SYNCHROINFO JOURNAL

www.media-publisher.eu/synchroinfo-journal

VISION-ASSISTED PROACTIVE HANDOVER AND RESOURCE ALLOCATION IN HYBRID LIFI/RF 6G NETWORKS USING EDGE AI

Ali Salah Mahdi Khafaga [1]

[1] Department of Computer Networks and Software, Al-Furat Al-Awsat Technical University, Babil, Iraq; ali.khafaja@atu.edu.iq
https://orcid.org/0000-0001-5792-1511
https://www.scopus.com/authid/detail.uri?authorId=59499331800

ABSTRACT

Light fidelity (LiFi) offers large unlicensed bandwidth for indoor 6G access, but its line-of-sight links are interrupted whenever a person walks between a luminaire and a device, or when the device tilts out of the receiver field of view. Handover schemes in hybrid LiFi/RF networks react only after the link is lost and therefore leave users without service for the detection and execution time of the handover. In this paper, we use the cameras that already monitor many smart buildings to anticipate these outages. An edge-AI pipeline tracks every person, converts the scene geometry into link-level features, fuses them with the measured signal-to-noise ratio, and predicts the probability that each LiFi link becomes unusable within the next second. A deep reinforcement learning agent, pre-trained from demonstrations and fine-tuned online, uses this forecast to select the access point of every user, while proportional-fair time sharing allocates resources inside each cell. In a room with six users and up to fifteen walking pedestrians, the learned vision-radio predictor raises the area under the ROC curve for 200-ms outage prediction from 0.79 (Kalman/Monte-Carlo forecast) to 0.89, and the proposed controller reduces link-interruption time by 42-47% with respect to a standard event-A3 handover, with no loss in mean throughput. Without retraining, the controller also reduces interruption by 53% when the pedestrians follow real trajectories from the ETH/UCY datasets.

KEYWORDS: *LiFi, hybrid LiFi/WiFi networks, handover, blockage prediction, computer vision, edge AI, deep reinforcement learning, 6G*.

DOI: 10.36724/2664-066X-2026-12-4-2-18

Received: 20.04.2026
Accepted: 22.06.2026

Citation: Ali Salah Mahdi Khafaga, "A. Vision-Assisted Proactive Handover and Resource Allocation in Hybrid LiFi/RF 6G Networks Using Edge AI," *Synchroinfo Journal* **2026**, vol. 12, no. 4, pp. 2-18.

Licensee IRIS, Vienna, Austria.

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Copyright: © 2026 by the authors.

1 Introduction

LiFi turns ceiling luminaires into access points (APs) and adds hundreds of megahertz of license-free spectrum to indoor networks [1]; the IEEE 802.11bb amendment now standardizes light communications in wireless local area networks [2]. Because each LiFi cell is small and relies on a line-of-sight (LoS) path, practical deployments combine LiFi with a radio-frequency (RF) network such as WiFi in a hybrid LiFi/WiFi network (HLWNet) [3]. In such a network, WiFi guarantees coverage and LiFi carries most of the traffic.

The optical link, however, is fragile. A person who walks between a luminaire and a device blocks the LoS path within tens of milliseconds, the body of the user can block the luminaires behind them, and a natural change of the device tilt can move an AP outside the receiver field of view (FOV) [4]. Conventional handover (HO) schemes in HLWNets, including the LTE-style event-A3 rule and learning-based variants [5-8], decide from radio measurements only. They detect an outage after it has happened and then spend a further HO execution time of hundreds of milliseconds before the user is served again. For latency-sensitive traffic, it is these interruptions, rather than the average rate, that users notice.

A camera, by contrast, can see a blocker approaching before the link is lost. Charan *et al.* showed that cameras at mmWave base stations can predict blockage and trigger a proactive hand-off [9], and later work has extended vision-aided beam and blockage prediction using real-world multi-modal datasets [10-12]. LiFi has been studied as a sensing modality together with cameras [13], but, to the best of our knowledge, vision has not been used to drive handover in hybrid LiFi/RF networks. These mmWave results do not carry over directly to LiFi. In mmWave the obstacle crosses a long, roughly horizontal beam. In LiFi the beam is short and steep, so only people very close to the user can block it; the user's own body is a frequent blocker; and the orientation of a hand-held device matters as much as the position of the people around it.

In this paper we address this gap. Our main contributions are as follows.

- **LiFi-specific vision-to-channel mapping.** We derive the *blocking zone* of a LiFi link, the part of the LoS segment low enough to be cut by a human body, and turn the tracked scene into link-level features: the clearance of every person from this zone now and extrapolated over the next second, self-blockage geometry, and the field-of-view margin implied by the estimated device tilt.
- **Learned bimodal outage predictor.** A lightweight network fuses the vision features with the measured signal-to-noise ratio (SNR) history and outputs, for every user–AP link, a monotone curve of outage probabilities over a 1-s horizon and the expected outage fraction. Training with randomized perception quality makes it robust to detector noise, missed detections and latency.
- **Joint proactive handover and resource allocation.** We formulate access-point selection with proportional-fair (PF) time allocation as a constrained stochastic control problem, show that PF allocation reduces to equal time sharing for a given association, and solve the association with a Double-DQN agent that is pre-trained from an outage-aware expert rule and fine-tuned with reinforcement learning.
- **Reproducible evaluation.** A frame-level simulator with measured orientation statistics, self-blockage, make-before-break handover and real pedestrian trajectories (ETH/UCY [14, 15]) compares the proposal with reactive, event-A3, vision-free DRL, physics-based vision and oracle schemes, and quantifies prediction accuracy, the throughput-reliability trade-off, robustness and complexity.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 presents the system model and the problem formulation. Section 4 describes the vision pipeline and the outage predictors, and Section 5 the handover controllers. Section 6 details the evaluation set-up, Section 7 discusses the results, Section 8 the limitations, and Section 9 concludes the paper.

2 Related work

2.1 Handover and Load Balancing in HLWNets

Handover in HLWNets is either horizontal (HHO), between two LiFi APs, or vertical (VHO), between LiFi and WiFi, and both incur signaling overhead that interrupts the service [3]. Wang and Haas combined load balancing with handover to account for this overhead [5]. Learning-based schemes have since been proposed: Frómeta Fonseca *et al.* classify the trajectory type of a user and use Q-learning for the HO decision, improving throughput by 59% over a smart HO baseline [6]; a dueling DQN jointly reduces delay and HO frequency [7]; and a sequential proximal policy optimization method improves sum throughput in user association [8]. Blockage has been handled with diversity receivers and rule-based resource allocation [16]. All of these use radio-side information only, so they cannot anticipate a blockage that has not yet affected the signal.

