Lu, H.; Wang, Y.; Feng, M.; Wang, D.; Yin, B.; and Ruan, X. 2017. Learning to Detect Salient Objects With Image-Level Supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, Y.; Wang, R.; Fan, X.; Wang, T.; and He, X. 2023. Pixels, regions, and objects: Multiple enhancement for salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10031–10040.
- Wei, J.; Wang, S.; and Huang, Q. 2020. F³Net: fusion, feedback and focus for salient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, 12321–12328.
- Wei, J.; Wang, S.; Wu, Z.; Su, C.; Huang, Q.; and Tian, Q. 2020. Label decoupling framework for salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13025–13034.
- Wei, J.; Zhou, S. K.; Cui, S.; and Li, Z. 2024. Weakpcsd: Overcoming the bias of box annotations for weakly supervised point cloud salient object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 5894–5902.
- Wu, Y.-H.; Liu, Y.; Zhang, L.; Cheng, M.-M.; and Ren, B. 2022. EDN: Salient object detection via extremely-downsampled network. *IEEE Transactions on Image Processing*, 31: 3125–3136.
- Yan, Q.; Xu, L.; Shi, J.; and Jia, J. 2013. Hierarchical Saliency Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yang, C.; Zhang, L.; Lu, H.; Ruan, X.; and Yang, M.-H. 2013. Saliency Detection via Graph-Based Manifold Ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yin, B.; Zhang, X.; Fan, D.-P.; Jiao, S.; Cheng, M.-M.; Van Gool, L.; and Hou, Q. 2024. Camoformer: Masked separable attention for camouflaged object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Yuan, R.; Wang, S.; Liu, G.; and Peng, M. 2024. Non-line-of-sight ultraviolet positioning using linearly-arrayed photon-counting receivers. *IEEE Journal on Selected Areas in Communications*, 42(9): 2535–2549.
- Zhang, X.; Yin, B.; Lin, Z.; Hou, Q.; Fan, D.-P.; and Cheng, M.-M. 2025. Referring camouflaged object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhao, H.; Hu, S.; Zhang, H.; Wang, Z.; Dong, H.; Del Hougne, P.; Cui, T. J.; and Li, L. 2023. Intelligent indoor metasurface robotics. *National Science Review*, 10(8): nwac266.
- Zhong, Y.; Li, B.; Tang, L.; Kuang, S.; Wu, S.; and Ding, S. 2022. Detecting camouflaged object in frequency domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4504–4513.
- Zhou, T.; Zhou, Y.; Gong, C.; Yang, J.; and Zhang, Y. 2022. Feature aggregation and propagation network for camouflaged object detection. *IEEE Transactions on Image Processing*, 31: 7036–7047.
- Zhu, J.; Qin, X.; and Elsaddik, A. 2025. Dc-net: Divide-and-conquer for salient object detection. *Pattern Recognition*, 157: 110903.