2.2 Device Orientation in LiFi

Soltani *et al.* measured how people hold their phones and showed that the polar angle follows a Laplace distribution for static users and a Gaussian distribution for walking users, with a significant impact on the LoS channel gain and on handover rate [4]. Existing HO schemes treat orientation as a random variable, whereas a camera can estimate it for each user individually.

2.3 Vision-Aided Wireless Communications

Vision-aided blockage prediction has been demonstrated for mmWave links with bimodal deep learning that combines RGB images and beam power [9], supported by real-world datasets [10]. Recent work reports over 98% accuracy three frames ahead for mmWave blockage [11] and vision-aided mobility management for terahertz links [12]. Camera, LiFi and RF data have also been fused for indoor sensing [13]. Our work differs in three ways: it targets the steep, short LiFi geometry including self-blockage and tilt; it predicts a probability *curve* over the horizon instead of a binary label; and it closes the loop with a controller that explicitly trades interruption against throughput in a multi-user hybrid network.

3 System model and problem formulation

3.1 Network and Geometry

We consider an indoor room of size L x W x H served by a set $A = \{1,\ldots,A\}$ of ceiling LiFi APs and one WiFi AP, indexed by 0 (Fig. 1). Neighbouring LiFi APs use orthogonal bands, so inter-cell interference is neglected. A set of users carries hand-held devices; a further set of non-connected pedestrians moves in the room. Time is slotted with the camera frame period $\Delta = 1/30^{o}$s, and decisions are taken every $D_s = 2$ frames.

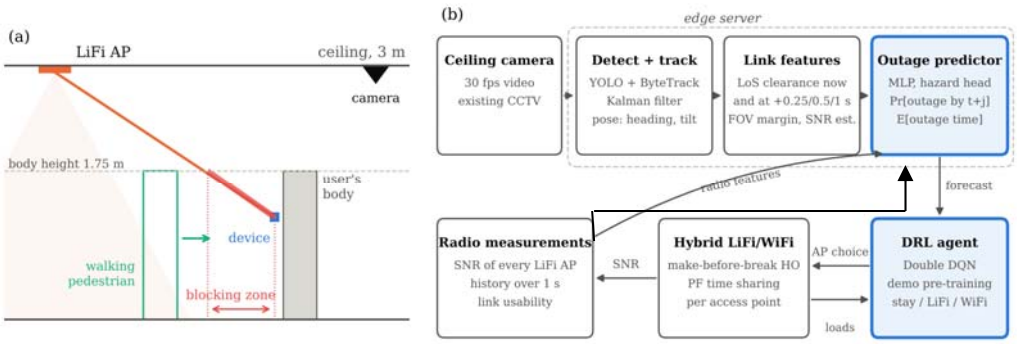


Fig. 1. (a) Side view of a LiFi link: only the segment of the LoS path lower than the body height (the blocking zone) can be cut by a person, so the relevant blockers are the user's own body and people within about a metre of the device. (b) Edge-AI pipeline: tracked people are turned into link features, fused with radio measurements, and the predicted outage curve drives the DRL handover agent.

Each person is modelled as a vertical cylinder of radius r_b and height h_b. User u stands at ground position $\mathbf{p}_u(t)$ facing direction $\phi_u(t)$ and holds the device at height hd, a distance δ in front of the body, i.e. $\mathbf{d}_u(t) = \mathbf{p}_u(t) + \delta\,[\cos\phi_u, \sin\phi_u]^T$.

3.2 LiFi Channel

The LoS channel gain between AP a at height H and the device of user u is [17]

$$G_{u,a} = \frac{(m+1)A_{\mathrm{PD}}}{2\pi d_{u,a}^2}\ \cos^m(\varphi_{u,a})\,T_s\,g(\psi_{u,a})\cos(\psi_{u,a}),\quad 0 \le \psi_{u,a} \le \Psi_c, \tag{1}$$

And $G_{u,a} = 0$ otherwise, where $m = -\ln 2/\ln(\cos\Phi_{1/2})$ is the Lambertian order, A_{PD} the photodiode area, $d_{u,a}$ the distance, $\varphi_{u,a}$ the irradiance angle, Ts the filter gain, $g(\Psi) = n^2_c/\sin^2\Psi_c$ the concentrator gain and Ψ_c the FOV. The incidence angle follows from the device normal

$$\cos\psi_{u,a} = \frac{\mathbf{n}_u^T(\mathbf{x}_a - \mathbf{x}_u)}{\|\,\mathbf{x}_a - \mathbf{x}_u\,\|},\qquad \mathbf{n}_u = \begin{bmatrix}\sin\theta_u\cos\omega_u\\ \sin\theta_u\sin\omega_u\\ \cos\theta_u\end{bmatrix}, \tag{2}$$

where x denotes 3-D positions, θ_u is the polar tilt and $\omega_u = \phi_u + \pi$ because the screen faces the user. Following the measurements in [4], θ_u is a Gaussian process with mean 29.67° and standard deviation 7.78° while walking (41.39°, 7.68° while static) and an AR(1) correlation with coherence time τ_θ.

The electrical SNR and the achievable rate of DC-biased optical OFDM are [18]

$$\gamma_{u,a} = \frac{\left(\kappa P_{\mathrm{opt}} G_{u,a}\right)^2}{N_0 B_{\mathrm{L}}},\qquad R_{u,a} = \frac{B_{\mathrm{L}}}{2}\log_2\left(1 + \frac{e}{2\pi}\gamma_{u,a}\right), \tag{3}$$

where k is the PD responsivity, P_{opt} the optical power, N_0 the noise power spectral density and B_L the LiFi bandwidth. A link is usable when it is not blocked and $10\log_{10}\gamma_{u,a} \geq \gamma_{\min}$.

3.3 Blocking Zone and Self-Blockage

Let s_a and e_u be the ground projections of AP a and device u. Along the LoS segment the height decreases linearly from H to h_d, so a body of height h_b can only intersect the fraction

$$s \in [s_0, 1], \qquad s_0 = \frac{H - h_b}{H - h_d}, \tag{4}$$

of the segment closest to the device. With H = 3°m, h_b = 1,75°m and h_d = 1,2 °m, s_0 = 0,69: only the last 31% of the horizontal projection, typically less than one metre, can be blocked. Link (u, a) is blocked at time t if any person k, including user u itself, satisfies

$$\beta_{u,a}(t) = \mathbb{1}\left\{\min_k \operatorname{dist}(\mathbf{c}_k(t), [\mathbf{s}_a + s_0(\mathbf{e}_u - \mathbf{s}_a), \mathbf{e}_u]) < r_b\right\}, \tag{5}$$

where c_k is the body centre and *dist* the point-to-segment distance. Including $k = u$ captures self-blockage of the APs behind the user. Equation (4) is the main geometric difference from mmWave blockage, where the beam is long and nearly horizontal.

3.4 WiFi Channel

The WiFi AP uses bandwidth B_R at 2.4 GHz with the IEEE 802.11 TGn path-loss model (breakpoint 5 m) [19]. Its single-user throughput is $R_{u,0} = \eta B_R \min\{\log_2(1 + \gamma_u^{RF}), c_{\max}\}$, where $c_{\max}$ caps the spectral efficiency at the highest modulation and coding scheme and η accounts for MAC overhead.

3.5 Handover and Resource Allocation

Let $x_{u,a}(t) \in \{0,1\}$ indicate that user u is associated with AP $a \in \{0\} \cup \mathcal{A}$, and $\tau_{u,a}(t) \in [0,1]$ the fraction of time AP a serves user u. A HO from a to a' takes T_{HHO} if both are LiFi APs and T_{VHO} otherwise. We assume make-before-break execution: the old link keeps serving the user until the HO completes, so a HO started early enough causes no interruption. If the serving LiFi link is lost, every scheme falls back to a reactive HO after a detection delay of T_{det}.¶

For a given association, the time allocation inside each cell maximizes the PF utility.¶

Lemma 1: For AP a serving the set $\mathcal{U}_a(t)$ of users with usable links, $\max_\tau \sum_{u \in \mathcal{U}_a} \log(\tau_{u,a} R_{u,a})$ subject to $\sum_u \tau_{u,a} \le 1$ is solved by $\tau^\star_{u,a} = 1/|\mathcal{U}_a(t)|$.¶

Proof: The objective equals $\sum_u \log \tau_{u,a} + \sum_u \log R_{u,a}$; the second term does not depend on τ, and maximizing $\sum_u \log \tau_{u,a}$ on the simplex gives equal shares by symmetry and concavity (KKT).¶

The rate of user u is then $r_u(t) = \sum_a x_{u,a}(t)\, R_{u,a}(t)/|\mathcal{U}_a(t)|$, and resource allocation is coupled to association through the loads $|\mathcal{U}_a|$.¶

3.6 Problem Formulation

We define the link-interruption indicator $\rho_u(t) = \mathbb{1}\{r_u(t) = 0\}$, which is one while the serving LiFi link is unusable and the user has not yet been handed over. The joint handover and resource-allocation problem is

$$\begin{aligned} \text{P1:} \max_{\{x(t)\}} \quad & \mathbb{E}\left[\sum_{t}\sum_{u\in\mathcal{U}}\left(\frac{r_u(t)}{R_{\text{ref}}} - \lambda\, o_u(t) - c_{\text{HO}}\, h_u(t)\right)\right] \\ \text{s.t.} \quad & \sum_{a} x_{u,a}(t) = 1, \quad x_{u,a}(t) \in \{0,1\}, \\ & \tau_{u,a}(t) = \tau^{*}_{u,a}(t) \text{ (Lemma 1)}, \\ & \text{HO execution constraints with } T_{\text{HHO}}, T_{\text{VHO}}, \end{aligned} \tag{6}$$

where $h_u(t)$ indicates a HO started by user u at time t, R_{ref} = 100 Mb/s normalizes throughput, and γ sets the cost of one unit of interruption time relative to throughput. P1 is a mixed-integer stochastic control problem whose state includes the future motion of every person in the room. It cannot be solved by per-slot optimization because the effect of a decision appears only after the HO execution time. We therefore (i) forecast link outages from vision and radio data (Section 4) and (ii) solve P1 as a Markov decision process (Section 5).

4 Vision-assisted outage prediction

4.1 Edge Perception

A ceiling camera feeds an edge server running a person detector [20] and a multi-object tracker [21]. Detections are projected to the floor plane through a calibrated homography, and pose keypoints give the facing direction of every person and the arm/hand pose of each user, from which the device tilt is estimated. In this study, the pipeline is modelled at the level of its outputs: floor positions with Gaussian error σ_p, independent missed detections with probability p_{miss}, a processing latency of L_v frames, heading error 10° and tilt error 5°. We assume that track identities are maintained and that the network knows which tracked person carries which device, for example from LiFi uplink positioning. All three perception parameters are swept in Section 7.

Each person is tracked with a constant-velocity Kalman filter with state $\mathbf{z}_k = [\mathbf{c}_k^T, \dot{\mathbf{c}}_k^T]^T$, transition $\mathbf{z}_k(t+1) = \mathbf{F}\mathbf{z}_k(t) + \mathbf{w}_k$ and process noise driven by a random acceleration of standard deviation σ_a. The filter output is propagated over the latency L_v to the current time.

4.2 Physics-Based Monte-Carlo Forecast

As a model-based reference, we draw K samples from the Kalman posterior of every person and of the device tilt, propagate them over the horizon J decision steps with the motion model, and evaluate (1)-(5) for every sample. The forecast for link (u,a) is

$$\hat{P}_{u,a}(j) = \frac{1}{K}\sum_{k=1}^{K} \mathbb{1}\left\{\exists\, i \le j:\ \bar{v}_{u,a}^{(k)}(i) = 1\right\}, \tag{7}$$

where $\bar{v}_{u,a}^{(k)}(i)$ indicates that the link is unusable in sample k at step i. This is the probability that the link fails at least once within j steps, together with the expected outage fraction $\hat{\eta}_{u,a}$ over the horizon. A deterministic variant uses the posterior mean only.

4.3 Learned Bimodal Predictor

In practice, the Monte-Carlo forecast is only as good as its motion model. Real people stop, turn and sway in ways that a constant-velocity model cannot capture, and the tilt process is only partly predictable. We therefore learn the outage curve directly. For every link and decision step, the input vector concatenates

- *vision features*: user speed and turn rate; distance and bearing of the AP relative to the facing direction; estimated tilt, FOV margin $\psi_{u,a} - \Psi_c$ and SNR from (1)–(3); number of people within 1.5 m; and the clearance of the user's own body and of the closest other person from the blocking zone (5), evaluated now and on constant-velocity extrapolations at $+0.25$, $+0.5$ and $+1$ s;
- *temporal context*: the change of the vision features over the last 133 and 267 ms;
- *radio features*: the measured SNR of the link, its change over the last 67–533 ms, and recent usability flags;
- *prior*: the Monte-Carlo forecast (7).

A three-layer perceptron with 192 units per layer maps this 77-dimensional input to per-step hazards $h_j \in (0,1)$ and marginal outage probabilities m_j. The cumulative forecast

$$\hat{P}_{u,a}(j) = 1 - \prod_{i=1}^{j} (1 - h_i) \tag{8}$$

is monotone in j by construction, and $\hat{q}_{u,a} = \frac{1}{J}\sum_j m_j$. The network is trained with the binary cross-entropy of $\hat{P}(j)$ against the realized first-outage indicator, masked to links that are usable at the decision time, plus one half of the cross-entropy of m_j.

Two design choices were needed to make the predictor usable in practice. First, it is trained on episodes with randomized perception quality (σ_p up to 0.3 m, p_{miss} up to 0.4, L_v up to 6 frames) in addition to nominal episodes, so that a degraded camera does not invalidate it. Second, the forecasts used to train the controller are produced out of fold (each training episode is predicted by a model that did not see it), which avoids training the controller on over-confident in-sample predictions.

5 Proactive handover control

5.1 Outage-Aware Expert Rule

Let s be the serving AP of user u, n_a the number of users on AP a and $\tilde{R}_{u,a} = R_{u,a}/(n_a + \mathbb{1}\{a \neq s\})$ the rate share the user would obtain. The value of a candidate AP over the forecast horizon is

$$V_{u,a} = \tilde{R}_{u,a}\left(1 - \hat{q}_{u,a}\right) - \lambda R_{\text{ref}}\hat{q}_{u,a}, \tag{9}$$

with $\hat{q}_{u,0} = 0$ for WiFi. Let g be the alternative with the largest value and j_g the number of decision steps needed to complete a HO to it. The rule starts a proactive HO when

$$\hat{P}_{u,s}(j_g) > \epsilon \quad \text{and} \quad V_{u,g} > V_{u,s}, \tag{10}$$

that is, when the serving link is likely to fail before a HO could complete and the alternative is worth it. Otherwise it applies event A3 (hysteresis and time-to-trigger), refusing targets that are themselves predicted to fail; WiFi users return to the LiFi AP with the largest value among those predicted to stay usable through a VHO. With the learned forecast this rule is called *Vision-Learned*; with (7) it is called *Vision-MC*.

5.2 Deep Reinforcement Learning Agent

The rule fixes the trade-off through two hand-tuned thresholds, whereas the DRL agent learns it from experience. P1 is cast as a semi-Markov decision process shared by all users.

- *State* $\mathbf{o}_u$ (16 features): association type; current, best-alternative and WiFi rate shares; loads of the serving, best and WiFi APs; previous rate share; time since the link was lost; forecast $\hat{P}_{u,s}$ at the HHO, VHO and full horizons and $\hat{q}_{u,s}$; and $\hat{P}_{u,b}$, $\hat{q}_{u,b}$ of the best alternative b.
- *Action* $a_u \in \{$stay, HO to b, HO to WiFi$\}$, taken only when no HO of user u is executing.
- *Reward*: the per-slot term of (6), accumulated and discounted over the slots that elapse until the user can act again, $r = \sum_{i=0}^{k-1} \gamma^i \rho_{t+i}$, with the next state discounted by γ^k.

A Double DQN [22, 23] with two hidden layers of 128 units approximates $Q(o, a)$ with target

$$y = r + \gamma^k Q_{\bar{\theta}}\left(\mathbf{o}', \operatorname{argmax}_{a'} Q_\theta(\mathbf{o}', a')\right). \tag{11}$$

When trained as a plain DQN, the agent learned to chase throughput but did not learn to avoid interruptions: the costly events are rare and appear only after the HO execution time, so the reward signal is sparse. We therefore follow deep Q-learning from demonstrations [24]: the agent first learns from the Vision-Learned rule with the TD loss plus the large-margin loss

$$\mathcal{L}_{\mathrm{E}} = \max_a [Q(\mathbf{o}, a) + \ell(a_{\mathrm{E}}, a)] - Q(\mathbf{o}, a_{\mathrm{E}}), \tag{12}$$

where a_{E} is the expert action and $\ell(a_{\mathrm{E}}, a) = 0.8$ for $a \neq a_{\mathrm{E}}$; it is then fine-tuned with ϵ-greedy interaction while one quarter of every mini-batch is still drawn from the demonstrations. The checkpoint with the best objective on separate validation episodes is kept. The resulting agent is called *DRL-Vision*; a vision-free agent trained the same way from event-A3 demonstrations serves as the *DRL (no vision)* baseline. Algorithm 1 summarizes one decision step.

Algorithm 1: Vision-assisted proactive handover (one decision step)

Track all persons; estimate heading and tilt of every user
for each user u and LiFi AP a
Build vision, temporal and radio features of link (u, a)
Predict $\hat{P}_{u,a}(1{:}J)$ by (8) and $\hat{q}_{u,a}$
end for
for each user u not executing a HO
if serving LiFi link lost for T_{det} : reactive HO to best usable AP
else $a_u \leftarrow \operatorname{argmax}_a Q_\theta(\mathbf{o}_u, a)$; start HO if $a_u \neq$ stay
end if
end for
Allocate time inside each cell by Lemma 1

6 Evaluation set-up

Table 1

Simulation Parameters

Parameter	Value
Room, LiFi APs	$8 \times 8 \times 3$ m, 2×2 grid at 2/6 m
WiFi AP	corner, 2.4 GHz, 20 MHz, 20 dBm
LED semi-angle, P_{opt}	$60°$, 3 W per AP
PD area, FOV, n_{c}, κ	1 cm^2, $80°$, 1.5, 0.53 A/W
N_0, B_{L}, $\gamma_{\min}$	10^{-21} A^2/Hz, 40 MHz, 10 dB
Body r_{b}, h_{b}; device h_{d}, δ	0.2 m, 1.75 m; 1.2 m, 0.3 m
Users / pedestrians	6 / 0–15
Walking speed, pauses	0.6–1.4 m/s, 1–8 s
Tilt statistics	from [4], $\tau_\theta = 0.5$ s
T_{HHO}, T_{VHO}, T_{det}	200 ms, 500 ms, 67 ms
A3 hysteresis, TTT	3 dB, 200 ms
Camera; σ_{p}, p_{miss}, L_{v}	30 fps; 6 cm, 0.05, 1 frame
Horizon J, MC samples K	1 s (15 steps), 32
λ, ϵ, c_{HO}, γ	2, 0.3, 0.1, 0.95

Table 1 lists the parameters. Every episode lasts 30 s. In Scenario A, users and pedestrians follow a random-waypoint model with pauses and a slow heading wobble. In Scenario B, the six users are static (seated or standing at random positions and orientations) and the pedestrians are real people: an 8 x 8 m window and a 30-s interval are cut from one of six ETH/UCY scenes [14, 15] (taken from the Trajectron++ release [25]), and their tracks are interpolated to 30 fps. All learned components are trained only on Scenario A; Scenario B is a zero-shot test.We compare eight schemes: *Reactive* (outage-triggered HO only), *STD-LTE (A3)*, *DRL (no vision)*, *Vision-Det* (deterministic Kalman forecast), *Vision-MC*, *Vision-Learned*, *DRL-Vision*, and an *Oracle* that applies the rule with perfect knowledge of future outages.

All schemes share the same channel realizations and the same reactive fallback, so differences are paired. Training uses 48 nominal and 48 perception-randomized episodes, model selection uses 6 separate episodes, and every test uses its own seeds (20 episodes per density, 10 per sweep point, 30 in Scenario B). We report the mean user throughput, the link-interruption time (fraction of time with r_u = 0), HOs per user per minute, the false-alarm ratio (proactive HOs away from a link that would not have failed within T_{HO} = 0,5 s), and 95% confidence intervals.

7 Results and discussion

7.1 Outage Prediction Accuracy

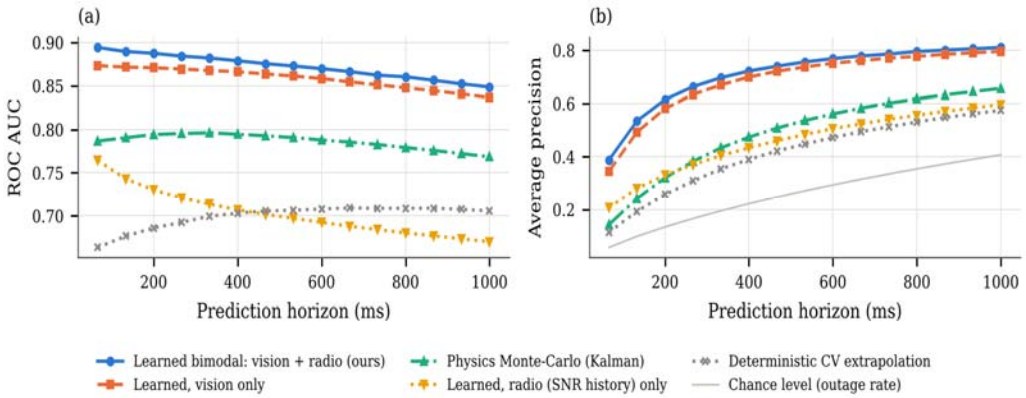


Figure 2. Accuracy of link-outage forecasts versus horizon (Scenario A, 0–15 pedestrians, 30 test episodes, links usable at the decision time). (a) Area under the ROC curve. (b) Average precision; the grey line is the outage rate, i.e. the precision of a random predictor.

Table 2

Outage-Prediction Ablation (ROC AUC / Average Precision)

Predictor	**200 ms**	**533 ms**	**1 s**
Deterministic Kalman	0.69 / 0.26	0.71 / 0.45	0.71 / 0.58
Monte-Carlo Kalman	0.79 / 0.32	0.79 / 0.54	0.77 / 0.66
Learned, radio only	0.73 / 0.33	0.70 / 0.48	0.67 / 0.60
Learned, vision only	0.87 / 0.58	0.86 / 0.74	0.84 / 0.80
Learned, vision + radio, no MC prior	0.88 / 0.59	0.87 / 0.74	0.84 / 0.80
Learned, vision + radio	0.88 / 0.60	0.87 / 0.75	0.84 / 0.80
+ perception randomization (final)	**0.89 / 0.62**	**0.87 / 0.76**	**0.85 / 0.81**
Outage rate (chance AP)	0.13	0.27	0.41

Figure 2 and Table 2 compare the forecasts. The deterministic Kalman extrapolation is the weakest, because a single trajectory either crosses the blocking zone or not and it ignores tilt uncertainty. Sampling the posterior (Monte-Carlo) raises the AUC at 200 ms from 0.69 to 0.79.

The learned predictor gains a further 0.10 in AUC and almost doubles the average precision (0.32 to 0.62), which matters because outages are rare at short horizons. According to the ablation, most of this gain comes from the vision features: radio-only learning reaches only 0.73, whereas vision-only learning reaches 0.87.

The radio history adds a small but consistent improvement, which we attribute mainly to tilt-induced FOV losses, which appear in the SNR trend before the link drops. Removing the Monte-Carlo prior does not change the accuracy, so the Kalman sampling can be dropped at deployment; training with randomized perception quality improves every horizon.

7.2 Network Performance Versus Crowd Density

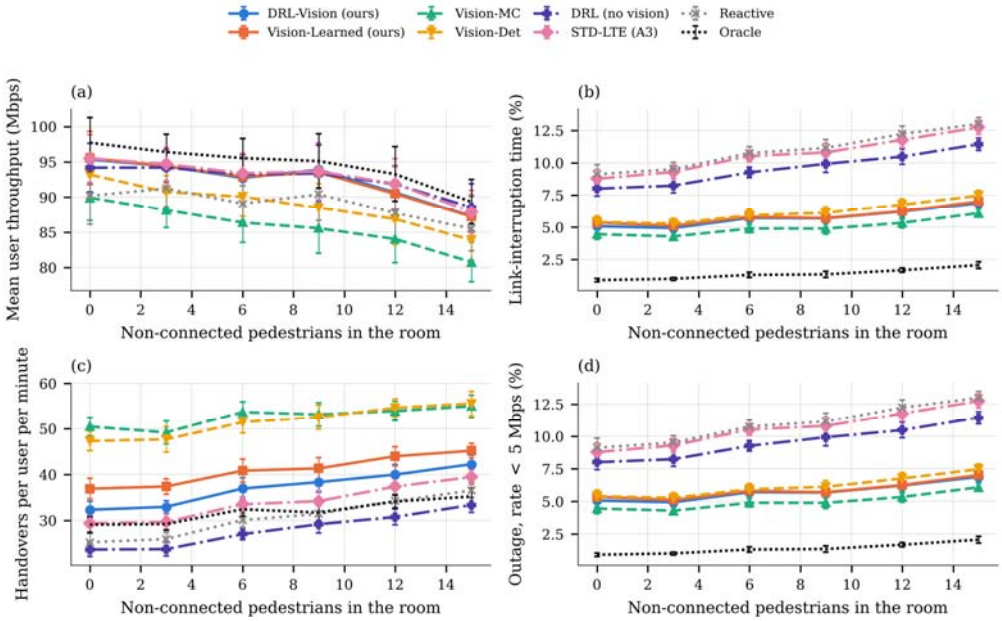


Figure 3. Scenario A, six users and 0–15 non-connected pedestrians (20 episodes per point, 95% confidence intervals). (a) Mean user throughput. (b) Link-interruption time. (c) Handovers per user per minute. (d) Outage probability (rate below 5 Mb/s).

Table 3

Scenario A With Nine Pedestrians (Mean ± 95% Confidence Interval Over 20 Episodes)

Scheme	Throughput (Mb/s)	Interruption (%)	No. of interruptions	Median length (ms)	HO/user/min	False alarms (%)	On WiFi (%)
Reactive	90.4 ± 3.8	11.15 ± 0.65	1522	233	31.4	–	21
STD-LTE (A3)	93.7 ± 4.0	10.81 ± 0.70	1480	233	34.2	–	22
DRL (no vision)	93.4 ± 4.0	9.93 ± 0.65	1345	233	29.2	60	29
Vision-Det	88.5 ± 3.4	6.12 ± 0.47	1086	167	52.5	43	34
Vision-MC	85.6 ± 3.5	4.89 ± 0.43	838	167	53.2	49	40
Vision-Learned	93.5 ± 3.9	5.71 ± 0.42	1000	167	41.3	26	31
DRL-Vision	**93.9 ± 3.7**	**5.69 ± 0.42**	1018	167	38.3	27	31
Oracle	95.1 ± 3.8	1.36 ± 0.24	387	100	31.7	12	32

Figure 3 and Table 3 report the network-level results. Interruption grows with crowd density for every scheme, because each additional person is a potential blocker; at zero pedestrians the outages are caused by self-blockage and tilt alone. The two radio-only baselines behave alike: event A3 improves throughput over the purely reactive scheme by moving users to stronger APs, but neither can act before a link is lost, and both leave users without service for about 11% of the time. The vision-free DRL agent saves less than one percentage point.

All vision-assisted schemes reduce interruption considerably, although at different costs. Vision-MC achieves the lowest interruption among the realizable schemes (4.9%) but loses 8.6% of throughput, because half of its proactive handovers are false alarms that move users to the slower WiFi network (40% of the time). The learned forecast halves the false-alarm ratio (26-27%). As a result, Vision-Learned and DRL-Vision reduce interruption by 39-47% and 42-47%, respectively, across all densities while keeping throughput within the confidence interval of STD-LTE. At nine pedestrians, DRL-Vision reduces interruption from 10.81% to 5.69% (p = 9 x 10^{-5}, paired Wilcoxon test) and changes throughput by +0,3% (p = 0,55). It also uses 7% fewer handovers than the rule with the same forecast (38.3 vs. 41.3 per minute), although it still hands over 12% more often than STD-LTE.

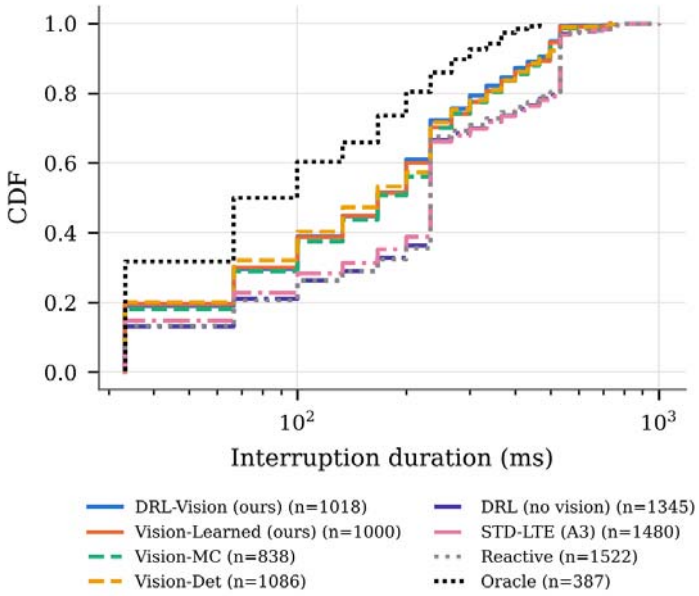


Figure 4. CDF of the duration of individual interruptions (nine pedestrians); *n* is the number of interruptions over 20 episodes.

Figure 4 shows how the interruptions themselves change. Proactive schemes have fewer interruptions (1018 instead of 1480) and shorter ones: the median duration falls from 233 ms, the detection plus horizontal-HO time of a reactive switch, to 167 ms. The remaining long interruptions are blockages that the forecast missed or that started while a handover was executing.

7.3 Throughput–Reliability Trade-Off

Because a single operating point can hide a trade-off, we also swept the parameters of every scheme (Fig. 5). No setting of STD-LTE brings interruption below 9.4%: without foresight, faster triggering only exchanges throughput for more ping-pong. The vision-assisted schemes reach operating points that are not available to radio-only control. The frontier of the learned forecast dominates both Kalman-based forecasts and every STD-LTE configuration: it reaches 96.1 Mb/s, above the best STD-LTE throughput (95.7 Mb/s), and it can bring interruption down to 2.6% if about 9% of throughput is sacrificed.

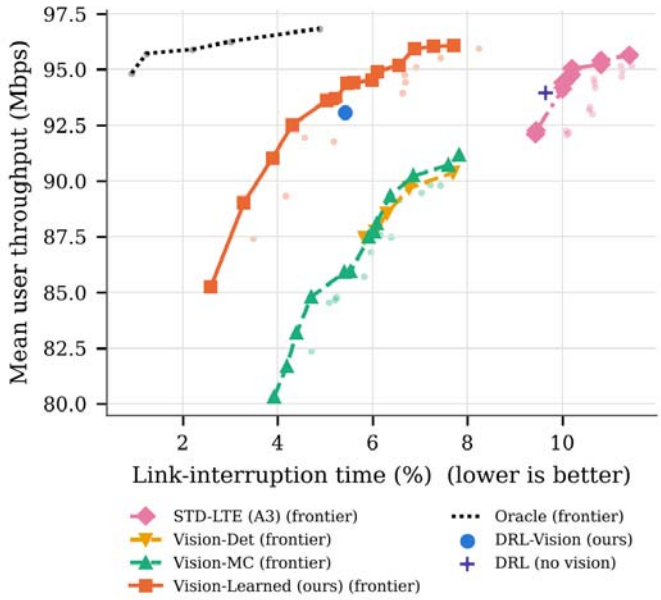


Figure 5. Throughput versus interruption on validation episodes (nine pedestrians). Each rule-based scheme is swept over $\epsilon \in \{0.1, \ldots, 0.7\}$ and $\lambda \in \{0, \ldots, 4\}$; STD-LTE over 27 combinations of hysteresis, time-to-trigger and detection delay. Lines connect Pareto-optimal points.

DRL-Vision, trained for one value of λ, lies within about 1 Mb/s of this frontier without any threshold tuning. It does not move beyond the frontier of the tuned rule, so in this study the main source of the gain is the forecast rather than the learning-based controller. The gap to the Oracle frontier (below 1% interruption at 95 Mb/s) shows how much further improvement is possible from better prediction.

7.4 Robustness to Perception Errors

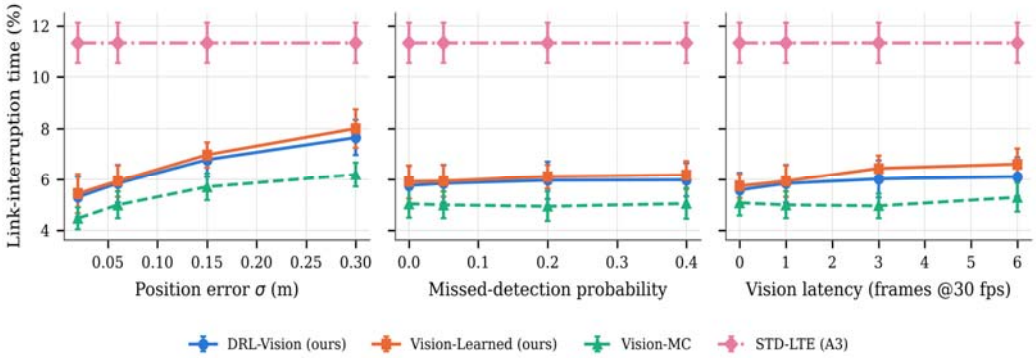


Figure 6. Interruption time versus perception quality (nine pedestrians, 10 episodes per point): floor-position error, missed-detection probability, and pipeline latency.

Table 4

Forecast ROC AUC at 200 ms Under Perception Shift

Test condition	MC Kalman	Learned	Learned + rand.
Nominal	0.80	0.88	**0.89**
σ_p = 0,15 m	0.75	0.85	**0.87**
σ_p = 0,30 m	0.71	0.81	**0.84**
pmiss = 0,4	0.78	0.87	**0.88**
L_v = 6 frames (200 ms)	0.76	0.86	**0.87**
Real trajectories (ETH/UCY)	0.84	0.93	**0.94**

In a real deployment, the camera pipeline will be noisier than in our nominal set-up. Figure 6 shows that missed detections up to 40% and latencies up to 200 ms raise the interruption of DRL-Vision by at most 0.3 percentage points. Position error matters more: at σ_p = 0,30 m, comparable to the radius of a body, interruption rises from 5.8% at the nominal 6 cm to 7.7%, still 32% below STD-LTE. The physics-based forecast degrades less in relative terms because its uncertainty model widens with the measurement noise, whereas the network has to learn this effect. Table 4 shows that training with randomized perception quality recovers part of the loss at the prediction level (AUC 0.81 to 0.84 at 0.3 m). Under all the conditions we tested, the learned forecast remained the more accurate of the two.

7.5 Sensitivity to Handover Execution Time

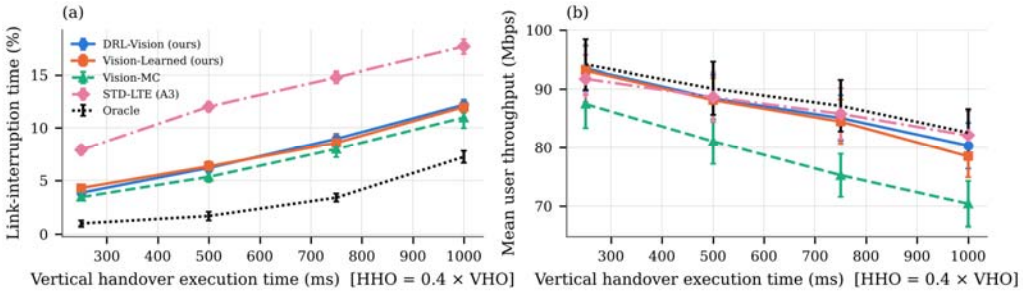


Figure 7. Effect of the handover execution time (nine pedestrians, 10 episodes per point; $T_{HHO} = 0{,}4T_{VHO}$). (a) Interruption time. (b) Throughput

The HO execution time in HLWNets depends on the protocol stack and is uncertain, so Figure 7 scales it from T_{VHO} = 250 ms to 1 s. Interruption grows for every scheme because a longer execution needs a longer, less accurate forecast and leaves less margin. DRL-Vision cuts interruption by 51% at 250 ms (7.9% to 3.8%) and by 31% at 1 s (17.7% to 12.2%), and keeps throughput within 2.2% of STD-LTE. Vision-MC loses up to 14% of throughput at long execution times because its false alarms become more expensive.

7.6 Zero-Shot Test With Real Pedestrian Trajectories

Table 5

Scenario B: Static Users, Real ETH/UCY Pedestrians
(30 Episodes, Mean 12.4 Pedestrians in the Room)

Scheme	Throughput (Mb/s)	Interruption (%)	HO/user/min
Reactive	104.7 ± 7.3	3.46 ± 0.74	10.6
STD-LTE (A3)	106.5 ± 7.3	3.57 ± 0.79	12.7
DRL (no vision)	109.4 ± 7.5	2.81 ± 0.66	8.9
Vision-Det	105.0 ± 7.5	1.74 ± 0.48	31.0
Vision-MC	103.4 ± 7.5	1.30 ± 0.28	35.5
Vision-Learned	107.6 ± 7.8	1.76 ± 0.41	16.3
DRL-Vision	107.8 ± 7.4	1.70 ± 0.36	14.2
Oracle	108.1 ± 7.7	0.60 ± 0.14	10.9

Scenario B tests whether models trained on synthetic motion transfer to real crowds and to a different usage pattern, with static users and pedestrians whose number in the room varies from 1 to 32 (median 5.7). Interruption is lower than in Scenario A because static users suffer no self-blockage changes. Table 5 shows that, without retraining, DRL-Vision reduces interruption by 53% relative to STD-LTE (p = 3 x 10^{-6}) with a +1,2% change in throughput (p = 0,75), and the learned forecast reaches an AUC of 0.94 (Table 4). Vision-MC again achieves the lowest interruption, but at the cost of 2.9% throughput (p = 0,015) and 2.5 times as many handovers as DRL-Vision.

7.7 Complexity

On one core of a general-purpose CPU, Kalman tracking of all persons takes 0.13 ms per frame. Building the link features and evaluating the predictor (94 878 parameters) for the 24 user-AP links takes 0.8 ms per decision step, and the DQN (19 075 parameters) a negligible further amount. The Monte-Carlo forecast with 32 samples takes 6.6 ms per decision and can be omitted, since it does not improve the learned predictor. The added computation is therefore small compared with the 67 ms decision period; the person detector dominates the edge workload and is accounted for through the latency L_v.

8 Limitations and future work

Several limitations of this study should be noted.

- *Emulated perception.* The detector and tracker are modelled by their output errors rather than run on video. The robustness sweeps bound the effect of these errors, but an end-to-end evaluation with recorded or rendered video, real detector failures and identity switches is needed.
- *Device association.* We assume that the network knows which tracked person carries which device. Practical options include LiFi uplink positioning and matching camera motion with inertial sensor data.
- *Channel model.* We consider the LoS channel with orthogonal LiFi bands. Diffuse reflections would soften some blockages, and inter-cell interference would couple the association decisions more strongly.
- *Controller.* The DRL agent matches but does not exceed a well-tuned rule using the same forecast. Its value here is that it learns the operating point from the reward.

Learning a joint policy for all users and extending the action space to per-user time allocation are natural next steps.

- *Privacy.* Processing at the edge and exporting only floor positions and pose angles reduces the exposure of video. A deployment would still require a privacy assessment.

9 Conclusion

We have shown that the cameras already installed in many smart buildings can be used to make LiFi handover proactive. The central element is a link-level forecast that respects LiFi geometry, including the short blocking zone, self-blockage and device tilt, and that is learned from vision and radio features. This forecast raised the 200-ms outage-prediction AUC from 0.79 for a Kalman/Monte-Carlo model to 0.89. Used by a DRL handover agent with proportional-fair resource allocation, it reduced link-interruption time by 42-47% across crowd densities, with no loss of throughput and with fewer false alarms than physics-based forecasting. The gain held, without retraining, with real pedestrian trajectories (53% lower interruption). The remaining gap to the oracle suggests that improving the prediction, rather than the controller, is the most promising direction for future work.

REFERENCES

[1] H. Haas, L. Yin, Y. Wang, and C. Chen, “What is LiFi?” *J. Lightw. Technol.*, vol. 34, no. 6, pp. 1533–1544, 2016, doi: 10.1109/JLT.2015.2510021.

[2] “IEEE standard for information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements–Part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications–Amendment 6: Light communications,” *IEEE Std 802.11bb-2023*, 2023.

[3] X. Wu, M. D. Soltani, L. Zhou, M. Safari, and H. Haas, “Hybrid LiFi and WiFi networks: A survey,” *IEEE Commun. Surveys Tuts.*, vol. 23, no. 2, pp. 1398–1420, 2021, doi: 10.1109/COMST.2021.3058296.

[4] M. D. Soltani, A. A. Purwita, Z. Zeng, H. Haas, and M. Safari, “Modeling the random orientation of mobile devices: Measurement, analysis and LiFi use case,” *IEEE Trans. Commun.*, vol. 67, no. 3, pp. 2157–2172, Mar. 2019.

[5] Y. Wang and H. Haas, “Dynamic load balancing with handover in hybrid Li-Fi and Wi-Fi networks,” *J. Lightw. Technol.*, vol. 33, no. 22, pp. 4671–4682, 2015, doi: 10.1109/JLT.2015.2480969.

[6] D. Frómeta Fonseca, B. Genovés Guzmán, G. L. Martena, R. Bian, H. Haas, and D. Giustiniano, “Prediction-model-assisted reinforcement learning algorithm for handover decision-making in hybrid LiFi and WiFi networks,” *J. Opt. Commun. Netw.*, vol. 16, no. 2, pp. 159–170, 2024, doi: 10.1364/JOCN.495234.

[7] S. Karimkhan and H. Shokrzadeh, “Reinforcement learning-based dueling DQN for access point selection and handover management in mobile hybrid LiFi/WiFi networks,” *Wireless Pers. Commun.*, 2026, doi: 10.1007/s11277-026-12215-w.

[8] P. Hou and N. Cen, “Deep reinforcement learning-based user association in hybrid LiFi/WiFi indoor networks,” *arXiv preprint* arXiv:2503.01803, 2025.

[9] G. Charan, M. Alrabeiah, and A. Alkhateeb, “Vision-aided 6G wireless communications: Blockage prediction and proactive handoff,” *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10193–10208, 2021, doi: 10.1109/TVT.2021.3104219.

[10] A. Alkhateeb, G. Charan, T. Osman, A. Hredzak, J. Morais, U. Demirhan, and N. Srinivas, “DeepSense 6G: A large-scale real-world multi-modal sensing and communication dataset,” *IEEE Commun. Mag.*, vol. 61, no. 9, pp. 122-128, 2023, doi: 10.1109/MCOM.006.2200730.

[11] S. Karunasena, E. Khordad, T. Drummond, and R. Senanayake, “Deep learning-based computer vision for beam selection and proactive blockage prediction,” *arXiv preprint* arXiv:2605.04514, 2026.

[12] Y. Ahn, J. Kim, S. Kim, and B. Shim, “Computer vision-aided proactive mobility management for 6G terahertz communications,” in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, 2023, pp. 1513-1518, doi: 10.1109/GLOBECOM54140.2023.10437835.

[13] S. Naser, O. Alhussein, and S. Muhaidat, “Towards multi-modal and cross-modal integration in LiFi-based sensing,” *Sci. Rep.*, vol. 16, Art. no. 6038, 2026, doi: 10.1038/s41598-026-36891-7.

[14] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, “You’ll never walk alone: Modeling social behavior for multi-target tracking,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2009, pp. 261-268.

[15] A. Lerner, Y. Chrysanthou, and D. Lischinski, “Crowds by example,” *Comput. Graph. Forum*, vol. 26, no. 3, pp. 655-664, 2007.

[16] S. Chatterjee, D. Sabui, and B. Roy, “Mobility aware blockage management of a multi-user multi-cell hybrid Li-Fi Wi-Fi system with freeform diversity receiver,” *Opt. Commun.*, vol. 560, Art. no. 130487, 2024, doi: 10.1016/j.optcom.2024.130487.

[17] J. M. Kahn and J. R. Barry, “Wireless infrared communications,” *Proc. IEEE*, vol. 85, no. 2, pp. 265-298, 1997.

[18] S. Dimitrov and H. Haas, “Information rate of OFDM-based optical wireless communication systems with nonlinear distortion,” *J. Lightw. Technol.*, vol. 31, no. 6, pp. 918–929, 2013.

[19] V. Erceg et al., “TGn channel models,” *IEEE 802.11 document 03/940r4*, 2004.

[20] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 779-788.

[21] Y. Zhang et al., “ByteTrack: Multi-object tracking by associating every detection box,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2022, pp. 1-21.

[22] V. Mnih et al., “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, pp. 529-533, 2015.

[23] H. van Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double Q-learning,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 30, 2016, pp. 2094-2100.

[24] T. Hester et al., “Deep Q-learning from demonstrations,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 32, 2018, pp. 3223-3230.

[25] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, “Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 683-700.